

1 **Title:** Minian an Open-source Miniscope Analysis Pipeline

2 **Authors:** Zhe Dong¹, William Mau¹, Yu Feng¹, Zachary T. Pennington¹, Lingxuan Chen¹, Yosif
3 Zaki¹, Kanaka Rajan¹, Tristan Shuman¹, Daniel Aharoni^{*2}, Denise J. Cai^{*1}

4 ^{*}Corresponding Authors

5 ¹Nash Family Department of Neuroscience, Icahn School of Medicine at Mount Sinai

6 ²Department of Neurology, David Geffen School of Medicine, University of California, Los
7 Angeles

8

9 Abstract

10 Miniature microscopes have gained considerable traction for *in vivo* calcium imaging in freely
11 behaving animals. However, extracting calcium signals from raw videos is a computationally
12 complex problem and remains a bottleneck for many researchers utilizing single-photon *in*
13 *vivo* calcium imaging. Despite the existence of many powerful analysis packages designed to
14 detect and extract calcium dynamics, most have either key parameters that are hard-coded or
15 insufficient step-by-step guidance and validations to help the users choose the best parameters.
16 This makes it difficult to know whether the output is reliable and meets the assumptions
17 necessary for proper analysis. Moreover, large memory demand is often a constraint for setting
18 up these pipelines since it limits the choice of hardware to specialized computers. Given these
19 difficulties, there is a need for a low memory demand, user-friendly tool offering interactive
20 visualizations of how altering parameters at each step of the analysis affects data output. Our
21 open-source analysis pipeline, Minian (Miniscope Analysis), facilitates the transparency and
22 accessibility of single-photon calcium imaging analysis, permitting users with little computational
23 experience to extract the location of cells and their corresponding calcium traces and
24 deconvolved neural activities. Minian contains interactive visualization tools for every step of the
25 analysis, as well as detailed documentation and tips on parameter exploration. Furthermore,
26 Minian has relatively small memory demands and can be run on a laptop, making it available to
27 labs that do not have access to specialized computational hardware. Minian has been validated
28 to reliably and robustly extract calcium events across different brain regions and from different
29 cell types. In practice, Minian provides an open-source calcium imaging analysis pipeline with
30 user-friendly interactive visualizations to explore parameters and validate results.

31 Introduction

32 *Overview of related works*

33 Open-source projects—hardware, software, training curricula—have changed science and
34 enabled significant advances across multiple disciplines. Neuroscience, in particular, has
35 benefitted tremendously from the open-source movement. Numerous open-source projects
36 have emerged [1,2], including various types of behavioral apparatus facilitating the design of
37 novel experiments [3,4,5,6,7], computational tools enabling the analysis of large scale datasets
38 [8,9,10,11,12,13,14,15,16,17,18,19,20,21], and recording devices allowing access to large
39 populations of neurons in the brain [22,23,24,25,26,27,28,29,30,31]. Miniature microscopy has
40 been an area of particular importance for the open-source movement in neuroscience. To
41 increase the usability, accessibility, and transparency of this remarkable technology originally
42 developed by Schnitzer and colleagues [32,33], a number of labs innovated on top of the
43 original versions with open-source versions [26,27,28,29,30,31]. The UCLA Miniscope project, a
44 user-friendly miniature head-mounted microscope for *in vivo* calcium imaging in freely behaving
45 animals, is one such project that has been accessible to a large number of users [22,34,35,36].

46 With the increasing popularity of miniature microscopes, there is a growing need for analysis
47 pipelines that can reliably extract neuronal activities from recording data. To address this need,
48 numerous algorithms have been developed and made available to the neuroscience community.
49 The principal component analysis or independent component analysis (PCA-ICA)-based
50 approach [13], and region-of-interest (ROI)-based approach [34] were among the earliest

51 algorithms that reliably detected the locations of neurons and extract their overall activities
52 across pixels. However, one of the limitations of these approaches is that activities from cells
53 that are spatially overlapping cannot be demixed. A subsequent constrained non-negative
54 matrix factorization (CNMF) approach was shown to reliably extract neuronal activity from both
55 two-photon and single-photon calcium imaging data [37], and demix the activities of overlapping
56 cells. The CNMF algorithm models the video as a product of a ‘spatial’ matrix containing
57 detected neuronal footprints (locations of cells) and a ‘temporal’ matrix containing the temporal
58 calcium traces of each detected cell. This approach is particularly effective at addressing
59 crosstalk between neurons, which is of particular concern in single-photon imaging, where the
60 fluorescence from overlapping or nearby cells contaminates each other. Moreover, by
61 deconvolving calcium traces, the CNMF algorithm enables a closer exploration of the underlying
62 activity of interest, action potentials [19,38]. Originally developed for two-photon data, the CNMF
63 algorithm did not include an explicit model of the out-of-focus fluorescence which is often
64 present in single-photon miniature microscope recordings. This issue was addressed via the
65 CNMF-E algorithm [11], where a ring-model is used as a background term to account for out-of-
66 focus fluorescence. Later, an open-source python pipeline for calcium imaging analysis,
67 CalmAn, was published, which included both the CNMF and CNMF-E algorithms, as well as
68 many other functionalities [16]. The latest development in analysis pipelines for *in vivo* miniature
69 microscope data is MIN1PIPE [12], where a morphological operation is used to
70 remove background fluorescence during pre-processing of the data, and a seed-based
71 approach is used for initialization of the CNMF algorithm. Other approaches have also been
72 used to extract signals from calcium imaging data including an online approach [20], ℓ_0 -
73 penalization approach to infer spikes [14,21], robust modeling of noise [39], and source
74 detection using neural networks [15].

75 The open sharing of the algorithms necessary for the computation of neural activity has been
76 exceptionally important for the field. However, implementation of these tools can be complex as
77 many algorithms have numerous free parameters (those that must be set by the user) that can
78 influence the outcomes, without clear guidance on how these parameters should be set or to
79 what extent they affect results. Moreover, there is a lack of ground-truth data for *in*
80 *vivo* miniature microscope imaging, making it hard to validate algorithms and/or parameters.
81 Together, these obstacles make it challenging for neuroscience labs to adopt the analysis
82 pipelines, since it is difficult for researchers to adjust parameters to fit their data, or to trust the
83 output of the pipeline for downstream analysis. Thus, the next challenge in open-source
84 analysis pipelines for calcium imaging is to make the analysis tools more user-friendly and
85 underlying algorithms more accessible to neuroscience researchers so that they can more
86 easily understand the pipeline and interpret the results.

87 *Contributions of Minian*

88 To increase the accessibility of the mathematical algorithms, transparency into how altering
89 parameters alters the data output, and usability for researchers with limited computational
90 resources and experience, we developed Minian, an open-source analysis pipeline for single-
91 photon calcium imaging data inspired by previously published algorithms. We based Minian on
92 the CNMF algorithm [16,37], but also leverage methods from other pipelines, including those
93 originally published by Cai et al. [34] and MIN1PIPE [12]. To enhance compatibility with different
94 types of hardware, especially laptops or personal desktop computers, we implemented an

95 approach that supports parallel and out-of-core computation (i.e., computation on data that are
96 too large to fit a computer's memory). We then developed interactive visualizations for every
97 step in Minian and integrated these steps into annotated Jupyter Notebooks as an interface for
98 the pipeline. We have included detailed notes and discussions on how to adjust the parameters
99 from within the notebook and have included all free parameters in the code for additional
100 flexibility. The interactive visualizations will help users to intuitively understand and visually
101 inspect the effect of each parameter, which we hope will facilitate more usability, transparency,
102 and reliability in calcium imaging analysis.

103 Minian contributes to three key aspects of calcium image data analysis:

- 104 1. **Visualization.** For each step in the pipeline, Minian provides visualizations of inputs and
105 results. Thus, users can proceed step-by-step with an understanding of how the data are
106 transformed and processed. In addition, all visualizations are interactive and support
107 simultaneous visualization of the results obtained with different parameters. This feature
108 provides users with knowledge about the corresponding outcome for each parameter
109 value, and allow the users to choose the outcome that fits best with their expectation.
110 Hence, the visualizations also facilitate parameter exploration for each step, which is
111 especially valuable when analyzing data from heterogeneous origins that may vary by brain
112 region, cell type, species, and the extent of viral transfection.
- 113 2. **Memory demand.** One of the most significant barriers in adopting calcium imaging
114 pipelines is the memory demand of algorithms. The recorded imaging data usually take up
115 tens of gigabytes of space when converted to floating-point datatypes and often cannot fit
116 into the RAM of standard computers without spatially and/or temporally down-sampling.
117 CalmAn [16] addresses this issue by splitting the data into overlapping patches of pixels,
118 processing each patch independently, and merging the results together. This enables out-
119 of-core computation since at any given time only subsets of data are needed and loaded
120 into memory. In Minian, we extend this concept further by flexibly splitting the data either
121 spatially (split into patches of pixels) or temporally (split into chunks of frames). In this way,
122 we avoid the need to merge the results based on overlapping parts. The result is a pipeline
123 that supports out-of-core computation at each step, which gives nearly constant memory
124 demand with respect to input data size. Minian can process more than 20min of recording
125 (approximately 12.6 GB of raw data) with 8GB of memory, which makes Minian suitable to
126 be deployed on modern personal laptops.
- 127 3. **Accessibility.** Minian is an open-source Python package. In addition to the codebase,
128 Minian distributes several Jupyter Notebooks that integrate explanatory text with code and
129 interactive visualizations of results. For each step in the notebook, detailed instructions, as
130 well as intuition about the underlying mathematical formulation are provided, along with
131 code, which can be directly executed from within the notebook. Upon running a piece of
132 code within the notebook visualizations appear directly below. In this way, the notebooks
133 serve as a complement to traditional API documentations of each function. In addition,
134 users can easily rearrange and modify the pipeline notebook to suit their needs without
135 diving into the codebase and modifying the underlying functions. The notebooks distributed
136 by Minian can simultaneously function as a user guide, template, and production tool. We
137 believe the inclusion of these notebooks, in combination with Minian's other unique

138 features, can increase understanding of the underlying functioning of the algorithms and
139 greatly improve the accessibility of miniature microscopy analysis pipelines.

140 *Paper organization*

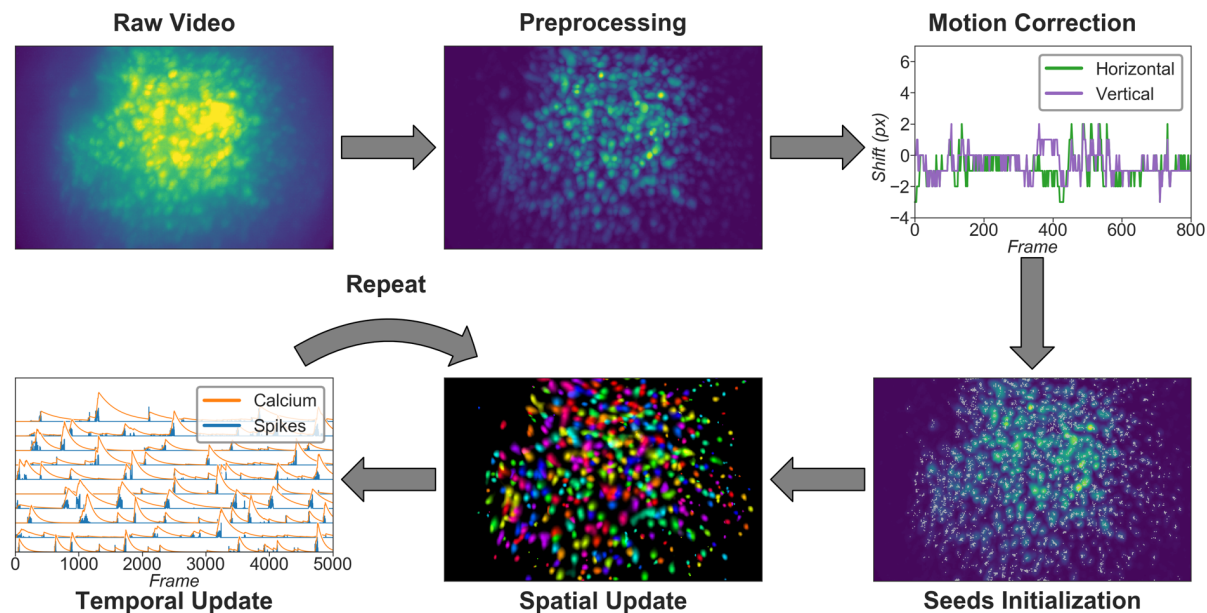
141 The paper is organized as follows: Since Minian's major contribution is usability and
142 accessibility, we first present the detailed steps in the analysis pipeline in Materials and
143 Methods section. Following a step-by-step description of the algorithms Minian adopted from
144 existing works, we present novel visualizations of the results, as well as how users can utilize
145 these visualizations. In the Results section, we benchmark Minian across two brain regions and
146 show that spatial footprints and the temporal activity of cells can be reliably extracted. We also
147 show that the cells extracted by Minian in hippocampal CA1 exhibit stable spatial firing
148 properties consistent with the existing literature.

149 **Materials and Methods**

150 Here, we present a detailed description of Minian. We begin with an overview of the Minian
151 pipeline. Then, we provide an explanation of each step, along with the visualizations. Lastly, we
152 provide information regarding hardware and dependencies.

153 *Overview of Minian*

154 Minian comprises five major stages, as shown in Figure 1. Raw videos are first passed into a
155 pre-processing stage. During pre-processing, the background caused by vignetting (in which the
156 central portion of the field of view is brighter) is corrected by subtracting a minimum projection of
157 the movie across time. Sensor noise, evident as granular specks, is then corrected with a
158 median filter. Finally, background fluorescence is corrected by the morphological process
159 introduced in MIN1PIPE [12]. The pre-processed video is then motion-corrected with a standard
160 template-matching algorithm based on cross-correlation between each frame and a reference
161 frame [40]. The motion-corrected and pre-processed video then serves as the input to
162 initialization and CNMF algorithms. The seed-based initialization procedure looks for local
163 maxima in max projections of different subsets of frames and then generates an over-complete
164 set of seeds, which are candidate pixels for detected neurons. Because this process is likely to
165 produce many false positives, seeds are then further refined based on various metrics, including
166 the amplitude of temporal fluctuations and the signal-to-noise ratio of temporal signals. The
167 seeds are transformed into an initial estimation of cells' spatial footprints based on the
168 correlation of neighboring pixels with each seed pixel, and the initial temporal traces are in turn
169 estimated based on the weighted temporal signal of spatial footprints. Finally, the processed
170 video, initial spatial matrix, and temporal matrix are fed into the CNMF algorithm. The CNMF
171 algorithm first refines the spatial footprints of the cells (spatial update). The algorithm then
172 denoises the temporal traces of each cell while simultaneously deconvolving the calcium trace
173 into estimated 'spikes' (temporal update). CNMF spatial and temporal updates are performed
174 iteratively and can be repeated until a satisfactory result is reached through visual inspection.
175 Typically, this takes two cycles of spatial, followed by temporal, updates. Minian also includes a
176 demo dataset which allows the user to run and test the pipeline comprised of the pre-made
177 Jupyter Notebook immediately after installation.



178

179 **Figure 1: Overview of the analysis pipeline.** The analysis is divided into five stages: Pre-
180 preprocessing, where sensor noise and background fluorescence from scattered light are removed;
181 Motion-correction, where rigid motion of the brain is corrected; Seeds-initialization, where the
182 initial spatial and temporal matrices for later steps are generated from a seed-based approach;
183 Spatial update, where the spatial footprints of cells are further refined; Temporal update, where
184 the temporal signals of cells are further refined. The last two steps of the pipeline are iterative
185 and can be repeated multiple times until a satisfactory result is reached.

186 **Setting up**

187 The first section in the pipeline includes house-keeping scripts to import packages and
188 functions, defining parameters, and setting up parallel computation and visualization. Most
189 notably, the distributed cluster that carries out all computations in Minian are set up in this
190 section. By default, the cluster runs locally with multi-core CPUs, however it can be easily
191 scaled up to run on distributed computers. The computation in Minian is optimized such that in
192 most cases the memory demand for each process/core can be as low as 2GB. However, in
193 some cases depending on the hardware, the state of operating system and data locality, Minian
194 might need more than 2GB per process to run. If a memory error (KilledWorker) is encountered,
195 it is common for users to increase the memory limit of the distributed cluster to get around the
196 error. Regardless of the exact memory limit per process, the total memory usage of Minian
197 roughly scales linearly with the number of parallel processes. The number of parallel processes
198 and memory usage of Minian are completely limited and managed by the cluster configuration
199 allowing users to easily change them to suit their needs.

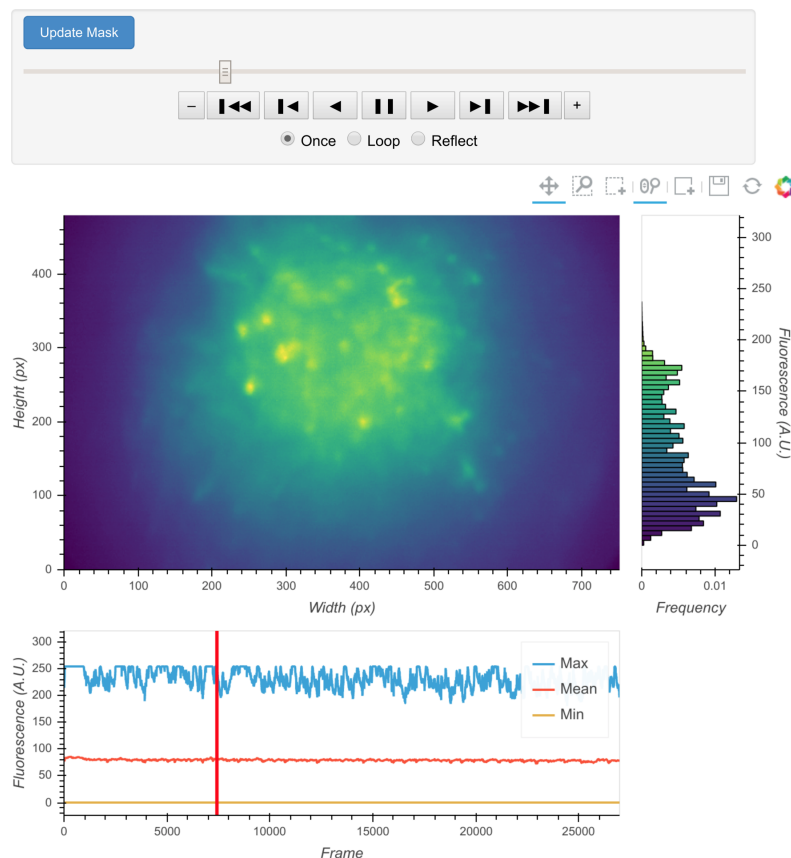
200 *Pre-processing*

201 Loading data and down-sampling

202 Currently Minian supports .avi movies, the default output from the UCLA Miniscopes, and .tif
203 stacks, the default output from Inscopix miniscopes. This functionality can be easily extended to
204 support more formats if desired. Users are required to organize their data so that each recording
205 session is contained in a single folder. Because Minian can extract relevant metadata from
206 folder nomenclature (e.g., animal name, group, date), we suggest organizing the video folders
207 based upon animal and other experiment-related groupings to facilitate the incorporation of
208 metadata into Minian output files.

209 Minian supports down-sampling on any of the three video dimensions (height, width, and
210 frames). Two down-sampling strategies are currently implemented: either sub-setting data on a
211 regular interval or calculating a mean for each interval. At this stage, users are required to
212 specify (1) the path to their data, (2) a pattern of file names to match all the videos to be
213 processed (e.g., all files containing 'msCam', a typical pattern resulting from Miniscope
214 recordings), (3) a Python dictionary specifying whether and how metadata should be pulled from
215 folder names, (4) another Python dictionary specifying whether and on which dimension down-
216 sampling should be carried out, and (5) the down-sampling strategy, if desired.

217 Once specified, the data can be immediately visualized through an interactive viewer, as shown
218 in Figure 2. Along with a player to visualize every frame in the video, the viewer also plots
219 summary values such as mean, maximum, or minimum fluorescence values across time. This
220 helps users to check their input data and potentially exclude any artifacts caused by technical
221 faults during experiments (e.g., dropped frames). Users can further subset data to exclude
222 specified frames, if necessary. Finally, restricting the analysis to a certain sub-region of the field
223 of view during specific steps could be beneficial. For example, if the video contains anchoring
224 artifacts resulting from dirt on the lenses, it is often better to avoid such regions during motion
225 correction. To facilitate this, the viewer provides a feature where users can draw an arbitrary
226 box within the field of view and have it recorded as a mask. This mask can be passed into later
227 motion correction steps to avoid the biases resulting from the artifacts.



228

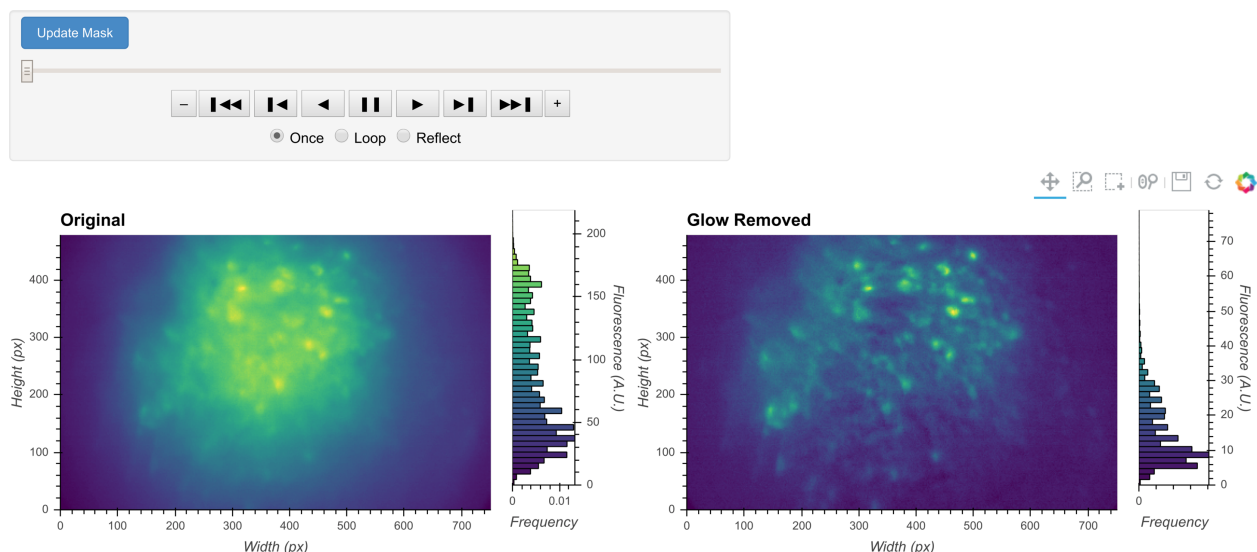
229 **Figure 2: Interactive visualization of raw input video.** One frame is shown in the central panel
230 of the visualization which can be interactively updated with the player toolbar on the top. A
231 histogram of fluorescence intensity of the current frame is shown on the right and will update in
232 response to zooming in on the central frame. A line plot of summary values across time is
233 shown on the bottom. Here the maximum, mean, and minimum fluorescence values are plotted.
234 These summaries are useful in checking whether there are unexpected artifacts or gaps in the
235 recording. Finally, the user can draw an arbitrary box in the central frame, and the position of
236 this boxed region can be recorded and used as a mask during later steps. For example, during
237 motion correction a sub-region of the data containing a stable landmark might provide better
238 information on the motion.

239 Vignetting correction

240 Single-photon miniature microscope data often suffer from a vignetting effect in which the
241 central portion of the field of view appears brighter than the periphery. Vignetting is deleterious
242 to subsequent processing steps and should be removed. We find that the effect can be easily
243 extracted by taking the minimum fluorescence value across time for each pixel and subtracting
244 this value from each frame, pixel-wise. One of the additional benefits of subtracting the
245 minimum is that it preserves the raw video's linear scale.

246 The result of this step can be visualized with the same video viewer used in the previous step.
247 In addition to visualizing a single video, the viewer can also show multiple videos side-by-side
248 (e.g., the original video and the processed video), as shown in Figure 3. The

249 operation/visualization is carried out 'on-the-fly' upon request for each frame, and users do not
250 have to wait for the operation to finish on the whole video to view the results.



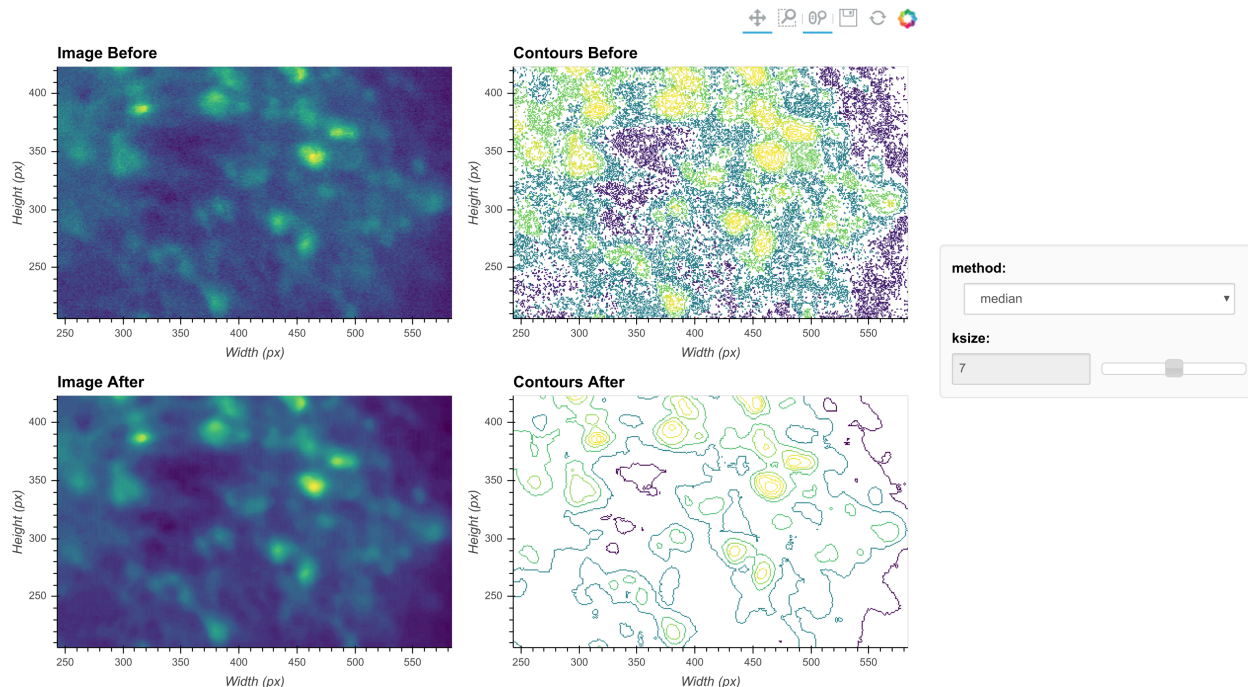
251
252 **Figure 3: General visualization of pre-processing.** The same visualization of input video can
253 be used to visualize the whole video before and after specific pre-processing steps side-by-side.
254 The effect of vignetting correction is visualized here. The image and accompanying histogram
255 on the left side show the original data; the data after vignetting correction are shown on the right
256 side. Any frame of the data can be selected with the player toolbar and histograms are
257 responsive to all updates in the image.

258 Denoising

259 Next, we correct for salt-and-pepper noise on each frame, which usually results from electronic
260 pixel noise. By default, we pass each frame through a median filter, which is generally
261 considered particularly effective at eliminating this type of noise, though other smoothing filters
262 like Gaussian filters and anisotropic filters can also be implemented. The critical parameter here
263 is the window size of the median filter. A window size that is too small will make the filter
264 ineffective at correcting outliers, while a window size that is too large will remove finer gradient
265 and edges that are much smaller than the window size, and can result in a failure to distinguish
266 between adjacent cells.

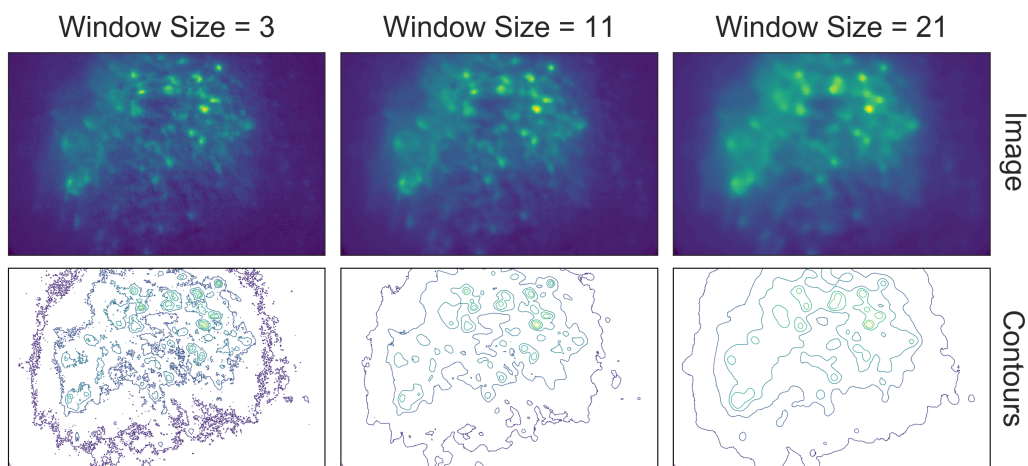
267 The effect of the window size can be checked with an interactive visualization tool used across
268 the pre-processing stage, as shown in Figure 4. Additionally, here we show an example of the
269 effect of window size on the resulting data in Figure 5. Users should see significantly reduced
270 amount of salt-and-pepper noise in the images, which should be made more obvious by the
271 contour plots. At the same time, users should keep the window size below the extent where
272 over-smoothing occurs. As a heuristic, the average cell radius in pixel units works well, since a
273 window of the same size as an average cell is unlikely to blend different cells together, while still
274 being able to adequately smooth the image.

method: median, ksize: 7



275

276 **Figure 4: Visualization of denoising.** Here, a single frame from the data is passed through the
277 background removal and both the image and a contour plot are shown for the frame before and
278 after the process. The contour plots show the iso-contour of 5 intensity levels spaced linearly
279 across the full intensity range of the corresponding image. The plots are interactive and
280 responsive to the slider of the window size on the right, thus the effect of different window sizes
281 for denoising can be visualized.



282

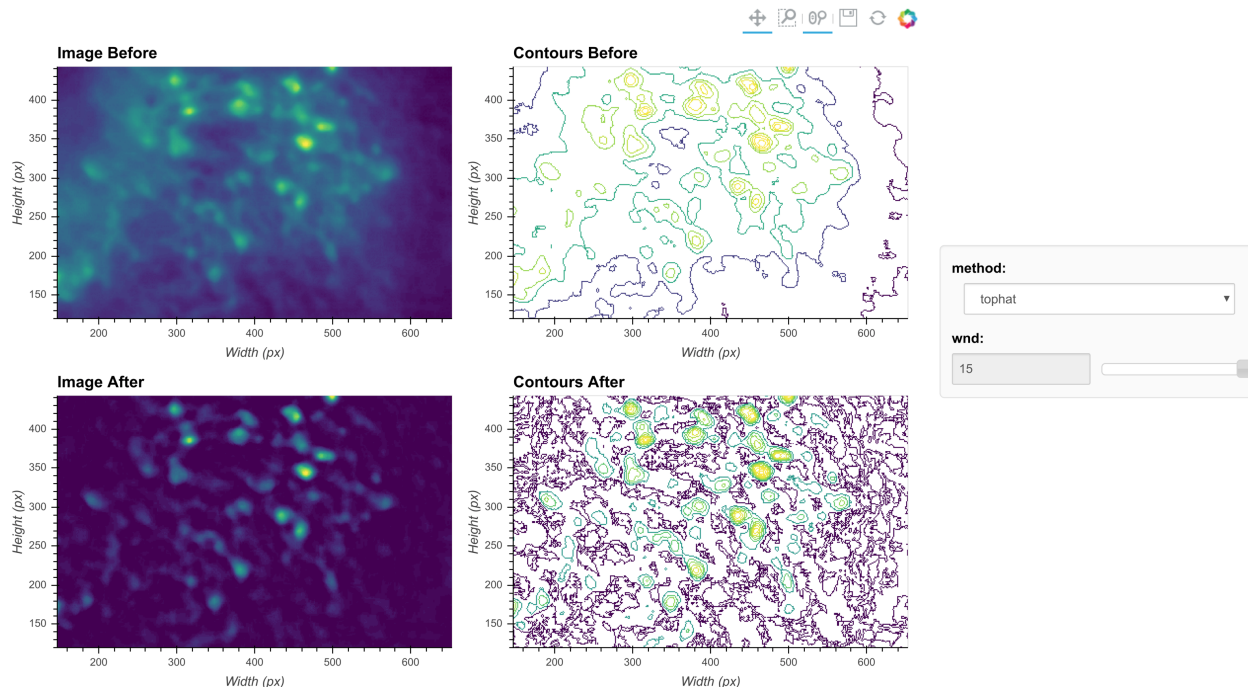
283 **Figure 5: Effect of window size on denoising.** One example frame is chosen from the data,
284 and the resulting images (top row) and contour plots (bottom row) are shown to demonstrate the
285 effect of window size on denoising. Here, a window size of 11 (middle column) is appropriate
286 while both smaller and larger window sizes result in artifacts.

287 Morphological background removal


288 Next, we remove any remaining background presumably caused by out-of-focus and tissue
289 fluorescence. To accomplish this we estimate the background using a morphological opening
290 process first introduced for calcium imaging analysis in MIN1PIPE [12], which acts as a size
291 filter that removes cell bodies. The morphological opening is composed of two stages: erosion
292 followed by dilation. In morphological erosion the image is passed through a filter where each
293 pixel will be substituted by the minimum value within the filter window. The effect of this process
294 is that any bright 'feature' that is smaller than the filter window will be 'eroded' away. Then the
295 dilation process accomplishes the reverse by substituting each pixel with the maximum value in
296 the window, which 'dilates' small bright features to the extent of the filter window size. The
297 combined effect of these two stages is that any bright 'feature' that is smaller than the filter
298 window is removed from the image. If we choose the window size to match the expected cell
299 diameter, performing a morphological opening will likely remove cells and provide a good
300 estimation of background. Hence, each frame is passed through the morphological opening
301 operation and the resulting image is subtracted from the original frame.

302 Although the window size parameter for the morphological opening can be pre-determined by
303 the expected cell diameter, it is helpful to visually inspect the effect of morphological
304 background removal. The effect of different window sizes can be visualized with the same tool
305 used in denoising, as shown in Figure 6. Additionally, here we show an example of the effect of
306 window size on the resulting data in Figure 7. In this case, a window size of 20 pixels is
307 considered appropriate because the resulting cells are appropriately sized and sharply defined.
308 In contrast, a smaller window results in limiting both the size and intensity of the cells. On the
309 other hand, residual out-of-focus fluorescence becomes visible when the window size is set too
310 large.

method: tophat, wnd: 15



311
312 **Figure 6: Visualization of background removal.** Here, a single frame from the data is passed
313 through background removal and both the image and a contour plot are shown for the frame
314 before and after the process. The plots are interactive and responsive to the slider of the
315 window size on the right, thus the effect of different window sizes for background removal can
316 be visualized.



317
318 **Figure 7: Effect of window size on background removal.** One example frame is chosen from
319 the data, and the resulting images (top row) and contour plots (bottom row) are shown to
320 demonstrate the effect of window size on background removal. The contour plots show the iso-
321 contour of 5 intensity levels spaced linearly across the full intensity range of the corresponding
322 image. Here a window size of 20 pixels (middle column) is appropriate while both smaller and
323 larger window sizes produce unsatisfactory results: a window size too small (left column)

324 *artificially limits the size of cells, and a window size too large (right column) does not remove the*
325 *background effectively.*

326 *Motion correction*

327 Estimate and apply translational shifts

328 We use a standard template-matching algorithm based on cross-correlation to estimate and
329 correct for translational shifts [40]. In practice, we found that this approach is sufficient to correct
330 for motion artifacts that could have a significant impact on the final outcome. Briefly, for a range
331 of possible shifts, a cross-correlation between each frame and a template frame is calculated.
332 The shift producing the largest cross-correlation is estimated to reflect the degree of movement
333 from the template and is corrected by applying a shift to the frame in that direction. We apply
334 this operation to the whole movie in a divide-and-conquer manner. We split the movie into
335 chunks of frames, within which we register both the first and last frame to the middle frame. We
336 then take the max projections of the three frames that have been registered in each chunk and
337 group every 3 chunks together and register them using the max projections as templates. After
338 the registration, the 3 chunks that have been registered are treated as a new single chunk and
339 we again take the max projection to use as a template for further registration. In this way, the
340 number of frames registered in each chunk keeps increasing in powers of three (3, 9, 27, 81
341 etc.), and we repeat this process recursively until all the frames are covered in a single chunk
342 and the whole movie is registered. Since the motion correction is usually carried out after
343 background removal, we essentially use cellular activity as landmarks for registration.
344 Sometimes this can be problematic when cellular activity is very sparse and different across two
345 chunks (for example, when only two different cells fired in two chunks), leading to false
346 estimation of shifts. To overcome this problem, every time shift is estimated using a max
347 projection from two chunks, we also estimate a shift with the two consecutive frames bordering
348 the chunks (that is, the last frame from the earlier chunk and the first frame from the latter
349 chunk). In most cases the shifts estimated with these two sets of templates should be close, in
350 which case we use the shifts estimated with the max projection as the final output. However,
351 when the two estimated shifts differ too much from each other, we use the shifts estimated with
352 consecutive frames as the final output. The reason we still favor using max projections in most
353 cases is that registering with consecutive frames can lead to very fast accumulation of error and
354 a slow drifting artifact in the estimated shifts. In practice, we find that such a process can
355 account for almost all motion in the brain, so currently we only implemented estimation of
356 translational shifts. If the user would like to take advantage of anatomical landmarks (such as
357 blood vessels) within the field of view and would like to implement motion correction before all
358 background subtraction steps have been performed, the pipeline can be easily modified to do
359 so. After the estimation of shifts, the shift in each direction is plotted across time and
360 visualization of the data before and after motion correction is displayed in Minian (see Figure 1,
361 top right).

362 *Seed initialization*

363 Generation of an over-complete set of seeds

364 The CNMF algorithm is a powerful approach to extract cells' spatial structure and corresponding
365 temporal activity. However, the algorithm requires an initial estimate of cell locations/activity,
366 which it then refines. We use a seed-based approach introduced in MIN1PIPE [12] to initialize
367 spatial and temporal matrices for CNMF. The first step is to generate an over-complete set of
368 seeds, representing the potential centroids of cells. We iteratively select a subset of frames,
369 compute a maximum projection for these frames, and find the local maxima on the projections.
370 This workflow is repeated multiple times and we take the union of all local maxima across
371 repetitions to obtain an over-complete set of seeds. In this way, we avoid missing cells that only
372 fire in short periods of time that might be masked by taking a maximum projection across the
373 whole video.

374 During seed initialization, the first critical parameter is the spatial window for defining local
375 maxima. Intuitively, this should be the expected diameter of cells. The other critical parameter is
376 an intensity threshold for a local maximum to be considered a seed. Since the spatial window
377 for local maxima is small relative to the field of view, a significant number of local maxima are
378 usually false positives and do not actually reflect the location of cells. Thresholding the
379 fluorescence intensity provides a simple way to filter out false local maxima, and usually a very
380 low value is enough to produce satisfactory results. We have found a value of 3 usually works
381 well (recall that the range of fluorescence intensity is usually 0-255 for unsigned 8-bit data). An
382 alternative strategy to thresholding the intensity is to model the distribution of fluorescence
383 fluctuations and keep the seeds with relatively higher fluctuations. This process is described
384 in Seeds refinement with a Gaussian-Mixture-Model, and is accessible if the user prefers explicit
385 modeling over thresholding.

386 Finally, the temporal sampling of frames for the maximum projections also impacts the result.
387 We provide two implementations here: either taking a rolling window of frames across time, or
388 randomly sampling frames for a user-defined number of iterations. For the rolling window
389 approach, users can specify a temporal window size (the number of successive frames for each
390 subset) and a step size (the interval between the start of subsets). For the random approach,
391 users can specify the number of frames in each subset and the total number of repetitions. We
392 use the rolling window approach as the default.

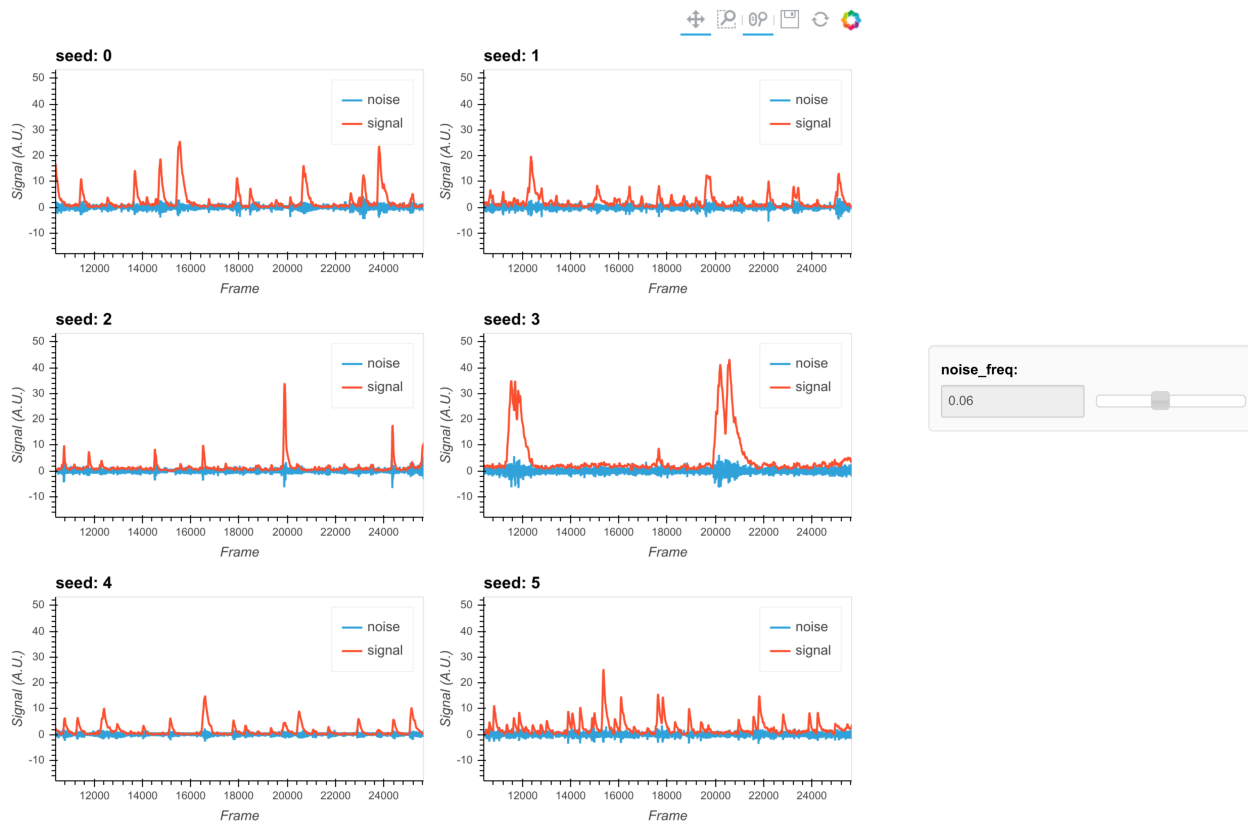
393 The resulting seeds are visualized on top of a maximum projection image (plot not shown).
394 Although the spatial window size of local maxima can be pre-determined, the parameters for
395 either the rolling window or random sampling of frames are hard to estimate intuitively. We
396 provide default parameters that generally provide robust results. However, the user is also free
397 to vary these parameters to obtain reasonable seeds. As long as the resulting seeds are not too
398 dense (populating almost every pixel) or too sparse (missing cells that are visible in the max
399 projection), subsequent steps can be performed efficiently and are fairly tolerable to the specific
400 ways the seeds are initialized.

401 Refinement with peak-to-noise ratio

402 Next, we refine the seeds by looking at what we call the peak-to-noise ratio of the temporal
403 traces and discard seeds with low peak-to-noise ratios. To compute this ratio, we first separate
404 the noise from the presumed real signal. Calcium dynamics are mainly composed of low
405 frequency fluctuations (from the slow kinetics of the calcium fluctuations) while noise is
406 composed of higher frequency fluctuations. Thus, to separate the noise from the calcium
407 dynamics we pass the fluorescence time trace of each seed through a low-pass and a high-
408 pass filter to obtain the 'signal' and 'noise' of each seed. We then compute the difference
409 between the maximum and minimum values (or peak-to-peak values) for both 'signal' and
410 'noise', and the ratio between the two difference values defines the peak-to-noise ratio. Finally,
411 we filter out seeds whose peak-to-noise value falls below a user-defined threshold.


412 The first critical parameter here is the cut-off frequency that separates 'signal' from 'noise'. This
413 parameter is also important for subsequent steps when implementing the CNMF algorithm. We
414 provide a visualization tool, shown in Figure 8, to help users determine cut-off frequency. In the
415 visualization, 6 seeds are randomly selected, and their corresponding 'signal' and 'noise' traces
416 are plotted. The user is then able to use a dynamic slider on the right side of the plots to adjust
417 the cut-off frequency and view the results. The goal is to select a frequency that best separates
418 signal from noise. A cut-off frequency that is too low will leave true calcium dynamics absorbed
419 in 'noise' (left panel in Figure 9), while a frequency that is too high will let 'noise' bleed into
420 'signal' (right panel in Figure 9). A suitable frequency is therefore the one where the 'signal'
421 captures all of the characteristics of the calcium indicator dynamics (i.e., large, fast rise, and
422 slow decay), while the 'noise' trace remains relatively uniform across time (middle panel in
423 Figure 9). The interactive plots make this easy to visualize. We also provide an example in
424 Figure 9 to show how cut-off frequency influences the separation of 'signal' from 'noise'. The
425 second parameter is the threshold of peak-to-noise ratio value. In practice, we have found a
426 threshold of 1 works well in most cases. An additional advantage of using 1 is that it reflects the
427 intuitive interpretation that fluctuations in a real 'signal' should be larger than fluctuations in
428 'noise'.

noise_freq: 0.06



429

430 **Figure 8: Visualization of noise frequency cut-off.** The cut-off frequency for noise is one of
431 the critical parameters in the pipeline that affects both the seed initialization process and
432 CNMF's temporal update steps. Here we help the user determine that parameter by plotting
433 temporal traces from six example seeds. In each plot the raw signal is passed through a high-
434 pass and low-pass filter at the chosen frequency, and the resulting signals are plotted
435 separately as "noise" and "signal". The plots are responsive to the chosen frequency controlled
436 by the slider on the right. In this way, the user can visually inspect whether the chosen
437 frequency can effectively filter out high frequency noise without deforming the calcium signal.



438

439 **Figure 9: Example of filtered traces with different frequency cut-offs.** Here the temporal
440 dynamics of three example seeds are chosen, and the low-pass and high-pass filtered traces
441 with different frequency cut-offs are shown. The low-pass filtered trace corresponds to 'signal',
442 while the high-pass filtered trace corresponds to 'noise'. Here a 1 Hz cut-off frequency is
443 considered appropriate, since calcium dynamics and random noise are cleanly separated. A
444 cut-off frequency smaller than 1 Hz left the calcium dynamics in the 'noise' trace, while a cut-off
445 frequency larger than 1 Hz let random noise bleed into the 'signal' trace (i.e., high frequency
446 fluctuations are presented in periods where the cells seem to be inactive).

447 Refinement with Kolmogorov-Smirnov tests

448 Finally, we refine the seeds with a Kolmogorov-Smirnov test. The Kolmogorov-Smirnov test
449 assesses the equality of two distributions and can be used to check whether the fluctuation of
450 values for each seed is non-normally distributed. We expect the noisy fluorescence values when
451 a cell is not firing to form a gaussian distribution with small mean value, and the fluorescence
452 values when a cell is firing should have a much higher mean value and frequency than expected
453 by the null gaussian distribution. Therefore, seeds corresponding to cells should be non-
454 normally distributed. We use a default significance threshold of 0.05. In some cases, this might
455 be too conservative or too liberal. Users can tweak this threshold or skip this step altogether
456 depending on the resulting seeds.

457 Merge seeds

458 There will usually be multiple seeds for a single cell and it is best to merge them whenever
459 possible. We implement two criteria for merging seeds: first, the distance between the seeds
460 must be below a given threshold, and second, the correlation coefficient of the temporal traces
461 between seeds must be higher than a given threshold. To avoid bias in the correlation due to
462 noise, we implement a smoothing operation on the traces before calculating the correlation. The
463 critical parameters are the distance threshold, the correlation threshold, and the cut-off

464 frequency for the smoothing operation. While the distance threshold is arbitrary and should be
465 explored, often the average radius of cells provides a good starting point. The cut-off frequency
466 should be the same as that used during the peak-to-noise-ratio refinement described above,
467 and the correlation should be relatively high (we typically use 0.8, but this can be refined by the
468 user). The resulting merged seeds can be visualized on the max projection. Since the main
469 purpose of this step is to alleviate computation demands for downstream steps, it is fine to have
470 multiple seeds for a single visually distinct cell. However, users should make sure each of the
471 visually distinct cells still has at least one corresponding seed after the merge.

472 Initialize spatial and temporal matrices from seeds

473 The last step before implementing CNMF is to initialize the spatial and temporal matrices for the
474 CNMF algorithm from the seeds. These matrices are generated with one dimension
475 representing each putative cell and the other representing each pixel or time, respectively. In
476 other words, the spatial matrix represents the spatial footprint for each cell at each pixel location
477 and the temporal matrix represents the temporal fluorescence value of each cell on each frame.
478 We assume each seed is the center of a potential cell, and we first calculate the spatial footprint
479 for each cell by taking the cosine similarity between the temporal trace of a seed and the pixels
480 surrounding that seed. In other words, we generate the weights in the spatial footprint by
481 computing how similar the temporal activities of each seed are to the surrounding pixels. Then,
482 we generate the temporal activities for each potential cell by taking the input video and
483 weighting the contribution of each pixel to the cell's temporal trace by the spatial footprint of the
484 cell. The final products are a spatial matrix and a temporal matrix.

485 Besides the two matrices representing neuronal signals, there are two additional terms in the
486 CNMF model that account for background fluorescence modeled as a spatial footprint for the
487 background and a temporal trace of background activity. To estimate these terms, we subtract
488 the matrix product of our spatial and temporal matrices, which represent cellular activities, from
489 the input data. We take the mean projection of this remainder across time as an estimation of
490 the spatial footprint of the background, and we take the mean fluorescence for each frame as
491 the temporal trace of the background.

492 Users can tweak two parameters to improve the outcome and performance of this step: a
493 threshold for cosine similarity and a spatial window identifying pixels on which to perform this
494 computation. To keep the resulting spatial matrix sparse and keep irrelevant pixels from
495 influencing the temporal traces of cells, we set a threshold for the cosine similarity of temporal
496 traces compared to the seed, where pixels whose similarity value falls below this threshold will
497 be set to zero in the spatial footprint of the cell. Cosine similarity is, in essence, a correlation
498 (the scale is 0-1) and thresholds of 0.5 and higher work well in practice. Computing many pair-
499 wise similarity measurements is computationally expensive, and it is unnecessary to compute
500 the similarities between pixels that are far apart because they are unlikely to have originated
501 from the same cell. We therefore set a window size to limit the number of pixel pairs to be
502 considered. This size should be set large enough so that it does not limit the size of spatial
503 footprints, but not unnecessarily large to the extent where it will impact performance. In practice,
504 a window size equal to the maximum expected cell diameter is reasonable.

505 *CNMF*

506 Estimate spatial noise

507 *CNMF* requires that we first estimate the spatial noise over time for each pixel in the input video.
508 The spatial noise of each pixel is simply the power of the high frequency signals in each pixel.
509 The critical parameter here is again the cut-off frequency for ‘noise’, and users should employ
510 the visualization tools as described above during peak-to-noise ratio refinement to determine
511 this frequency (see Refinement with peak-to-noise ratio).

512 Spatial update

513 Next, we proceed to the spatial update of the *CNMF* algorithm. The original paper describing
514 this algorithm [37] contains a detailed theoretical derivation of the model. Here, we provide only
515 a conceptual overview of the process so that users can understand the effect of each
516 parameter. The *CNMF* framework models the input video to be the product of the spatial and
517 temporal matrices representing signals contributed by real cells, a background term, and
518 random noise. In equation form, this is $\mathbf{Y} = \mathbf{A}\mathbf{C} + \mathbf{B} + \mathbf{E}$, where \mathbf{Y} represents the input video, \mathbf{A}
519 represents the spatial matrix containing the spatial footprints for all putative cells, \mathbf{C} represents
520 the temporal matrix containing the calcium dynamics for all putative cells, \mathbf{B} represents the
521 spatial-temporal fluctuation of background, and \mathbf{E} represents error or noise. Since the full
522 problem of finding proper \mathbf{A} and \mathbf{C} matrices is hard (non-convex), we break down the full
523 process into spatial update and temporal update steps, where iterative updates of \mathbf{A} and \mathbf{C} are
524 carried out, respectively. Each iteration will improve on previous results and eventually converge
525 on the best estimation.

526 During the spatial update, given an estimation of the temporal matrix and the background term,
527 we seek to update the spatial matrix so that it best fits the input data, along with the
528 corresponding temporal traces. To do so, we first subtract the background term from the input
529 data so that the remainder is composed only of signals from cells and noise. Then, for each
530 pixel, the algorithm attempts to find the weights for each cell’s spatial footprint that best
531 reproduces the input data (\mathbf{Y}) with the constraint that individual pixels should not weigh on too
532 many cells (controlled through what is called a sparseness penalty). To reduce computational
533 demand, we do this for each pixel independently and in parallel to improve performance, while
534 retaining the ‘demixing’ power of the *CNMF* algorithm by updating the weights for all cells
535 simultaneously. In the optimization process, the function to be minimized contains both
536 a squared error term to assess error, and an ℓ_1 -norm term to promote sparsity [16]. The
537 optimization process can be expressed formally as:

$$\begin{aligned} 538 \quad & \underset{\mathbf{A}, \mathbf{b}}{\text{minimize}} && \|\mathbf{Y}(p, :) - \mathbf{A}(p, :)\mathbf{C} - \mathbf{b}\mathbf{f}\| + \lambda \|\mathbf{A}(p, :)\|_1 \\ & \text{subject to} && \mathbf{A}, \mathbf{b} \geq 0 \end{aligned}$$

539 Where $\mathbf{Y}(p, :)$ denotes the input movie data indexed at p -th pixel, $\mathbf{A}(p, :)$ denotes the spatial
540 matrix indexed at p -th pixel across all putative cells, and \mathbf{C} , \mathbf{b} , \mathbf{f} denotes the temporal matrix, the
541 spatial footprint of background term, and the temporal fluctuation of background term,

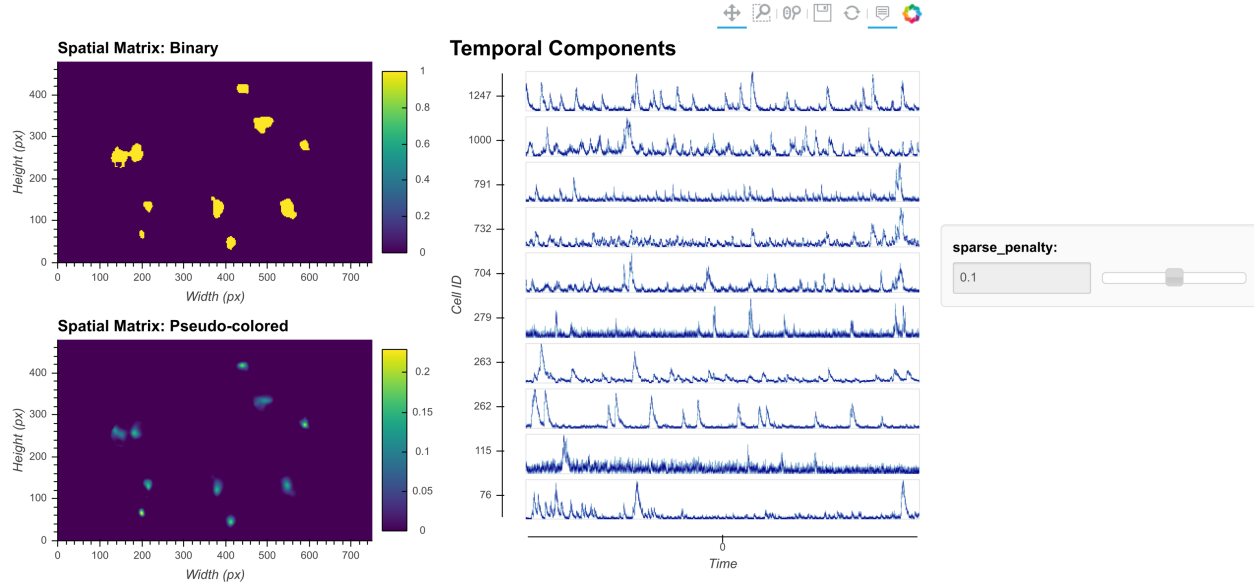
542 respectively. The scalar λ represents the sparse penalty that controls the balance between the
543 error term and sparsity term.

544 Lastly, the spatial footprint of the background term is updated in the exact same way, together
545 with other putative cells. However, the background term the temporal activity used in the spatial
546 update is not constrained by the autoregressive model. After the spatial footprint of the
547 background term is updated, we subtract the neural activity (AC) from the input data to get
548 residual background fluctuations. Then the temporal activity of background term is calculated as
549 the projection of residual onto the new background spatial footprint, where the raw activities of
550 each pixel is weighted by the spatial footprint.

551 In other CNMF implementations, the estimated spatial noise is used to determine the scaling of
552 the ℓ_1 -norm term in the target function and control the balance between error and sparsity of the
553 result. However, in practice we find that it does not always give the best result for all types of
554 datasets. For example, sometimes the estimated spatial noise is too large, which results in an
555 overly-conservative estimation of spatial footprints. Hence, we have introduced a sparseness
556 penalty on top of the estimated scaling factor for the ℓ_1 -norm term. This parameter gives users
557 more control over how sparsity should be weighted in the updating process. The higher the
558 number, the higher the penalty imposed by the ℓ_1 -norm, and the more sparse the spatial
559 footprints will become. The effect of this parameter can be visualized with the tool shown in
560 Figure 10. Users can employ this tool to determine the best sparseness penalty for their data,
561 where the binarized spatial footprint representing non-zero terms should approach the visible
562 part of the spatial footprint as much as possible, without reducing the amplitude of spatial
563 footprints to the extent that cells are discarded in the spatial update. Figure 11 shows an
564 example of the effect of changing the sparseness penalty on the resulting spatial footprints. A
565 sparseness penalty of 0.1 is considered appropriate in this case. When the sparseness penalty
566 is set much lower, many of the additional 'fragments' begin to appear in the binarized spatial
567 footprint, even if they are not part of the cell. On the other hand, when the sparseness penalty is
568 set too high, some cells are discarded. In the interactive visualization tool, users can inspect the
569 temporal dynamics of these discarded cells. In general, however, we do not recommend
570 exploiting the sparseness penalty during the spatial update to filter cells since this step does not
571 have an explicit model of the temporal signal and thus has no power to differentiate real cells
572 from noise.

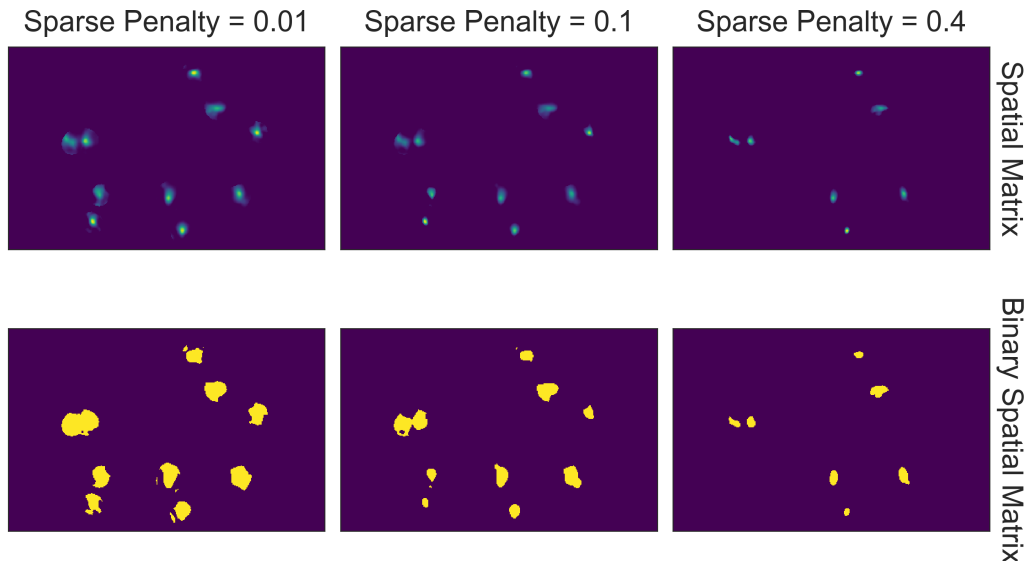
573 In addition, a dilation window parameter must be specified by the user. To reduce the amount of
574 computation when calculating how each pixel weighs onto each cell, we only update weights for
575 cells that are close to each pixel. For each cell, an ROI is computed by performing a
576 morphological dilation process on the previous spatial footprints of that cell. If a pixel lies outside
577 of a cell's region of interest, this cell will not be considered when updating the pixel's weight.
578 Thus, the dilation window parameter determines the maximum distance a cell is allowed to grow
579 during the update compared to its previous spatial footprints. This parameter should be set large
580 enough so that it does not interfere with the spatial update process, but at the same time not so
581 large as to impact performance. The expected cell diameter in pixels is a good starting point.

sparse_penalty: 0.1



582

583 **Figure 10: Visualization of spatial updates.** Here 10 cells are randomly chosen to pass
584 through spatial update with different parameters. The resulting spatial footprints, as well as
585 binarized footprints, are plotted. In addition, the corresponding temporal traces of cells are
586 plotted. The user can visually inspect the size and shape of the spatial footprints and at the
587 same time easily determine whether the results are sparse enough by looking at the binarized
588 footprints.



589

590 **Figure 11: Effect of sparseness penalty in spatial update.** Here the sum projection of the
591 spatial matrix and binarized spatial matrix are shown for 3 different sparse penalties. A
592 sparseness penalty of 0.1 is considered appropriate in this case. When the sparseness penalty
593 is set lower, artifacts begin to appear. On the other hand, when the sparseness penalty is set
594 higher, cells are dropped out.

595 Temporal update

596 Next, we proceed to the temporal update of the CNMF algorithm. Please refer to the original
597 paper for the detailed derivation [37]. Here, given the spatial matrix and background terms,
598 we update the temporal matrix so that it best fits the input data (Y). First, we subtract the
599 background term from the input data, leaving only the noisy signal from cells. We then project
600 the data onto the spatial footprints of cells, obtaining the temporal activity for each cell. Next we
601 estimate a contribution of temporal activity from neighboring overlapping cells using the spatial
602 footprints of cells, and subtract it from the temporal activity of each cell. This process results in a
603 two-dimensional matrix representing the raw temporal activity of each cell [41].

604 The CNMF algorithm models the relationship between the underlying ‘spiking’ and the calcium
605 dynamics of a cell as an auto-regressive (AR) process. It should be noted that although the
606 underlying process that drives calcium influx is presumably cell firing, the ‘spiking’ signal is
607 modeled as a continuous variable rather than a binary variable, and strictly speaking, it is only a
608 de-convolved calcium signal. Following convention, we will refer to this variable as ‘spike signal’,
609 an approximation of the underlying cellular activity that drives calcium influx. It should be
610 understood, however, that the exact relationship between this variable and the actual firing rate
611 of cells is unclear, since the absolute amount of fluorescence generated by a single spike, as
612 well as the numerical effect of integrating multiple spikes on the resulting calcium signal, is
613 unknown.

614 We first estimate the coefficients for the AR model. The coefficients of the AR model can be
615 conveniently estimated from the autocorrelation of the estimated temporal activity. In addition,
616 noise power for each cell is also estimated directly from the signal. In practice, we find that
617 during the estimation of the AR model parameters, it is helpful to first smooth the signal,
618 otherwise the time constant of the AR model tends to be biased by high frequency noise. Users
619 should again use the peak-to-noise-refinement cut-off frequency for both estimation of the noise
620 power and smoothing of the signals. Finally, we update the temporal matrix by minimizing a
621 target function for different cells, similar to what was done with the spatial matrix. Again, the
622 target function contains a squared error term and a ℓ_1 -norm term. We also introduce a
623 sparseness penalty parameter to control the balance between the two terms. The squared error
624 term contains the difference between input signal and estimated calcium dynamics, while
625 the ℓ_1 -norm term regulates the sparsity of the “spiking” signal. Pre-estimated AR coefficients
626 allow for a determined relationship between the ‘spiking’ signal and calcium dynamics for a
627 given cell. Thus, the problem can be transformed and simplified as minimizing the target
628 function over ‘spiking’ signals of different cells.

629 In practice, it is computationally more efficient to break down the minimization problem into
630 smaller pieces and update subsets of cells independently and in parallel. To do so, we first
631 identify non-overlapping cells using a Jaccard index, which measures the amount of overlap
632 between the spatial footprints of different cells. Once we identify these individual cells, we can
633 update them independently so that an optimization problem and target function are formulated
634 for each cell independently. Here, we set a cutoff Jaccard index where cells above this amount
635 of overlap are updated in parallel. During the updating process, two additional terms are
636 introduced: a baseline term to account for constitutive non-zero activity of cells and an initial
637 calcium concentration to account for a ‘spiking’ that started just prior to recording. The initial

638 calcium concentration term is a scalar that is recursively multiplied by the same AR coefficient
639 estimated for the cell. The resulting time trace, modeling the decay process of a ‘spiking’ event
640 prior to the recording, is added on top of the calcium trace. The baseline activity term is also a
641 scalar that is simply added on top of all the modeled signals. Both terms are often zero, but they
642 are nevertheless saved and visualized. For each cell, the optimization process can be
643 expressed formally as:

$$644 \quad \begin{array}{ll} \text{minimize} & \| \mathbf{y}_{ra} - \mathbf{c} - b_0 - c_0 \mathbf{d} \| + \lambda \| \mathbf{Gc} \|_1 \\ \text{subject to} & \mathbf{c}, \mathbf{Gc} \geq 0 \end{array}$$

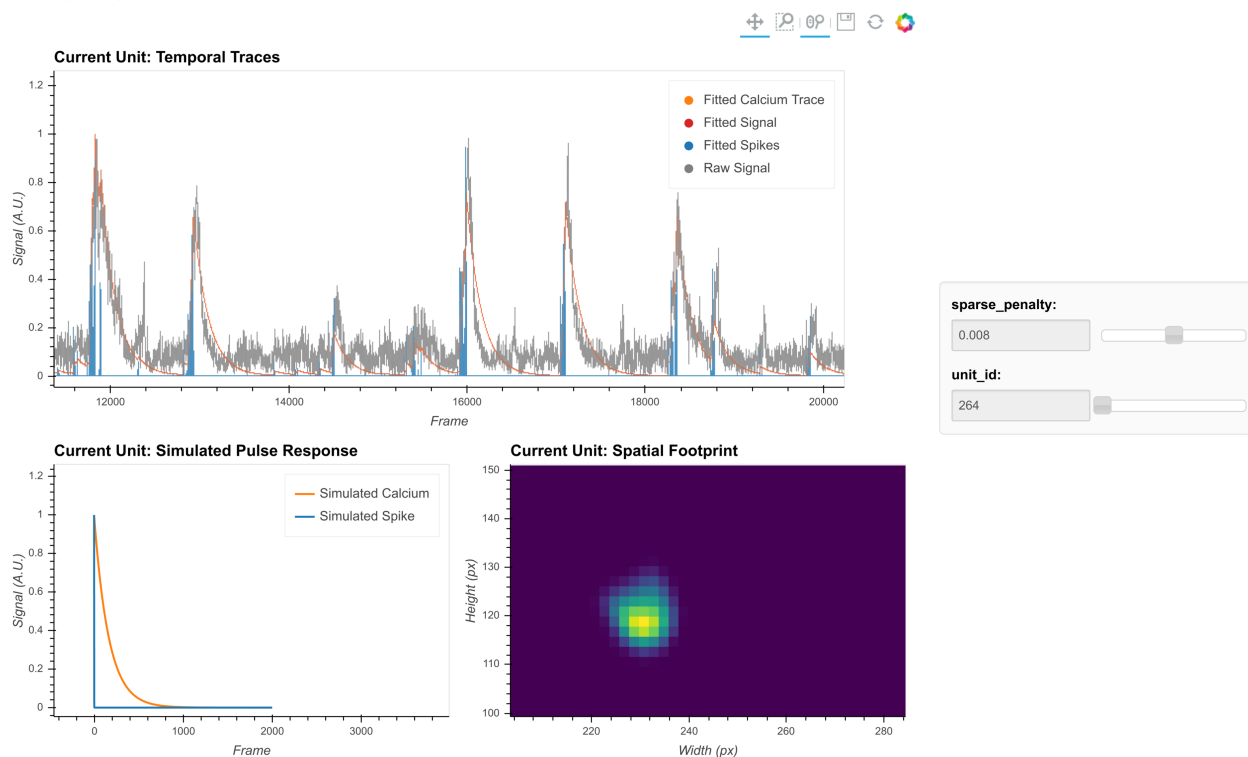
645 Where \mathbf{y}_{ra} denotes the input movie data projected onto the spatial footprint of the given cell, \mathbf{c}
646 denotes the estimated calcium dynamic of the given cell, b_0 denotes the constant baseline
647 fluorescent activity, c_0 denotes the initial calcium concentration, G represent a matrix of AR
648 coefficients such that \mathbf{Gc} is the estimated ‘spike’ signal, \mathbf{d} is a vector representing the temporal
649 decay of a single spike based on the estimated AR coefficients, such that the term $c_0 \mathbf{d}$
650 represent the contribution of initial calcium concentration. Similar to spatial update, the scalar λ
651 represents the sparse penalty and controls the balance between the error term and sparsity
652 term.

653 The ℓ_1 -norm in the optimization problem is known to reduce not only the number of non-zero
654 terms (i.e., promotes sparsity), but also the amplitude/value of non-zero terms. This effect is
655 unwanted, since in some cases the numerical the spatial update step in CNMF algorithm
656 and value of the resulting ‘spike’ signal can become too small as a side-effect of promoting
657 sparsity, making it hard to interpret and compare the ‘spike’ signal for downstream analysis. To
658 counteract this phenomenon, we introduce a *post hoc* scaling process. After the temporal
659 update, each cell is assigned a scaling factor to scale all the fitted signals to the appropriate
660 values. The scaling factor is solved by least square minimizing the error between the fitted
661 calcium signal and the projected raw signal.

662 The critical parameters in temporal updates are as follows: (1) The order of the AR model,
663 usually 1 or 2. Users should choose 1 if near-instantaneous rise time is presented in the calcium
664 dynamics of the input data (i.e., from the relatively slow sampling rate) and should choose 2
665 otherwise. (2) The cut-off frequency for noise used for both noise power estimation and pre-
666 smoothing of the data during AR coefficients estimation. Users should use the values set during
667 peak-to-noise ratio refinement. (3) The threshold for the Jaccard index determining which cells
668 can be updated independently. Users should use a value as low as possible, as long as the
669 speed of this step is acceptable (with large amounts of cells packed closely together, a low
670 threshold may dramatically slow down this step), or visually inspect how sparse the spatial
671 footprints are and determine what amount of overlap between spatial footprints results in
672 significant crosstalk between cells. (4) The sparseness penalty is best set through visualization
673 tools. The effect of any parameter on the temporal update can be visualized through the tool
674 shown in Figure 12, where the result of the temporal update for 10 randomly selected cells are
675 plotted as traces. There are a total of 4 traces shown for each cell: the calcium signal, the
676 deconvolved ‘spiking’ signal, the projected raw signal, and the ‘fitted signal’. The ‘fitted signal’ is
677 very similar to the calcium signal and is often indistinguishable from the latter. The difference
678 between them is that the ‘fitted signal’ also includes the baseline term and the initial calcium

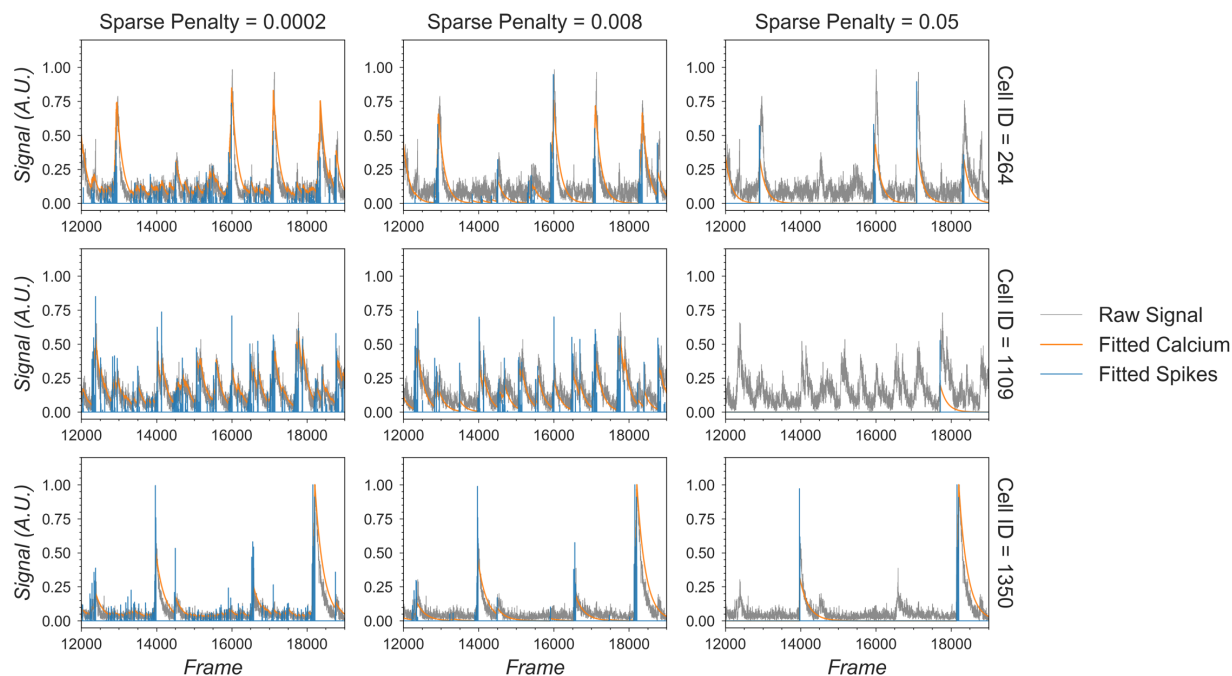
679 concentration term. Hence, the ‘fitted signal’ should better follow the projected raw signal, but it
680 may be less interesting for downstream analysis. Toggling between different parameters
681 triggers the dynamic update of the plots, helping the user to determine the best parameters for
682 their data. Additionally, we highlight the effect of the sparseness penalty on resulting fitted
683 calcium signals and spike signals in Figure 13. The effect is most evident in the ‘fitted spikes’
684 trace, which corresponds to the spike signal and can arguably be interpreted as a measure of
685 the underlying neural activity per frame scaled by an unknown scalar. Here, a sparseness
686 penalty of 0.008 is considered most appropriate. A lower sparseness penalty will introduce
687 many false positive signals which do not correspond to real calcium dynamics, as can be seen
688 in the plots. On the other hand, too high a sparseness penalty will produce false negatives
689 where clear rises in the raw signal are not accompanied by spikes.

sparse_penalty: 0.008, unit_id: 264



690

691 **Figure 12: Visualization of temporal update.** Here, a subset of cells is randomly chosen to
692 pass through temporal updates with different parameters. Only one cell is visualized at a given
693 time and the cell can be selected using the slider on the right. The raw signal, the fitted signal,
694 the fitted calcium traces, and the spike signals are overlaid in the same plot. In addition, a
695 simulated pulse-response based on the estimated auto-regressive parameters is plotted with
696 the same time scale. Furthermore, the corresponding spatial footprint of the cell is plotted for
697 cross-reference. With a given set of parameters, the user can visually inspect whether the
698 pulse-response captures the typical calcium dynamics of the cell, and whether the timing and
699 sparsity of the spike signal fit well with the raw data. The data shown here was acquired with a
700 framerate of 30 fps.



701

702 **Figure 13: Effect of the sparseness penalty in temporal update.** Here, 3 example cells are
703 selected and passed to the temporal update with different sparseness penalties. The “Raw
704 Signal” corresponds to the input video projected onto predetermined spatial footprints. The
705 “Fitted Calcium” and “Fitted Spikes” correspond to the resulting model-fitted calcium dynamics
706 and spike signals. A sparseness penalty of 0.008 (middle column) is considered appropriate in
707 this case. The data shown here was acquired with a framerate of 30 fps.

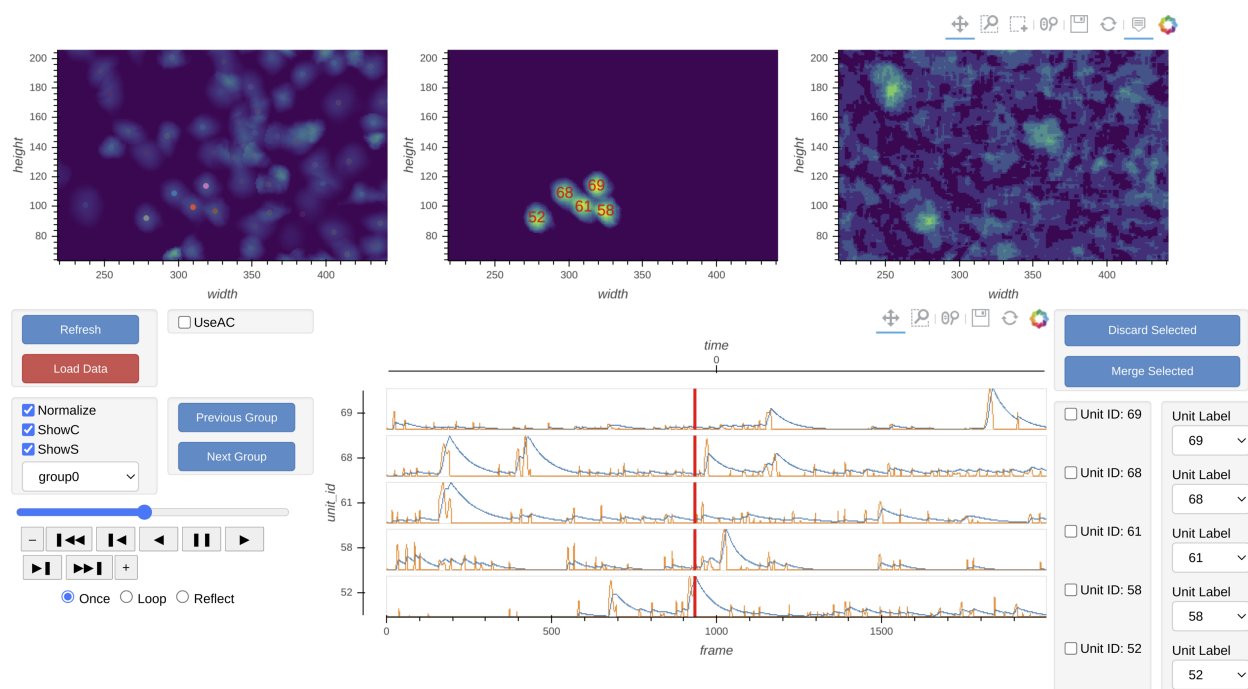
708 Merging cells

709 The CNMF algorithm can sometimes misclassify a single cell as multiple cells. To counteract
710 this phenomenon, we implement a step to merge cells based on their proximity and temporal
711 activity. All cells with spatial footprints sharing at least one pixel are considered candidates for
712 merging, and the pair-wise correlation of their temporal activity is computed. Users can then
713 specify a threshold where cell pairs with activity correlations above the threshold are merged.
714 Merging is done by taking the sum of the respective spatial footprints and the mean of all of the
715 temporal traces for all cells to be merged. Since this is only a simple way to correct for the
716 number of estimated cells and does not fit numerically with what the model CNMF assumes,
717 merging is only done between iterations of CNMF, but not at the end.

718 Manual curation

719 Minian provides an interactive visualization to help the users manually inspect the quality of
720 putative cells and potentially merge or drop cells. At any given time, the visualization shows
721 spatial temporal activities (top row, middle panel in Figure 14) and temporal dynamics of a
722 selected subset of cells (bottom row in Figure 14). The spatial temporal activities are shown side
723 by side with the spatial footprints of all cells and the pre-processed movie (input to CNMF
724 algorithm) at a given frame (top row of Figure 14). The field of view is synchronized across the

725 three images on the top, so that the users can easily zoom in and compare the estimated spatial
726 footprints of cells to the input data. The spatial temporal images in the middle show the product
727 of spatial footprints and calcium dynamics, which represent the model estimated image of a
728 subset of cells at a given frame. This spatial temporal product is calculated on-the-fly and
729 synchronized with the frame indicators on the temporal dynamic plots. In this way users can
730 easily pick times of interest (for example, when a cell has a calcium event), and validate
731 whether the estimated spatial temporal activities match the input data. Lastly, this interactive
732 visualization allows the user to either drop false positive cells or merge multiple cells together
733 via dropdown menus. The result of manual curation is saved as an array with a label for each
734 unit indicating whether a cell should be discarded or how several cells should be merged. In this
735 way, only the new label is saved and no data is modified, allowing the user to repeat or correct
736 the manual curation process if needed.



738 **Figure 14: Interactive visualization of Minian output.** The three images on the top show the
739 spatial footprints of all the cells (left), the spatial temporal activities of selected subset of cells
740 (middle), and the pre-processed data. The bottom row shows the display control panel (left), the
741 temporal dynamics of selected subset of cells (middle), and the manual curation panel (right).
742 The field of view, current frame, and selection of cells are all synced across different plots to
743 help user focus on a specific region and time. The users can use the control panel to select
744 groups of cells, change display options for temporal dynamics and spatial temporal activities,
745 change the current frame or play the movie. In addition, the users can directly select cells from
746 the spatial footprints plot on the top left. The users can also directly jump to frames by double-
747 clicking on the temporal dynamic plots. These interactive features help the users quickly focus
748 on region and time of interests. The manual curation menu on bottom right can be used to
749 assign unit labels to each cell, which indicate whether a cell should be dropped or merged.

750 *Cross registration*

751 After completing the analysis of individual recording sessions, users can register cells across
752 sessions. While more complex approaches are proposed in other pipelines [16,17], here, our
753 intention is simplicity. To account for shifts in the field of view from one session to the next, we
754 first align the field of view from each session based upon a summary frame. Users can either
755 choose a max projection of each pre-processed and motion-corrected video, or a summed
756 projection of the spatial footprints of all cells. Users can also choose which session should be
757 used as the template for registration, to which every other session should be aligned. We use a
758 standard cross-correlation based on a template-matching algorithm to estimate the translational
759 shifts for each session relative to the template and then correct for this shift. The weighted
760 centroid of each cell's spatial footprint is then calculated and pair-wise centroid distances are
761 used to cross-register cells. A distance threshold (maximum pixel distance) is set. Users should
762 choose this threshold carefully to reflect the maximum expected displacement of cells across
763 sessions after registration. We found that a threshold of 5 pixels works well. Finally, a pair of
764 cells must be the closest cells to each other in order to be considered the same cell across
765 sessions.

766 To extend this method to more than two sessions, we first cross-register all possible session
767 pairs. We then take the union of all these pair-wise results and transitively extend the cross-
768 registration across more than two sessions. At the same time, we discard all matches that result
769 in conflicts. For example, if cell A in the first session is matched with cell B in the second
770 session, and cell B is in turn matched with cell C in the third session, but cells A and C are not
771 matched when directly registering the first and third sessions, all of these matches are
772 discarded and all three cells are treated as individual cells. We recognize that this approach
773 might be overly conservative. However, we believe that this strategy provides an easy-to-
774 interpret result that does not require users to make decisions about whether to accept cell pairs
775 that could conflict across sessions.

776 To save computation time, we implement a moving window where centroid distances are only
777 calculated for cell pairs within these windows. Users should set the size of windows to be much
778 larger than the expected size of cells.

779 *Hardware and dependencies*

780 Minian has been tested using OSX, Linux, and Windows operating systems. Additionally,
781 although we routinely use Minian on specialized analysis computers, the pipeline works on
782 personal laptops for many common length (~30min) miniature microscope experiments.
783 Specifications of all of the computers that have been tested can be found in Tested hardware
784 specifications. We anticipate that any computer with at least 16GB of memory will be capable of
785 processing at least 20 minutes of recording data, although increased memory and CPU power
786 will speed up processing. Moreover, due to the read-write processes involved in out-of-core
787 computation, we recommend that the videos to be processed are held locally at the time of
788 analysis, preferably on a solid-state drive. The relatively slow speed of transfer via ethernet
789 cables, Wi-Fi, or USB cables to external drives will severely impair analysis times.

790 Minian is built on top of project Jupyter [42], and depends heavily on packages provided by the
791 open-source community, including numpy [43], scipy [44], xarray [45], holoviews [46],
792 bokeh [47], opencv [48], and dask [49]. A complete list of direct dependencies for Minian can be
793 found in List of dependencies. Of note, the provided install instructions handle the installation of
794 all dependencies.

795 **Results**


796 To validate the accuracy as well as benchmark the performance of Minian, we ran the Minian
797 pipeline on a series of simulated and experimental datasets and compare the output and
798 performance to those obtained with CalmAn, which is one of most widely-adopted calcium
799 imaging analysis pipeline in the field. In addition, we also validated the full workflow of Minian by
800 applying the pipeline to several recordings of animals running on a linear track and looked at the
801 stability of place cells. These results are presented in sections below.

802 *Validation with simulated datasets*

803 We first validated Minian with simulated datasets. We synthesized different datasets with
804 varying number of cells and signal levels based on existing works [11,12]. The simulated
805 datasets contain local background fluctuations, noise, and motions similar to experimental
806 datasets (See Generation of simulated datasets for details). The field of view contains 512 x 512
807 pixels and 20000 frames, corresponding to roughly 10 minutes of recording at 30 fps. We
808 processed the data with both Minian and CalmAn. For Minian, we utilized the visualization
809 described here to optimize the parameters. For CalmAn, we used the same parameters as
810 Minian whenever the implementations were equivalent. Otherwise, we followed the suggested
811 parameters and tweaked them based on the knowledge of simulated ground truth.

812 To compare the results objectively, we first matched the resulting putative cells from the output
813 of Minian or CalmAn to the simulated ground truth (See Matching neurons for validation for
814 details). We then calculated three metrics to measure the quality of output: F1 score, spatial
815 footprints correlation, and temporal dynamics correlation. The F1 score is defined as the
816 harmonic mean of precision (proportion of detected neurons that are true) and recall (proportion
817 of ground truth neurons that has been detected). Hence the F1 score measures the overall
818 accuracy of neuron detection. For each detected neuron that has been matched to ground truth,
819 we compute Pearson correlation between the estimated and ground truth spatial footprint, as
820 well as the Pearson correlation between the estimated calcium dynamic and the ground truth
821 calcium dynamic. We then take the median correlation across all the matched neurons to
822 measure the overall quality of estimated spatial footprints and temporal dynamics.

823 As shown in Figure 15, both Minian and CalmAn achieve similar and near perfect levels (> 0.95)
824 of F1 score across all conditions. Similarly, the spatial footprints remain nearly perfect (> 0.95)
825 for both pipelines across all conditions. At the lowest signal level (0.2), both pipelines suffer from
826 decreased correlation of temporal dynamics. This is likely due to noise and background
827 contaminating the true signal. Overall, these results show that the Minian and CalmAn pipelines
828 perform similarly well in terms of output accuracy on simulated datasets.

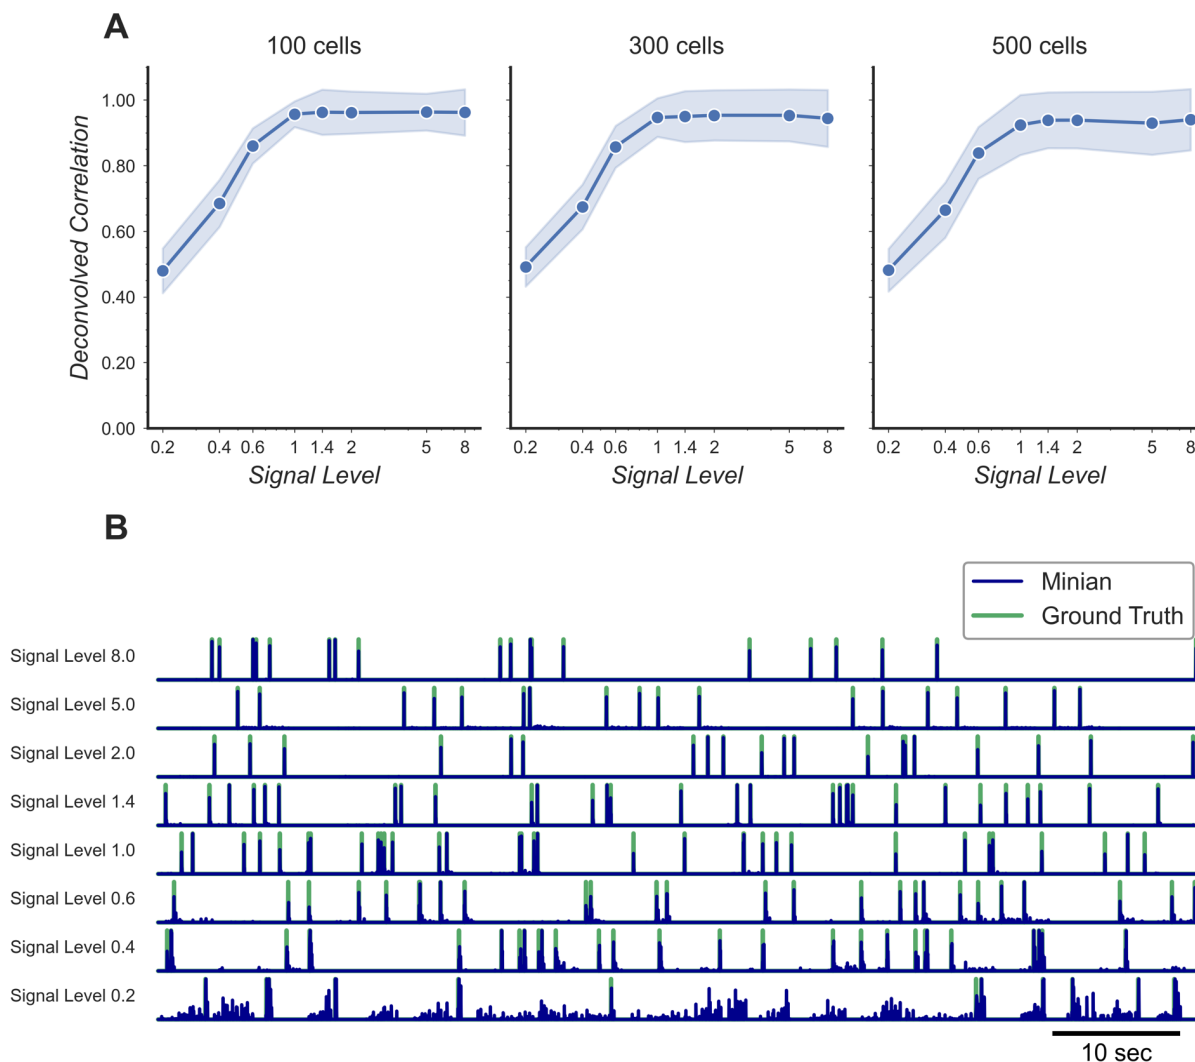


829

830 **Figure 15: Validation of Minian with simulated datasets.** Simulated datasets with varying
 831 signal level and number of cells are processed through Minian and CalmAn. The F1 score (top),
 832 median correlation of spatial footprints (middle), and median correlation of temporal dynamics
 833 (bottom) are plotted as a function of signal level. Both pipelines achieve near perfect (> 0.95) F1
 834 scores and spatial footprint correlation across all conditions. The correlation of temporal
 835 dynamics are lower when the signal level is 0.2, but remains similar across the two pipelines
 836 overall.

837 Additionally, we want to validate the deconvolved signal from Minian output, since this is usually
 838 the most important output for downstream analysis. Our ground truth spikes are simulated as
 839 binary signals. However, in reality calcium activity often reflect the integration of several spikes,
 840 and the deconvolved signals from Minian output are real-valued. Because of this, we down-
 841 sampled both the ground truth spikes and deconvolved signals by 5 times, and then calculated
 842 Pearson correlation for all matched cells. The resulting correlation is summarized in Figure 16
 843 A. Our results indicate that the deconvolved output from Minian is highly similar to ground truth
 844 spikes when signal level is high, and the correlation asymptote and approach 1 when signal
 845 level is higher than 1. The lower correlation corresponding to low signal level is likely due to the

846 background and noise contamination being stronger than signal. In line with this idea, the
847 detected “spikes” from the deconvolved signals closely match those from ground truth, as
848 shown by the example traces in Figure 16 B. The main difference between the two traces is the
849 amplitude of the deconvolved signals, which is prone to be influenced by local background and
850 noise. Overall, these results suggest that Minian can produce deconvolved signals that are
851 faithful to ground truth and suitable for downstream analysis.



852


853 **Figure 16: Validation of deconvolved signal from Minian.** (A) Correlation of deconvolved
854 signals from Minian output with simulated ground truth. The mean correlation across all cells
855 (blue line) and the standard deviation (light blue shade) are shown separately for different signal
856 levels and number of cells. The correlation asymptote and approach 1 when signal level is
857 higher than 1. (B) Example deconvolved traces from Minian output overlaid with simulated
858 ground truth. One representative cell is drawn from each signal level. The binary simulated
859 spikes are shown in green, with the real-valued Minian deconvolved output overlaid on top in
860 blue. The deconvolved signals closely match the ground truth and the main difference between

861 *the two signals is in the amplitude of the deconvolved signals, which tend to be influenced by*
862 *local background.*


863 *Validation with experimental datasets*

864 We next validated Minian with experimental datasets. The data was collected from hippocampal
865 CA1 regions in animals performing a spatial navigation task. 6 animals with different density of
866 cells were included in the validation dataset. The recordings are collected with 608 x 608 pixels
867 at 30 fps and lasts 20 min (~36000 frames). Due to difficulties in obtaining ground truth for
868 experimental data, we choose to validate Minian with CalmAn, which has been established as
869 one of the most accurate existing pipelines. To evaluate the results objectively, we matched
870 resulting ROIs from Minian with those from CalmAn using the same approach as in the
871 Validation with simulated datasets section. We then calculated correlation of spatial footprints
872 and temporal activity between matched ROIs from the two pipelines. Across the 6 datasets, the
873 mean F1 score is 0.73 (sem +/-0.03). The mean spatial footprints correlation is 0.84 (sem +/-
874 0.02), and the mean temporal activity correlation is 0.86 (sem +/-0.02). An example field of view
875 and temporal activity from matched ROIs are shown in Figure 17. Our results indicate that most
876 of the ROIs detected by Minian and CalmAn correspond to the same population of putative
877 cells, and the resulting spatial footprints and temporal activity are nearly identical. These cells
878 tend to cluster near the center of the field of view, which usually have better signal-to-noise
879 ratio. However, the cells near the edge of the field of view usually have low intensity and spatial
880 consistency due to the optical property of GRIN lens. As a consequence, Minian and CalmAn
881 might detect different population of cells near the border of field of view, due to differences in
882 pre-processing and initialization between the two pipelines. We have chosen to use the same
883 set of parameters across all datasets so that the results are easier to interpret, hence the
884 parameters we used were relatively conservative. In practice, the users can further fine-tune the
885 parameters for each recording so that Minian would be able to capture all the low signal cells in
886 the field of view. Overall, these results suggest that the output of Minian is highly similar to
887 CalmAn when analyzing experimental datasets.

A



B



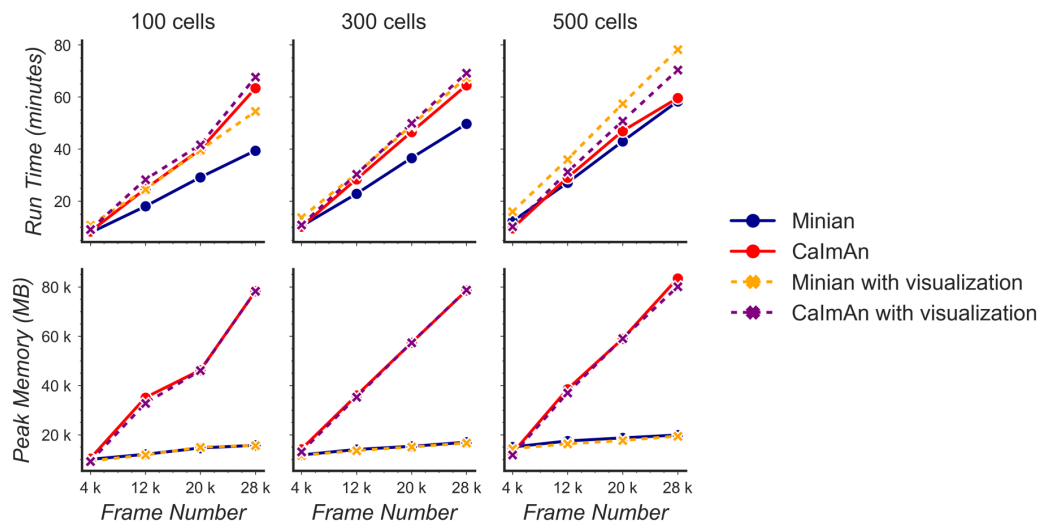
888

889 **Figure 17: Example output of Minian and CalmAn with experimental datasets. (A)** An
890 example field-of-view from one of the experimental datasets. The spatial footprints from Minian
891 and CalmAn are colored as blue and red respectively, and overlaid on top of each other. Most
892 of the spatial footprints from both pipelines overlap with each other. **(B)** 5 example matched
893 temporal activity from Minian and CalmAn overlaid on top of each other. The extracted temporal
894 activity are highly similar across the two pipelines.

895 *Benchmarking computational performance*


896 To see how the performance of Minian scales with different input data size, we synthesized
897 datasets with varying number of cells and number of frames (recording length). The field of view
898 contains 512 x 512 pixels (same as those used in validation of accuracy), and the signal level
899 was held constant at 1 to make sure both Minian and CalmAn can detect roughly equal number
900 of neurons during the pipeline. To this end, we tracked two metrics of performance: the total
901 running time of the pipeline and the peak memory usage during running. The running time was
902 obtained by querying operating system time during the pipeline. The memory usage was
903 tracked with an independent process that queries memory usage of the pipeline from the
904 operating system on a 0.5 seconds interval. Both pipelines were set to utilize 4 parallel
905 processes during the run across all conditions. All benchmarking are carried out on a custom-
906 built linux machine (Model "Carbon" under Tested hardware specifications)

907 As shown in Figure 18, the run time of both Minian and CalmAn scales linearly as a function of
908 input recording length. The exact running times vary depending on number of cells as well as
909 whether visualization is included in the processing, but in general the running time is similar
910 across both pipelines. On the other hand, the peak memory usage of CalmAn scales linearly
911 with recording length when the number of parallel processes was set to be constant. At the
912 same time, the peak memory usage of Minian stays mostly constant across increasing number
913 of frames. This is likely due to the flexible chunking implementation of Minian (See Parallel and
914 out-of-core computation with dask), where Minian was able to break down computations into
915 chunks in both the spatial and the temporal dimensions depending on which way is more
916 efficient. In contrast, CalmAn only splits data into different spatial chunks (patches), resulting in
917 a linear scaling of memory usage with recording length for each chunk-wise computation.
918 Additionally, we run Minian and CalmAn with different number of parallel processes on the
919 simulated dataset with 28000 frames and 500 cells. As expected, with more parallel processes
920 the performance improves and the run time decreases but at the same time the total peak
921 memory usage increases. The tradeoff between run time and peak memory usage are shown in
922 Figure 19. In conclusion, these results show that in practice, Minian is able to perform as fast as
923 CalmAn, while maintaining near constant memory usage regardless of input data size. This
924 allows the users to process much longer recordings with limited RAM resources.



925

926 **Figure 18: Benchmarking of computational performance.** Data with varying number of cells
 927 and frames were processed through Minian and CalmAn. The run time (top) and peak memory
 928 usage (bottom) were recorded and plotted as a function of frame number. For both pipelines,
 929 the run time scales linearly as a function of the number of frames and remains similar across the
 930 pipelines. However, the peak memory usage for CalmAn also scales linearly as the number of
 931 frames increases, while Minian maintains a relatively constant peak memory usage across
 932 different frame numbers and cell numbers.



933

934 **Figure 19: Tradeoff between run time and memory usage.** Simulated data with 500 cells and
 935 28000 frames were processed through Minian and CalmAn with different numbers of parallel
 936 processes. We varied the number of parallel processes from 2 to 10, and the resulting memory
 937 usage is plotted as a function of run time. For both pipelines, the curve takes a hyperbola
 938 shape, showing the tradeoff between run time and memory usage.

939 **Validation with hippocampal CA1 place cells**


940 In addition to direct validation of the output for single session, we wanted to validate the
 941 scientific significance of the spike signal, as well as the quality of the cross-session registration,
 942 and ensure that Minian is capable of generating meaningful results consistent with the existing

943 literature. We leveraged the extensively documented properties of place cells in rodent
944 hippocampal CA1 [50]. Place cells have been shown to have consistent place fields across at
945 least two days [33,51] with only a minority of detected cells undergoing place field remapping.
946 Here, we looked at place field stability across two linear track sessions (Figure 20 A). Briefly,
947 animals were trained to run back and forth on a 2 m linear track while wearing a Miniscope to
948 obtain water rewards available at either end [35]. The time gap between each session was 2
949 days. We record calcium activity in dorsal CA1 region with a FOV of 480 x 752 pixels collected
950 at 30 fps. Each recording session lasts 15 min (~27000 frames). Calcium imaging data were
951 analyzed with Minian, while the location of animals was extracted with an open-source
952 behavioral analysis pipeline ezTrack [18]. The resulting calcium dynamics and animal behavior
953 were aligned with the timestamps recorded by Miniscope data acquisition software
954 (miniscope.org). We used the spike signal for our downstream analysis. To calculate average
955 spatial activity rate, we binned the 2-meters long track into 100 spatial bins. In addition,
956 we separated the epochs when the animals are running in opposite directions, resulting in a
957 total of 200 spatial bins. We then smoothed both the binned activity rate and animal's
958 occupancy with a Gaussian kernel with a standard deviation of 5 cm. We classified place cells
959 based on three criteria: a spatial information criterion, a stability criterion, and a place field size
960 criterion [35]. (See Classification of place cells for more detail.) Finally, we analyzed cells that
961 are cross-registered by Minian and are classified as place cells in both sessions. We then
962 calculated the Pearson correlation for the average spatial firing rate for each cross-registered
963 cell. We found that, on average, place cells have a correlation of ~0.6, which is consistent with
964 the existing literature [35].

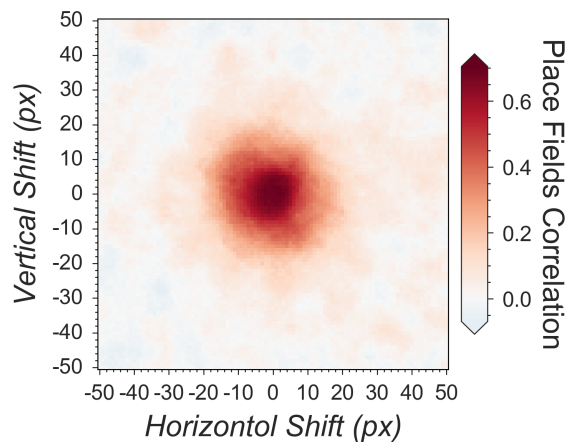
965 Next, we validated the cross-session registration to verify that the correct cells were being
966 matched across days. We translated the spatial footprints of the second session in both
967 directions up to 50 pixels and registered the cells with the shifted spatial footprints. We then
968 carried out the same analysis with the registration results from shifted spatial footprints. We
969 found that the average correlations between spatial firing patterns have higher values when the
970 shifts are close to zero (Figure 20 B).

971 In conclusion, Minian can reliably process *in vivo* calcium imaging data and produce results that
972 are in agreement with the known properties of rodent CA1. Minian can thus help neuroscience
973 labs easily implement and select the best parameters for their calcium analysis pipeline by
974 providing detailed instructions and visualizations.

A



B



975

976 **Figure 20: Validation of Minian with hippocampal CA1 place cells. (A)** Matching place cells
977 from two recording sessions. The cells are matched from one session to the other using the
978 cross-session registration algorithm and sorted based on place field in the first session. In both
979 sessions, animals run on a 2-meter-long linear track with water reward at both ends. The track
980 is divided into 200 spatial bins. The mean “firing” rate calculated from the spike signal for each
981 cell is shown. Cell IDs are assigned by Minian when each session is analyzed independently.
982 **(B)** Averaged correlations of spatial firing rates with different artificial shifts. We artificially shifted
983 the spatial footprints of the second linear track session, then carried out registration and
984 calculated a mean correlation of spatial firing rates for all place cells. The artificial shifts were
985 relative to the aligned spatial footprints and range from -50 to 50 pixels.

986 **Discussion**

987 *Making open science more accessible*

988 Neuroscience has benefitted tremendously from open-source projects, ranging from do-it-
989 yourself (DIY) hardware [1] to sophisticated algorithms [2]. Open-source projects are impactful
990 because they make cutting-edge technologies available to neuroscience labs with limited
991 resources, as well as opening the door for innovation on top of previously established methods.
992 We believe that openly sharing knowledge and tools is just the first step. Making knowledge
993 accessible even to non-experts should be one of the ultimate goals of open-source projects.

994 With the increasing popularity of miniaturized microscopes [36], there has been significant
995 interest in analysis pipelines that can reliably extract neural activities from the data. Numerous
996 algorithms have been developed to solve this problem [11,13,15,19,20,37], and many of them
997 are implemented as open-source packages that can function as a one-stop pipeline [12,14,16].
998 However, one of the biggest obstacles for neuroscience labs in adopting analysis pipelines is
999 the difficulty in understanding the exact operation of the algorithms, leading to two notable
1000 challenges: first, researchers face difficulties adjusting the parameters when the data they have
1001 collected are out of the expected scope of the pipeline's default parameters. Second, even after
1002 neural activity data is obtained, it is hard for researchers to be sure that they have chosen the
1003 best approaches and parameters for their dataset. Indeed, it has been found that depending on
1004 the features of the data and the metric used, more sophisticated algorithms do not always out-
1005 perform simpler algorithms [52], making it even harder for researchers to interpret the results
1006 obtained from some analysis pipelines. Researchers therefore often have to outsource data
1007 analysis to experts with strong computational backgrounds or simply trust the output of the
1008 algorithms being used. Minian was created to address these challenges. By providing not only
1009 detailed documentation of all functions, but also by providing rich interactive visualizations,
1010 Minian helps researchers to develop an intuitive understanding of the operations of algorithms
1011 without expertise in mathematics or computer science. These insights help researchers choose
1012 the best parameters, as well as to become more confident in their interpretation of results.
1013 Furthermore, transparency regarding the underlying algorithms enables researchers to develop
1014 in-house modifications of the pipeline, which is a common practice in neuroscience labs. We
1015 believe that Minian will contribute to the open science community by making the analysis of
1016 calcium imaging data more accessible and understandable to neuroscience labs.

1017 *Limitations*

1018 Although Minian provides users with insights into the parameter tuning process across different
1019 brain regions, these insights are achieved mainly through visual inspection. However, the
1020 performance of an analysis pipeline should be measured objectively. While calcium imaging has
1021 been validated with electrophysiology under *ex vivo* settings [53], ground-truth data for single-
1022 photon *in vivo* calcium imaging are lacking, making objective evaluation of the algorithms
1023 difficult. Therefore, here we have provided only indirect validations of the pipeline by
1024 recapitulating well-established biological findings.

1025 **Supplemental information**

1026 *Parallel and out-of-core computation with dask*

1027 In Minian, we use a modern parallel computing library called dask to implement parallel and out-
1028 of-core computation. Dask divides the data into small chunks along all dimensions, then flexibly
1029 merges the data along some dimensions in each step. We leverage the fact that each step in
1030 our pipeline can be carried out chunk by chunk independently along either the temporal (frame)
1031 dimension or the spatial (height and width) dimensions, thus requiring no interpolation or special
1032 handling of borders when merged together, producing results as if no chunking had been done.
1033 For example, motion correction and most pre-processing steps that involve frame-wise filtering
1034 can be carried out on independent temporal chunks, whereas computation of pixel correlations
1035 can be carried out on independent spatial chunks. Similarly, during the core CNMF computation
1036 steps, spatial chunking can be used during update of spatial footprints, since spatial update is
1037 carried out pixel by pixel. Meanwhile, temporal chunking can be used when projecting the input
1038 data onto spatial footprints of cells, which is usually the most memory-demanding step.
1039 Although the optimization step during the temporal update is computed across all frames and no
1040 temporal chunking can be used, we can still chunk across cells, and in practice the memory
1041 demand in this step is much smaller comparing to other steps involving raw input data.
1042 Consequently, our pipeline fully supports out-of-core computation, and memory demand is
1043 dramatically reduced. In practice, a modern laptop can easily handle the analysis of a full
1044 experiment with a typical recording length of up to 20 minutes. Dask also enables us to carry out
1045 lazy evaluation of many steps where the computation is postponed until the result is needed, for
1046 example, when a plot of the result is requested. This enables selective evaluation of operations
1047 only on the subset of data that will become part of the visualization and thus helps users to
1048 quickly explore a large space of parameters without committing to the full operation each time.

1049 *Seeds refinement with a Gaussian-Mixture-Model*

1050 As described in the main text, an alternative strategy to thresholding fluorescence intensity
1051 during seeds initialization is to explicitly model the distribution of fluorescence fluctuations of all
1052 candidate seeds and select those with relatively higher fluctuation. Here, we describe this
1053 process and the rationale. Since the seeds are generated from local maxima, they include noise
1054 from relatively empty regions with no actual cells. The seeds from these regions usually have
1055 low fluctuations in fluorescence across time and can be classified as spurious. To identify these
1056 cases, we compute a range of fluctuation for each seed (range of min-max across time), and
1057 model these ranges with a Gaussian-Mixture-Model of two components. The fluctuations from
1058 'noise' seeds compose a Gaussian distribution with low fluctuation, while seeds from actual cells
1059 assume a higher degree of fluctuation and form another Gaussian distribution with a higher
1060 mean. Any seed whose fluctuations belong to the lower Gaussian distribution is discarded in
1061 this step. To compute the range of fluctuation for each seed, we compute the difference
1062 between the 99.9 and 0.1 percentile of all fluorescence values across time, which is less biased
1063 by outliers than the actual maximum and minimum values.

1064 Normally, this step is parameter-free. In rare cases, there are regions containing noise while
1065 other regions are almost completely dark. Thus, seeds from these two regions will form two
1066 peaks in the distribution of what the user would consider 'bad seeds', and a Gaussian-Mixture-

1067 Model with two components will no longer be valid. In such cases users can tweak the number
1068 of components (number of modeled Gaussian distributions), as well as the number of
1069 components to be considered as composed of real signal. However, because the two noise
1070 distributions are likely to overlap to some degree, using two components will likely suffice. The
1071 distribution of fluctuations, the Gaussian-Mixture-Model fit, and the resulting seeds, are
1072 visualized, enabling the user to judge the appropriateness and accuracy of this step. It should
1073 be noted that in practice, we have found this process to depend heavily on the relative
1074 proportion of the 'good' and 'bad' seeds and can easily result in a significant amount of false
1075 negatives if the proportion of the 'bad' seed is too low. This makes the Gaussian-Mixture-Model
1076 approach less stable and in general less preferable to simple thresholding unless a good
1077 threshold of fluorescence intensity cannot be easily determined.

1078 *Generation of simulated datasets*

1079 We use a pipeline modified from [11] and [12] to generate simulated data for validation and
1080 benchmarking of Minian. Specifically, we generate a 512 x 512 pixels field of view with varying
1081 number of frames and neurons. The neurons are simulated as spherical 2-D Gaussian. The
1082 center of neurons are drawn uniformly from the whole field of view, and the Gaussian widths σ_x
1083 and σ_y for each neuron are drawn from $\mathcal{N}(15, 5^2)$, with a minimum value of 3. Spikes are
1084 simulated from a Bernoulli process with a 0.01 probability of spiking per frame. Calcium
1085 dynamics are simulated by convolving the spikes with a temporal kernel $g(t) = \exp(-t/\tau_d) -$
1086 $\exp(-t/\tau_r)$, with rise time $\tau_r = 5$ frame and decay time $\tau_d = 60$ frame. We simulate the spatial
1087 footprints of backgrounds as spherical 2-D Gaussian distributed uniformly across field of view.
1088 In total 300 independent background terms are used for all simulation. The Gaussian widths are
1089 drawn from $\mathcal{N}(900, 50^2)$ The temporal dynamic of backgrounds are simulated from a
1090 constrained Gaussian random walk process with steps drawn from $\mathcal{N}(0, 2^2)$, then clipped to be
1091 non-negative and gaussian smoothed temporally with a variance of 60 frames. We also simulate
1092 motion of the field of view as 2-D translations. The translational shift in each direction is
1093 simulated from a constrained Gaussian random walk process with steps drawn from
1094 $\mathcal{N}(-0.2d, 1)$, where d is the current amount of shift. Lastly, we add a $\mathcal{N}(0, 0.1^2)$ Gaussian noise
1095 to the entire simulated data. The activity of neurons are multiplied by a scalar before combining
1096 with the background activity and noise. We call this scalar 'signal level'.

1097 To validate the accuracy of Minian output, we simulate data with different signal level and
1098 number of cells. The signal levels we use are 0.2, 0.4, 0.6, 1.0, 1.4, 1.8. The number of cells we
1099 use are 100, 300, 500. On the other hand, to benchmark the performance of Minian, we
1100 simulate data with different number of frames and cells. The number of frames vary from 4000
1101 to 28000 with a step size of 8000. The number of cells we use are 100, 300, 500.

1102 *Matching neurons for validation*

1103 To compute different metrics of the accuracy of Minian output, we first need to match the
1104 putative neurons from Minian output with neurons from ground truth. To obtain this mapping we
1105 first compute the max projection of spatial footprints across all neurons. We then register the
1106 max projection of putative spatial footprints to the max projection of ground truth spatial
1107 footprints, by estimating a translational shift between the two max projection images. After
1108 correcting for translational shifts, we compute the center-of-mass for all neurons, from which we

1109 obtain a $N \times M$ pairwise distance matrix, where N and M are number of neurons detected by
1110 Minian and number of ground truth neurons, respectively. We then calculate an optimal
1111 mapping by solving the linear assignment problem of minimizing the total cost (distance) of a
1112 particular cell mapping. Lastly, we threshold the resulting mapping by discarding any matched
1113 cells that has a distance larger than 15 pixels.

1114 *Classification of place cells*

1115 We use the spatially-binned averaged ‘firing’ rate calculated from spike signals to classify
1116 whether each cell is a place cell. A place cell must simultaneously satisfy three criteria: a spatial
1117 information criterion, a stability criterion, and a place field size criterion. To determine whether a
1118 cell has significant spatial information or stability, we obtain a null distribution of the
1119 measurements (spatial information and stability) with a bootstrap strategy, where we roll the
1120 timing of activity by a random amount for each cell 1000 times. The observed spatial information
1121 or stability is defined as significant if it exceeds the 95th percentile of its null distribution ($p <$
1122 0.05). For the spatial information criterion, we use the joint information between ‘firing’ rate and
1123 an animal’s location measured in bits per ‘spike’. For the stability criterion, we calculate the
1124 Fisher z -transformation of the Pearson correlation coefficient between spatial ‘firing’ patterns
1125 across different trials within a recording session. A trial is defined as the time which the animal
1126 runs from one end of the linear track to the other and returns to the starting location. We
1127 calculate the z -transformed correlation between the odd number of trials and the even number
1128 of trials, as well as between the first half of the trials and the second half of the trials. We then
1129 average these two measures of correlations and use that as the measure of stability for a cell.
1130 Lastly, For the place field size criterion, we define the place field of each cell as the longest
1131 contiguous spatial bin where the averaged ‘firing’ rate exceeded the 95th percentile of all
1132 averaged firing rate bins. A cell must have a place field larger than 4 cm (i.e., 2 spatial bins) to
1133 pass the place field size criterion.

1134 *Animals*

1135 Adult male C57/BL6J mice from Jackson Laboratories were used for all testing. Animals were
1136 housed in a temperature, humidity and light controlled vivarium down the hall from the
1137 experimental testing rooms with lights on at 7 a.m. and off at 7 p.m. Water was restricted to
1138 maintain a body weight of 85–90%. Water deprivation consisted of allotting the animal ~1 mL of
1139 water per day, including water obtained during testing. Water not obtained during testing was
1140 given after the testing period. Animals were acclimated to handling for 5–7 days prior to
1141 training/testing. All experiments were performed in accordance with relevant guidelines and
1142 regulations approved by the Institutional Animal Care and Use Committee of Icahn School of
1143 Medicine at Mount Sinai (Reference #: IACUC-2017-0361, Protocol #: 17-1994).

1144 *Tested hardware specifications*

1145 The hardware specifications of computers that have effectively run Minian are summarized in
1146 the table below.

1147 *Table 1: A list of computers tested with Minian with specifications. Listed roughly by*
 1148 *increasing computation power.*

Manufacture	Model	CPU	RAM	Storage	Operating System
custom-built	Carbon	AMD Ryzen Threadripper 2950X 4.4GHz x 16	128GB	2TB SSD	Ubuntu 18.04
Microsoft	Surface Pro 6	Intel Core i5-8250U 1.6GHz x 4	8GB	256GB SSD	Windows 10
Dell	Precision 5530	Intel Core i5-8400H 2.5GHz x 4	16GB	256GB SSD	Ubuntu 18.04
Apple	MacBook Pro 152	Intel Core i7-8559U 2.7GHz x 4	16GB	1TB SSD	macOS 10.14 Mojave
custom-built	Amethyst	Intel Xeon E5-1650 3.6GHz x 6	128GB	6TB HDD	Ubuntu 17.1

1149 *List of dependencies*

1150 *Table 2: A list of open-source packages and the specific versions on which Minian*
 1151 *depends.*

Package	Version
av	7.0
bokeh	1.4
bottleneck	1.3
cairo	1.16
cvxpy	1.0
dask	2.11
datashader	0.1
distributed	2.11
ecos	2.0
ffmpeg	4.1
fftw	3.3
holoviews	1.12
ipython	7.12
ipywidgets	7.5
jupyter	1.0
matplotlib	3.1
natsort	7.0
netcdf4	1.5
networkx	2.4
nodejs	13.9
numba	0.48

numpy	1.18
opencv	4.2
pandas	1.0
panel	0.8
papermill	2.0
param	1.9
pip	20.0
pyfftw	0.12
python	3.8
scipy	1.4
scs	2.1
statsmodels	0.11
tiff file	2020.2
tqdm	4.43
xarray	0.15
zarr	2.4
medpy	0.4
simpleitk	1.2

1152 *Comparison of algorithms in related pipelines*

1153 *Table 3: List of algorithm implementations in different pipelines. For a lot of steps different*
 1154 *algorithm implementation can be chosen by the user based on features of the data. In such*
 1155 *cases we only list the default and most commonly used algorithms here.*

Step	Minian implementation	CalmAn implementation	MIN1PIPE implementation	Critical parameters
Denoising	Median filter	None	Anisotropic filter	Spatial window size of the filter
Background removal	Morphological top-hat transform	None	Morphological top-hat transform	Spatial window size of the top-hat transform
Motion correction	FFT-based translational motion correction	Non-rigid patch-wise translational motion correction (NoRMCorre)	Mix of translational motion correction and Demons diffeomorphic motion correction	Different
Initialization	Seed-based with peak-noise-ratio and KS-test refinement	Pixel-wise correlation and peak-noise-ratio thresholding	Seed-based with GMM, peak-noise-ratio and KS-test refinement	Threshold for correlation and peak-noise-ratio
Spatial and temporal updates	CNMF with cvxpy as deconvolution backend	CNMF-E with oasis as deconvolution backend	CNMF with cvx matlab package as deconvolution	Noise cut-off frequency. Expected size of

			backend	neurons. Sparse penalty
--	--	--	---------	-------------------------

1156 *Source data*

1157 **Table 4: List of source data related to validation figures.**

Title	Description
Figure 15 - source data 1	Raw validation performance with simulated data.
Figure 16 - source data 1	Raw correlations between Minian deconvolved traces and simulated ground truth.
Figure 16 - source data 2	Raw example traces from Minian and simulated ground truth. Filenames indicate signal level and source of trace.
Figure 17 - source data 1	Raw spatial footprint values shown in the overlay plot.
Figure 17 - source data 2	Raw example traces from Minian and Caiman. Filenames indicate cell id and source of trace.
Figure 18 - source data 1	Raw memory usage and running time with different datasets for both pipelines.
Figure 19 - source data 1	Raw memory usage and running time with different parallel processes for both pipelines.
Figure 20 - source data 1	Raw correlation of spatial firing pattern with different shifts in field-of-view.
Figure 20 - source data 2	Raw spatial firing activity for the two sessions shown.

1158

1159 **Conflict of interest**

1160 The authors declare that they have no competing financial interests.

1161 **Funding**

1162 *NIH F32AG067640*

1163 • William Mau

1164 *NIH BRAIN Initiative (R01 EB028166)*

1165 • Kanaka Rajan

1166 *James S. McDonnell Foundation's Understanding Human Cognition Scholar Award*

1167 • Kanaka Rajan

1168 *NSF FOUNDATIONS Award (NSF1926800)*

1169 • Kanaka Rajan

1170 *CURE Epilepsy Taking Flight Award*

1171 • Tristan Shuman

1172 *American Epilepsy Society Junior investigator Award*

1173 • Tristan Shuman

1174 *R03 NS111493*

1175 • Tristan Shuman

1176 *R21 DA049568*

1177 • Tristan Shuman

1178 *R01 NS116357*

1179 • Tristan Shuman

1180 *U01 NS094286-01*

1181 • Daniel Aharoni

1182 *1700408 Neurotech Hub*

1183 • Daniel Aharoni

1184 *NIH DP2MH122399*

1185 • Denise J. Cai

1186 *R01 MH120162*

1187 • Denise J. Cai

1188 *Botanical Center Pilot Award from P50 AT008661-01 from the NCCIH and the ODS (Pasinetti PI)*

1190 • Denise J. Cai

1191 *One Mind Otsuka Rising Star Award*

1192 • Denise J. Cai

1193 *McKnight Memory and Cognitive Disorders Award*

1194 • Denise J. Cai

1195 *Klingenstein-Simons Fellowship Award in Neuroscience*

1196 • Denise J. Cai

1197 *Mount Sinai Distinguished Scholar Award*

1198 • Denise J. Cai

1199 *Brain Research Foundation Award*

1200 • Denise J. Cai

1201 *NARSAD Young Investigator Award*

1202 • Denise J. Cai

1203 **Acknowledgements**

1204 We thank Eftychios A Pnevmatikakis, Andrea Giovannucci, and Liam Paninski for establishing
1205 the theoretical foundation and providing helpful insights for the pipeline. We thank Pat Gunn for
1206 helping with benchmarking with CalmAn pipeline. We thank Taylor Francisco and Denisse
1207 Morales-Rodriguez for helping with data analysis and revision. We thank MetaCell (Stephen
1208 Larson, Giovanni Idili, Zoran Sinnema, Dan Knudsen, and Paolo Bazzigaluppi) for contributing
1209 to the documentation and continuous integration of the pipeline. We thank Brandon Wei, Mimi
1210 La-Vu, and Christopher Lee for contributing to the dataset used in Minian development and
1211 testing. The authors acknowledge support from following funding sources: WM is supported by
1212 NIH F32AG067640. KR is supported by NIH BRAIN Initiative (R01 EB028166), James S.
1213 McDonnell Foundation's Understanding Human Cognition Scholar Award, and NSF

1214 FOUNDATIONS Award (NSF1926800). TS is supported by CURE Epilepsy Taking Flight
1215 Award, American Epilepsy Society Junior investigator Award, R03 NS111493, R21 DA049568,
1216 and R01 NS116357. DA is supported by U01 NS094286-01, and 1700408 Neurotech Hub. DJC
1217 is supported by NIH DP2MH122399, R01 MH120162, Botanical Center Pilot Award from P50
1218 AT008661-01 from the NCCIH and the ODS (Pasinetti PI), One Mind Otsuka Rising Star Award,
1219 McKnight Memory and Cognitive Disorders Award, Klingenstein-Simons Fellowship Award in
1220 Neuroscience, Mount Sinai Distinguished Scholar Award, Brain Research Foundation Award,
1221 and NARSAD Young Investigator Award.

1222 **References**

1223 **1. The Future Is Open: Open-Source Tools for Behavioral Neuroscience Research**

1224 Samantha R. White, Linda M. Amarante, Alexxai V. Kravitz, Mark Laubach

1225 *eneuro* (2019-07) <https://doi.org/ggcmcv>

1226 DOI: 10.1523/eneuro.0223-19.2019 · PMID: 31358510 · PMCID: PMC6712209

1227 **2. Open source tools for large-scale neuroscience**

1228 Jeremy Freeman

1229 *Current Opinion in Neurobiology* (2015-06) <https://doi.org/ghqn37>

1230 DOI: 10.1016/j.conb.2015.04.002 · PMID: 25982977

1231 **3. Open source modules for tracking animal behavior and closed-loop stimulation based 1232 on Open Ephys and Bonsai**

1233 Alessio Paolo Buccino, Mikkel Elle Lepperød, Sverre-Arne Dragly, Philipp Häfliger, Marianne
1234 Fyhn, Torkel Hafting

1235 *Journal of Neural Engineering* (2018-10-01) <https://doi.org/ggp3mh>

1236 DOI: 10.1088/1741-2552/aac45 · PMID: 29946057

1237 **4. An open source automated two-bottle choice test apparatus for rats**

1238 Jude A. Frie, Jibrán Y. Khokhar

1239 *HardwareX* (2019-04) <https://doi.org/ghr3cd>

1240 DOI: 10.1016/j.ohx.2019.e00061 · PMID: 31245655 · PMCID: PMC6594565

1241 **5. Bonsai: an event-based framework for processing and controlling data streams**

1242 Gonçalo Lopes, Niccolò Bonacchi, João Frazão, Joana P. Neto, Bassam V. Atallah, Sofia
1243 Soares, Luísa Moreira, Sara Matias, Pavel M. Itskov, Patrícia A. Correia, ... Adam R. Kampff

1244 *Frontiers in Neuroinformatics* (2015-04-08) <https://doi.org/ggbj87>

1245 DOI: 10.3389/fninf.2015.00007 · PMID: 25904861 · PMCID: PMC4389726

1246 **6. Feeding Experimentation Device (FED): A flexible open-source device for measuring 1247 feeding behavior**

1248 Katrina P. Nguyen, Timothy J. O'Neal, Olurotimi A. Bolunduro, Elecia White, Alexxai V. Kravitz

1249 *Journal of Neuroscience Methods* (2016-07) <https://doi.org/f8rcmm>

1250 DOI: 10.1016/j.jneumeth.2016.04.003 · PMID: 27060385 · PMCID: PMC4884551

1251 **7. An open-source device for measuring food intake and operant behavior in rodent 1252 home-cages**

- 1253 Bridget A Matikainen-Ankney, Thomas Earnest, Mohamed Ali, Eric Casey, Justin G Wang, Amy
1254 K Sutton, Alex A Legaria, Kia M Barclay, Laura B Murdaugh, Makenzie R Norris, ... Alexxai V
1255 Kravitz
1256 *eLife* (2021-03-29) <https://doi.org/gj6mqj>
1257 DOI: 10.7554/elife.66173 · PMID: 33779547 · PMCID: PMC8075584
- 1258 **8. JAABA: interactive machine learning for automatic annotation of animal behavior**
1259 Mayank Kabra, Alice A Robie, Marta Rivera-Alba, Steven Branson, Kristin Branson
1260 *Nature Methods* (2013-01) <https://doi.org/gg66kh>
1261 DOI: 10.1038/nmeth.2281 · PMID: 23202433
- 1262 **9. DeepLabCut: markerless pose estimation of user-defined body parts with deep**
1263 **learning**
1264 Alexander Mathis, Pranav Mamidanna, Kevin M. Cury, Taiga Abe, Venkatesh N. Murthy,
1265 Mackenzie Weygandt Mathis, Matthias Bethge
1266 *Nature Neuroscience* (2018-09) <https://doi.org/gd249k>
1267 DOI: 10.1038/s41593-018-0209-y · PMID: 30127430
- 1268 **10. Automated classification of self-grooming in mice using open-source software**
1269 Bastijn J. G. van den Boom, Pavlina Pavlidi, Casper J. H. Wolf, Adriana H. Mooij, Ingo Willuhn
1270 *Journal of Neuroscience Methods* (2017-09) <https://doi.org/gb2wxk>
1271 DOI: 10.1016/j.jneumeth.2017.05.026 · PMID: 28648717
- 1272 **11. Efficient and accurate extraction of in vivo calcium signals from microendoscopic**
1273 **video data**
1274 Pengcheng Zhou, Shanna L Resendez, Jose Rodriguez-Romaguera, Jessica C Jimenez, Shay
1275 Q Neufeld, Andrea Giovannucci, Johannes Friedrich, Eftychios A Pnevmatikakis, Garret D
1276 Stuber, Rene Hen, ... Liam Paninski
1277 *eLife* (2018-02-22) <https://doi.org/gfxbdp>
1278 DOI: 10.7554/elife.28728 · PMID: 29469809 · PMCID: PMC5871355
- 1279 **12. MIN1PIPE: A Miniscope 1-Photon-Based Calcium Imaging Signal Extraction Pipeline**
1280 Jinghao Lu, Chunyuan Li, Jonnathan Singh-Alvarado, Zhe Charles Zhou, Flavio Fröhlich,
1281 Richard Mooney, Fan Wang
1282 *Cell Reports* (2018-06) <https://doi.org/gdpc2z>
1283 DOI: 10.1016/j.celrep.2018.05.062 · PMID: 29925007 · PMCID: PMC6084484
- 1284 **13. Automated Analysis of Cellular Signals from Large-Scale Calcium Imaging Data**
1285 Eran A. Mukamel, Axel Nimmerjahn, Mark J. Schnitzer
1286 *Neuron* (2009-09) <https://doi.org/bhwqvc>
1287 DOI: 10.1016/j.neuron.2009.08.009 · PMID: 19778505 · PMCID: PMC3282191
- 1288 **14. Suite2p: beyond 10,000 neurons with standard two-photon microscopy**
1289 Marius Pachitariu, Carsen Stringer, Mario Dipoppa, Sylvia Schröder, L. Federico Rossi, Henry
1290 Dalgleish, Matteo Carandini, Kenneth D. Harris
1291 *Cold Spring Harbor Laboratory* (2017-07-20) <https://doi.org/ggdxxm>
1292 DOI: 10.1101/061507

- 1293 **15. Fast, Simple Calcium Imaging Segmentation with Fully Convolutional Networks**
1294 Aleksander Klibisz, Derek Rose, Matthew Eicholtz, Jay Blundon, Stanislav Zakharenko
1295 *Lecture Notes in Computer Science* (2017) <https://doi.org/ghm58x>
1296 DOI: 10.1007/978-3-319-67558-9_33
- 1297 **16. CalmAn an open source tool for scalable calcium imaging data analysis**
1298 Andrea Giovannucci, Johannes Friedrich, Pat Gunn, Jérémie Kalfon, Brandon L Brown, Sue
1299 Ann Koay, Jiannis Taxidis, Farzaneh Najafi, Jeffrey L Gauthier, Pengcheng Zhou, ... Eftychios
1300 A Pnevmatikakis
1301 *eLife* (2019-01-17) <https://doi.org/gf4v82>
1302 DOI: 10.7554/elife.38173 · PMID: 30652683 · PMCID: PMC6342523
- 1303 **17. Tracking the Same Neurons across Multiple Days in Ca²⁺ Imaging Data**
1304 Liron Sheintuch, Alon Rubin, Noa Brande-Eilat, Nitzan Geva, Noa Sadeh, Or Pinchasof, Yaniv
1305 Ziv
1306 *Cell Reports* (2017-10) <https://doi.org/ghdnqz>
1307 DOI: 10.1016/j.celrep.2017.10.013 · PMID: 29069591 · PMCID: PMC5670033
- 1308 **18. ezTrack: An open-source video analysis pipeline for the investigation of animal**
1309 **behavior**
1310 Zachary T. Pennington, Zhe Dong, Yu Feng, Lauren M. Vetere, Lucia Page-Harley, Tristan
1311 Shuman, Denise J. Cai
1312 *Scientific Reports* (2019-12-27) <https://doi.org/ghm6dp>
1313 DOI: 10.1038/s41598-019-56408-9 · PMID: 31882950 · PMCID: PMC6934800
- 1314 **19. Fast online deconvolution of calcium imaging data**
1315 Johannes Friedrich, Pengcheng Zhou, Liam Paninski
1316 *PLOS Computational Biology* (2017-03-14) <https://doi.org/f9tsn9>
1317 DOI: 10.1371/journal.pcbi.1005423 · PMID: 28291787 · PMCID: PMC5370160
- 1318 **20. OnACID: Online Analysis of Calcium Imaging Data in Real Time***
1319 Andrea Giovannucci, Johannes Friedrich, Matt Kaufman, Anne Churchland, Dmitri Chklovskii,
1320 Liam Paninski, Eftychios A. Pnevmatikakis
1321 *Cold Spring Harbor Laboratory* (2017-10-02) <https://doi.org/ghqn38>
1322 DOI: 10.1101/193383
- 1323 **21. Exact spike train inference via ℓ_0 optimization**
1324 Sean Jewell, Daniela Witten
1325 *The Annals of Applied Statistics* (2018-12-01) <https://doi.org/ghm589>
1326 DOI: 10.1214/18-aos1162 · PMID: 30627301 · PMCID: PMC6322847
- 1327 **22. All the light that we can see: a new era in miniaturized microscopy**
1328 Daniel Aharoni, Baljit S. Khakh, Alcino J. Silva, Peyman Golshani
1329 *Nature Methods* (2019-01) <https://doi.org/ghdvnz>
1330 DOI: 10.1038/s41592-018-0266-x · PMID: 30573833 · PMCID: PMC8320687

- 1331 **23. An open-source control system for in vivo fluorescence measurements from deep-**
1332 **brain structures**
1333 Scott F. Owen, Anatol C. Kreitzer
1334 *Journal of Neuroscience Methods* (2019-01) <https://doi.org/ghvk8p>
1335 DOI: 10.1016/j.jneumeth.2018.10.022 · PMID: 30342106 · PMCID: PMC6258340
- 1336 **24. Open Ephys: an open-source, plugin-based platform for multichannel**
1337 **electrophysiology**
1338 Joshua H Siegle, Aarón Cuevas López, Yogi A Patel, Kirill Abramov, Shay Ohayon, Jakob
1339 Voigts
1340 *Journal of Neural Engineering* (2017-08-01) <https://doi.org/gfvmzq>
1341 DOI: 10.1088/1741-2552/aa5eea · PMID: 28169219
- 1342 **25. Open Source Tools for Temporally Controlled Rodent Behavior Suitable for**
1343 **Electrophysiology and Optogenetic Manipulations**
1344 Nicola Solari, Katalin Sviatkó, Tamás Laszlovszky, Panna Hegedüs, Balázs Hangya
1345 *Frontiers in Systems Neuroscience* (2018-05-15) <https://doi.org/gdns24>
1346 DOI: 10.3389/fnsys.2018.00018 · PMID: 29867383 · PMCID: PMC5962774
- 1347 **26. A wireless miniScope for deep brain imaging in freely moving mice**
1348 Giovanni Barbera, Bo Liang, Lifeng Zhang, Yun Li, Da-Ting Lin
1349 *Journal of Neuroscience Methods* (2019-07) <https://doi.org/ghtkfs>
1350 DOI: 10.1016/j.jneumeth.2019.05.008 · PMID: 31116963 · PMCID: PMC6636826
- 1351 **27. A Compact Head-Mounted Endoscope for In Vivo Calcium Imaging in Freely Behaving**
1352 **Mice**
1353 Alexander D. Jacob, Adam I. Ramsaran, Andrew J. Mocle, Lina M. Tran, Chen Yan, Paul W.
1354 Frankland, Sheena A. Josselyn
1355 *Current Protocols in Neuroscience* (2018-07) <https://doi.org/gdr76d>
1356 DOI: 10.1002/cpns.51 · PMID: 29944206
- 1357 **28. An open source, wireless capable miniature microscope system**
1358 William A Liberti, L Nathan Perkins, Daniel P Leman, Timothy J Gardner
1359 *Journal of Neural Engineering* (2017-08-01) <https://doi.org/gf73sj>
1360 DOI: 10.1088/1741-2552/aa6806 · PMID: 28514229 · PMCID: PMC5955387
- 1361 **29. NINscope, a versatile miniscope for multi-region circuit investigations**
1362 Andres de Groot, Bastijn JG van den Boom, Romano M van Genderen, Joris Coppens, John
1363 van Veldhuijzen, Joop Bos, Hugo Hoedemaker, Mario Negrello, Ingo Willuhn, Chris I De Zeeuw,
1364 Tycho M Hoogland
1365 *eLife* (2020-01-14) <https://doi.org/ghsb8m>
1366 DOI: 10.7554/elife.49987 · PMID: 31934857 · PMCID: PMC6989121
- 1367 **30. High-speed volumetric imaging of neuronal activity in freely moving rodents**
1368 Oliver Skocek, Tobias Nöbauer, Lukas Weillguny, Francisca Martínez Traub, Chuying Naomi
1369 Xia, Maxim I. Molodtsov, Abhinav Grama, Masahito Yamagata, Daniel Aharoni, David D. Cox,
1370 ... Alipasha Vaziri

- 1371 *Nature Methods* (2018-06) <https://doi.org/gf2n7z>
1372 DOI: 10.1038/s41592-018-0008-0 · PMID: 29736000 · PMCID: PMC7990085
- 1373 **31. Imaging Cortical Dynamics in GCaMP Transgenic Rats with a Head-Mounted**
1374 **Widefield Macroscope**
1375 Benjamin B. Scott, Stephan Y. Thiberge, Caiying Guo, D. Gowanlock R. Tervo, Carlos D. Brody,
1376 Alla Y. Karpova, David W. Tank
1377 *Neuron* (2018-12) <https://doi.org/gfgk25>
1378 DOI: 10.1016/j.neuron.2018.09.050 · PMID: 30482694 · PMCID: PMC6283673
- 1379 **32. Miniaturized integration of a fluorescence microscope**
1380 Kunal K Ghosh, Laurie D Burns, Eric D Cocker, Axel Nimmerjahn, Yaniv Ziv, Abbas El Gamal,
1381 Mark J Schnitzer
1382 *Nature Methods* (2011-10) <https://doi.org/cv75qh>
1383 DOI: 10.1038/nmeth.1694 · PMID: 21909102 · PMCID: PMC3810311
- 1384 **33. Long-term dynamics of CA1 hippocampal place codes**
1385 Yaniv Ziv, Laurie D Burns, Eric D Cocker, Elizabeth O Hamel, Kunal K Ghosh, Lacey J Kitch,
1386 Abbas El Gamal, Mark J Schnitzer
1387 *Nature Neuroscience* (2013-02-10) <https://doi.org/gdh98h>
1388 DOI: 10.1038/nn.3329 · PMID: 23396101 · PMCID: PMC3784308
- 1389 **34. A shared neural ensemble links distinct contextual memories encoded close in time**
1390 Denise J. Cai, Daniel Aharoni, Tristan Shuman, Justin Shobe, Jeremy Biane, Weilin Song,
1391 Brandon Wei, Michael Veshkini, Mimi La-Vu, Jerry Lou, ... Alcino J. Silva
1392 *Nature* (2016-06) <https://doi.org/f8pp28>
1393 DOI: 10.1038/nature17955 · PMID: 27251287 · PMCID: PMC5063500
- 1394 **35. Breakdown of spatial coding and interneuron synchronization in epileptic mice**
1395 Tristan Shuman, Daniel Aharoni, Denise J. Cai, Christopher R. Lee, Spyridon Chavlis, Lucia
1396 Page-Harley, Lauren M. Vetere, Yu Feng, Chen Yi Yang, Irene Mollinedo-Gajate, ... Peyman
1397 Golshani
1398 *Nature Neuroscience* (2020-02) <https://doi.org/ghm6dn>
1399 DOI: 10.1038/s41593-019-0559-0 · PMID: 31907437 · PMCID: PMC7259114
- 1400 **36. Circuit Investigations With Open-Source Miniaturized Microscopes: Past, Present and**
1401 **Future**
1402 Daniel Aharoni, Tycho M. Hoogland
1403 *Frontiers in Cellular Neuroscience* (2019-04-05) <https://doi.org/ghqn39>
1404 DOI: 10.3389/fncel.2019.00141 · PMID: 31024265 · PMCID: PMC6461004
- 1405 **37. Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data**
1406 Eftychios A. Pnevmatikakis, Daniel Soudry, Yuanjun Gao, Timothy A. Machado, Josh Merel,
1407 David Pfau, Thomas Reardon, Yu Mu, Clay Lacefield, Weijian Yang, ... Liam Paninski
1408 *Neuron* (2016-01) <https://doi.org/f8g23x>
1409 DOI: 10.1016/j.neuron.2015.11.037 · PMID: 26774160 · PMCID: PMC4881387

- 1410 **38. Fast Nonnegative Deconvolution for Spike Train Inference From Population Calcium**
1411 **Imaging**
1412 Joshua T. Vogelstein, Adam M. Packer, Timothy A. Machado, Tanya Sippy, Baktash Babadi,
1413 Rafael Yuste, Liam Paninski
1414 *Journal of Neurophysiology* (2010-12) <https://doi.org/fpddqn>
1415 DOI: 10.1152/jn.01073.2009 · PMID: 20554834 · PMCID: PMC3007657
- 1416 **39. Fast and statistically robust cell extraction from large-scale neural calcium imaging**
1417 **datasets**
1418 Hakan Inan, Claudia Schmuckermair, Tugce Tasci, Biafra O. Ahanonu, Oscar Hernandez,
1419 Jérôme Lecoq, Fatih Dinç, Mark J. Wagner, Murat A. Erdogdu, Mark J. Schnitzer
1420 *Neuroscience* (2021-03-25) <https://doi.org/gjs4d5>
1421 DOI: 10.1101/2021.03.24.436279
- 1422 **40. Template matching techniques in computer vision: theory and practice**
1423 Roberto Brunelli
1424 *Wiley* (2009)
1425 ISBN: 9780470517062
- 1426 **41. Online analysis of microendoscopic 1-photon calcium imaging data streams**
1427 Johannes Friedrich, Andrea Giovannucci, Eftychios A. Pnevmatikakis
1428 *PLOS Computational Biology* (2021-01-28) <https://doi.org/gp2fsj>
1429 DOI: 10.1371/journal.pcbi.1008565 · PMID: 33507937 · PMCID: PMC7842953
- 1430 **42. Jupyter Notebooks – a publishing format for reproducible computational workflows**
1431 Thomas Kluyver, Benjamin Ragan-Kelley, Pé, Fernando Rez, Brian Granger, Matthias
1432 Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica Hamrick, Jason Grout, ... Jupyter
1433 Development Team
1434 *Positioning and Power in Academic Publishing: Players, Agents and Agendas* (2016)
1435 <https://ebooks.iospress.nl/doi/10.3233/978-1-61499-649-1-87>
1436 DOI: 10.3233/978-1-61499-649-1-87
- 1437 **43. Array programming with NumPy**
1438 Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen,
1439 David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, ... Travis E.
1440 Oliphant
1441 *Nature* (2020-09-17) <https://doi.org/ghbzf2>
1442 DOI: 10.1038/s41586-020-2649-2 · PMID: 32939066 · PMCID: PMC7759461
- 1443 **44. SciPy 1.0: fundamental algorithms for scientific computing in Python**
1444 Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David
1445 Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, ... Yoshiki
1446 Vázquez-Baeza
1447 *Nature Methods* (2020-03-02) <https://doi.org/ggj45f>
1448 DOI: 10.1038/s41592-019-0686-2 · PMID: 32015543 · PMCID: PMC7056644

- 1449 **45. xarray: N-D labeled Arrays and Datasets in Python**
1450 Stephan Hoyer, Joseph J. Hamman
1451 *Journal of Open Research Software* (2017-04-05) <https://doi.org/gdqdmw>
1452 DOI: 10.5334/jors.148
- 1453 **46. holoviz/holoviews: Version 1.13.3**
1454 Philipp Rudiger, Jean-Luc Stevens, James A. Bednar, Bas Nijholt,, Andrew, Chris B, Achim
1455 Randelhoff, Jon Mease, Vasco Tenner, Maxalbert, ... Kbowen
1456 *Zenodo* (2020-06-23) <https://doi.org/ghm6dq>
1457 DOI: 10.5281/zenodo.3904606
- 1458 **47. Bokeh: Python library for interactive visualization**
1459 Bokeh Development Team
1460 (2020) <https://bokeh.org/>
- 1461 **48. The OpenCV Library**
1462 G. Bradski
1463 *Dr. Dobb's Journal of Software Tools* (2000)
- 1464 **49. Dask: Library for dynamic task scheduling**
1465 Dask Development Team
1466 (2016) <https://dask.org>
- 1467 **50. The hippocampus as a spatial map. Preliminary evidence from unit activity in the**
1468 **freely-moving rat**
1469 J. O'Keefe, J. Dostrovsky
1470 *Brain Research* (1971-11) <https://doi.org/bwdqcb>
1471 DOI: 10.1016/0006-8993(71)90358-1
- 1472 **51. Long-term stability of the place-field activity of single units recorded from the dorsal**
1473 **hippocampus of freely behaving rats**
1474 L. T. Thompson, P. J. Best
1475 *Brain Research* (1990-02) <https://doi.org/cp6bjf>
1476 DOI: 10.1016/0006-8993(90)90555-p
- 1477 **52. Robustness of Spike Deconvolution for Neuronal Calcium Imaging**
1478 Marius Pachitariu, Carsen Stringer, Kenneth D. Harris
1479 *The Journal of Neuroscience* (2018-09-12) <https://doi.org/gd9mcx>
1480 DOI: 10.1523/jneurosci.3339-17.2018 · PMID: 30082416 · PMCID: PMC6136155
- 1481 **53. Ultrasensitive fluorescent proteins for imaging neuronal activity**
1482 Tsai-Wen Chen, Trevor J. Wardill, Yi Sun, Stefan R. Pulver, Sabine L. Renninger, Amy Baohan,
1483 Eric R. Schreiter, Rex A. Kerr, Michael B. Orger, Vivek Jayaraman, ... Douglas S. Kim
1484 *Nature* (2013-07) <https://doi.org/gcz68k>
1485 DOI: 10.1038/nature12354 · PMID: 23868258 · PMCID: PMC3777791