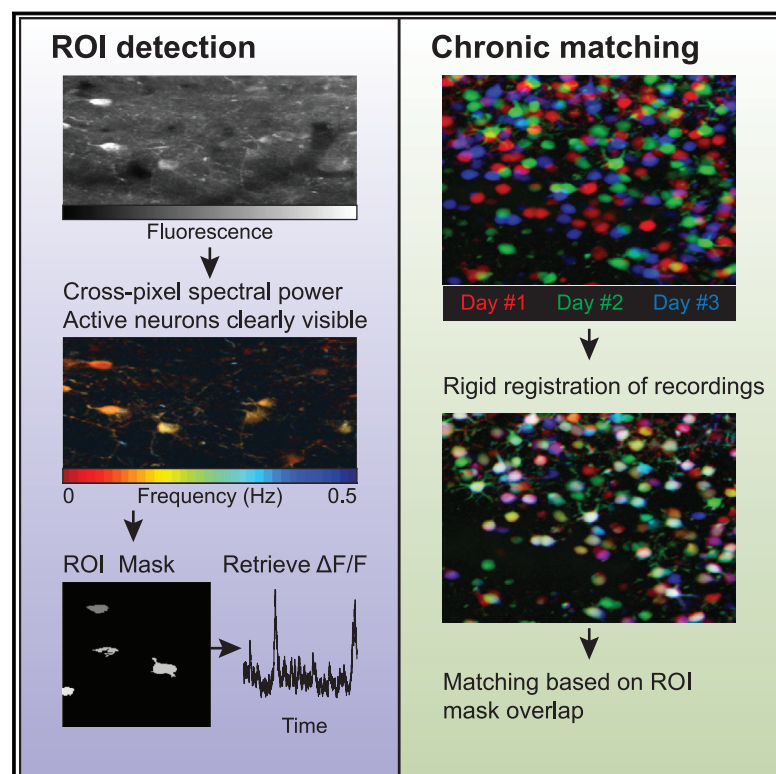


SpecSeg is a versatile toolbox that segments neurons and neurites in chronic calcium imaging datasets based on low-frequency cross-spectral power

Graphical abstract



Authors

Leander de Kraker, Koen Seignette, Premnath Thamizharasu, ..., Ingo Willuhn, Christiaan N. Levelt, Chris van der Togt

Correspondence

c.levelt@nin.knaw.nl

In brief

de Kraker et al. present SpecSeg, a user-friendly and versatile open-source toolbox for calcium imaging data acquired with miniscopes or multi-photon microscopes. It selects regions of interest (ROIs) around neurons and neurites, performs neuropil correction and signal extraction, and matches ROIs in sequential recordings.

Highlights

- SpecSeg finds ROIs and extracts signals in Ca^{2+} imaging data
- Neurons and neurites are identified by cross-correlations of low-frequency elements
- SpecSeg works intuitively for single- (miniscope) and multi-photon microscopy data
- SpecSeg can identify the same ROIs through entire series of chronic datasets

Article

SpecSeg is a versatile toolbox that segments neurons and neurites in chronic calcium imaging datasets based on low-frequency cross-spectral power

Leander de Kraker,¹ Koen Seignette,^{1,5} Premnath Thamizharasu,^{1,5} Bastijn J.G. van den Boom,^{3,4} Ildfonso Ferreira Pica,¹ Ingo Willuhn,^{3,4} Christiaan N. Levelt,^{1,2,6,7,*} and Chris van der Togt^{1,6}

¹Netherlands Institute for Neuroscience, Molecular Visual Plasticity Group, Royal Netherlands Academy of Arts and Sciences, Meibergdreef 47, 1105 BA Amsterdam, the Netherlands

²Department of Molecular and Cellular Neurobiology, Center for Neurogenomics and Cognitive Research, VU University Amsterdam, de Boelelaan 1085, 1081 HV Amsterdam, the Netherlands

³Netherlands Institute for Neuroscience, Neuromodulation and Behavior Group, Royal Netherlands Academy of Arts and Sciences, Meibergdreef 47, 1105 BA Amsterdam, the Netherlands

⁴Department of Psychiatry, Amsterdam UMC, University of Amsterdam, Meibergdreef 5, 1105 AZ Amsterdam, the Netherlands

⁵These authors contributed equally

⁶These authors contributed equally

⁷Lead contact

*Correspondence: c.levelt@nin.knaw.nl

<https://doi.org/10.1016/j.crmeth.2022.100299>

MOTIVATION Currently available pipelines for the analysis of calcium imaging data are of high quality. However, most of them have limitations when it comes to selecting regions of interest (ROIs) around dendrites or axons, use settings and parameters that are not intuitive for the end user and do not enable immediate and extensive quality control of the ROIs. In addition, different packages are required for analyzing one-photon and multi-photon data. No pipelines enable the identification of the same ROIs through entire series of chronically recorded calcium imaging datasets. Because these are all important features necessary for the in-depth analyses of calcium imaging data, we developed SpecSeg, a versatile and user-friendly pipeline for calcium data analysis that encompasses all the features described above.

SUMMARY

Imaging calcium signals in neurons of animals using single- or multi-photon microscopy facilitates the study of coding in large neural populations. Such experiments produce massive datasets requiring powerful methods to extract responses from hundreds of neurons. We present SpecSeg, an open-source toolbox for (1) segmentation of regions of interest (ROIs) representing neuronal structures, (2) inspection and manual editing of ROIs, (3) neuropil correction and signal extraction, and (4) matching of ROIs in sequential recordings. ROI segmentation in SpecSeg is based on temporal cross-correlations of low-frequency components derived by Fourier analysis of each pixel with its neighbors. The approach is user-friendly, intuitive, and insightful and enables ROI detection around neurons or neurites. It works for single- (miniscope) and multi-photon microscopy data, eliminating the need for separate toolboxes. SpecSeg thus provides an efficient and versatile approach for analyzing calcium responses in neuronal structures imaged over prolonged periods of time.

INTRODUCTION

The advances in *in vivo* dual-photon fluorescence microscopy (Engert and Bonhoeffer, 1999; Svoboda et al., 1996) and the development of genetically encoded fluorescent biosensors (Chen et al., 2013; Mank et al., 2008) have revolutionized neurobiological research over the last two decades. The combination

of these techniques has enabled the imaging of neuronal activity in awake behaving animals over timespans up to many months. This provides combined anatomical and functional information at the cellular level for hundreds of neurons at the same time, or at the dendritic, axonal, or synaptic level in more restricted numbers of neurons (Cichon and Gan, 2015; Gambino et al., 2014; Iacaruso et al., 2017; Jaepel et al., 2017; Jia et al., 2010;



Petreaunu et al., 2009; Szalay et al., 2016; Wilson et al., 2016; Winubst et al., 2015). More recently, miniaturized single-photon fluorescence microscopes, or miniscopes, are further expanding the possibilities for analyzing neural activity during behavior in freely moving animals (Aharoni et al., 2019; Cai et al., 2016; Ghosh et al., 2011; de Groot et al., 2020; Liberti et al., 2017; Resendez et al., 2016).

The most-used approach is the imaging of changes in intracellular calcium levels as a proxy for neuronal activity. This is achieved by making use of genetically encoded calcium sensors, such as GCaMP6 (Chen et al., 2013), whose fluorescent properties change upon binding calcium. The mere size of the obtained datasets, which consist of movies of calcium-indicator fluorescence images, forms a considerable challenge for data analysis. An important step in the analysis of calcium imaging data is the identification of cell bodies or neurites in the image sequences. Ideally, one can identify the same structures in recordings performed at different days over prolonged periods of time, enabling the assessment of changes in neuronal responses during learning at the single-cell level. Identification of these regions of interest (ROIs) is preferentially done in an automated fashion, as manual segmentation is neither reproducible nor scalable. Moreover, human annotators tend to include non-active ROIs and miss active ROIs with low background fluorescence (Giovannucci et al., 2019; Pachitariu et al., 2017). Automated ROI identification requires robust detection algorithms with minimal assumptions on the properties of ROIs to detect the circumferences of individual cells and neurites.

Various software packages have been published that accomplish this task, using different methods. Cell boundaries may be detected by multiple coupled active contours (Reynolds et al., 2017). Matrix factorization approaches (Giovannucci et al., 2019; Maruyama et al., 2014; Petersen et al., 2018; Pnevmatikakis et al., 2016; Zhou et al., 2018) determine the activity of neurons and their delineation by considering fluorescence as a spatiotemporal pattern that can be expressed as the product of a matrix encoding location and a matrix encoding time. Deep learning approaches (Apthorpe et al., 2016; Klíbisz et al., 2017; Mukamel et al., 2009) define ROIs based on neuron features learned from data in which cells were manually detected. Dictionary learning approaches make use of neuron templates to identify ROIs (Pachitariu et al., 2017). Finally, correlation-based approaches define ROIs based on activity correlations between pixels (Kaifosh et al., 2014; Mishne et al., 2018; Smith and Häusser, 2010; Spaen et al., 2019). All these methods have their strengths but also some weaknesses. For example, deep learning approaches that select ROIs based on shapes must be trained for different types of data and may include neurons from which no signal can be extracted. Identifying ROIs with highly variable shapes is complicated for most approaches except those based on activity correlation. For most methods, it is necessary to introduce significant adaptations to the software to make it suitable for specific experimental settings, and the underlying calculations that define the ROIs can be difficult or impossible to track, making it difficult to select the optimal settings. Correlation-based approaches are easy to understand and require few predefined constraints but can be severely limited by noise in the recordings.

In this paper, we describe SpecSeg, an open-source pipeline (Figure 1) for calcium imaging data analysis. It makes use of an ROI detection process that is based on cross-correlations of low-frequency components, derived through Fourier transforms, of each pixel trace with its eight adjacent neighbors. It makes use of our finding that spiking neurons can be identified by low-frequency fluctuations below 0.4 Hz in the calcium signal. Our approach is insensitive to noise, straightforward, and highly insightful for the end user as it enables the visualization of the ROI detection process. It enables the detection of ROIs around irregular structures, such as dendrites and axons, and the parameters set by the users are intuitive (e.g., size, roundedness). The pipeline also includes a user interface for quality control and the manual splitting, addition, or rejection of ROIs, and a tool to match ROIs in sequential imaging sessions (Figure 1). Finally, the pipeline can be used for the analysis of data acquired using multi-photon microscopes and single-photon miniscopes. Together, this pipeline provides an efficient and user-friendly approach for analyzing calcium responses in neuronal structures imaged over prolonged periods of time.

RESULTS

Pipeline

Figure 1 shows the components of our pipeline for the analysis of chronic calcium imaging datasets. The pipeline requires calcium imaging data in SBX (NeuroLabware) format. MATLAB functions to convert series of TIFF or H5 files to SBX format are provided. We use an adapted version of NoRMCorre (Pnevmatikakis and Giovannucci, 2017) to motion correct SBX files.

The first step in the pipeline is ROI selection, which involves (1) reorganization and temporal downsampling of the data to speed up memory retrieval, (2) extraction of frequency components of fluorescence traces for each pixel (“pixel trace”) by Fourier analysis and the creation of images representing their correlation with those of neighboring pixels, (3) identification of peaks within the images and the construction of preliminary contours (ROIs), and (4) further ROI refinement based on correlations of raw fluorescence traces within each contour. The second step in the pipeline is the ROI manager, a user interface for ROI inspection that has tools for manually rejecting, splitting, or adding ROIs. The third step is neuropil subtraction and signal extraction, and the final, fourth, step is deconvolution of the signals, using maximum likelihood spike estimation (MLspike) (Deneux et al., 2016).

In addition, a separate toolbox is included for matching ROIs in sequential imaging sessions by aligning the images and measuring the overlap of ROIs between different sessions. The sensitivity of ROI matching can be changed easily with an overlap threshold and matching results can be evaluated and edited. Below, each step in the pipeline is described in detail.

The MATLAB implementation of the toolbox and instructions on how to install and use the software can be found at <https://github.com/Leveltlab/SpectralSegmentation>.

ROI selection based on cross-spectral power

We established an approach for the automated segmentation of ROIs in calcium imaging data based on the cross-spectral power of the pixel trace. To develop the method, we made

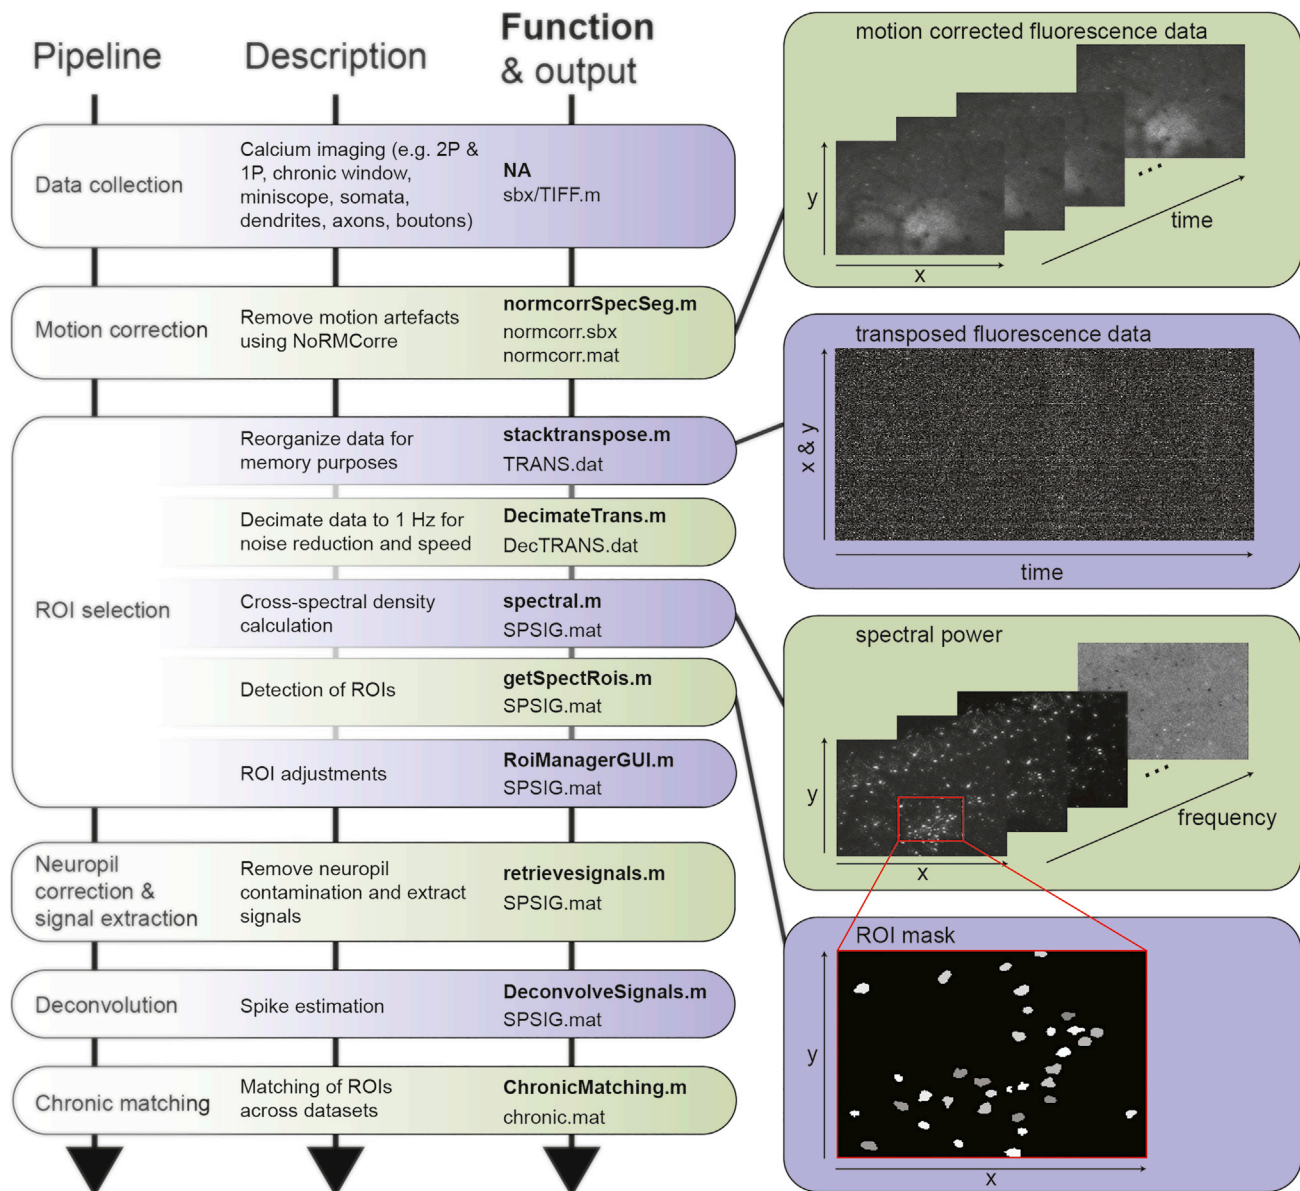


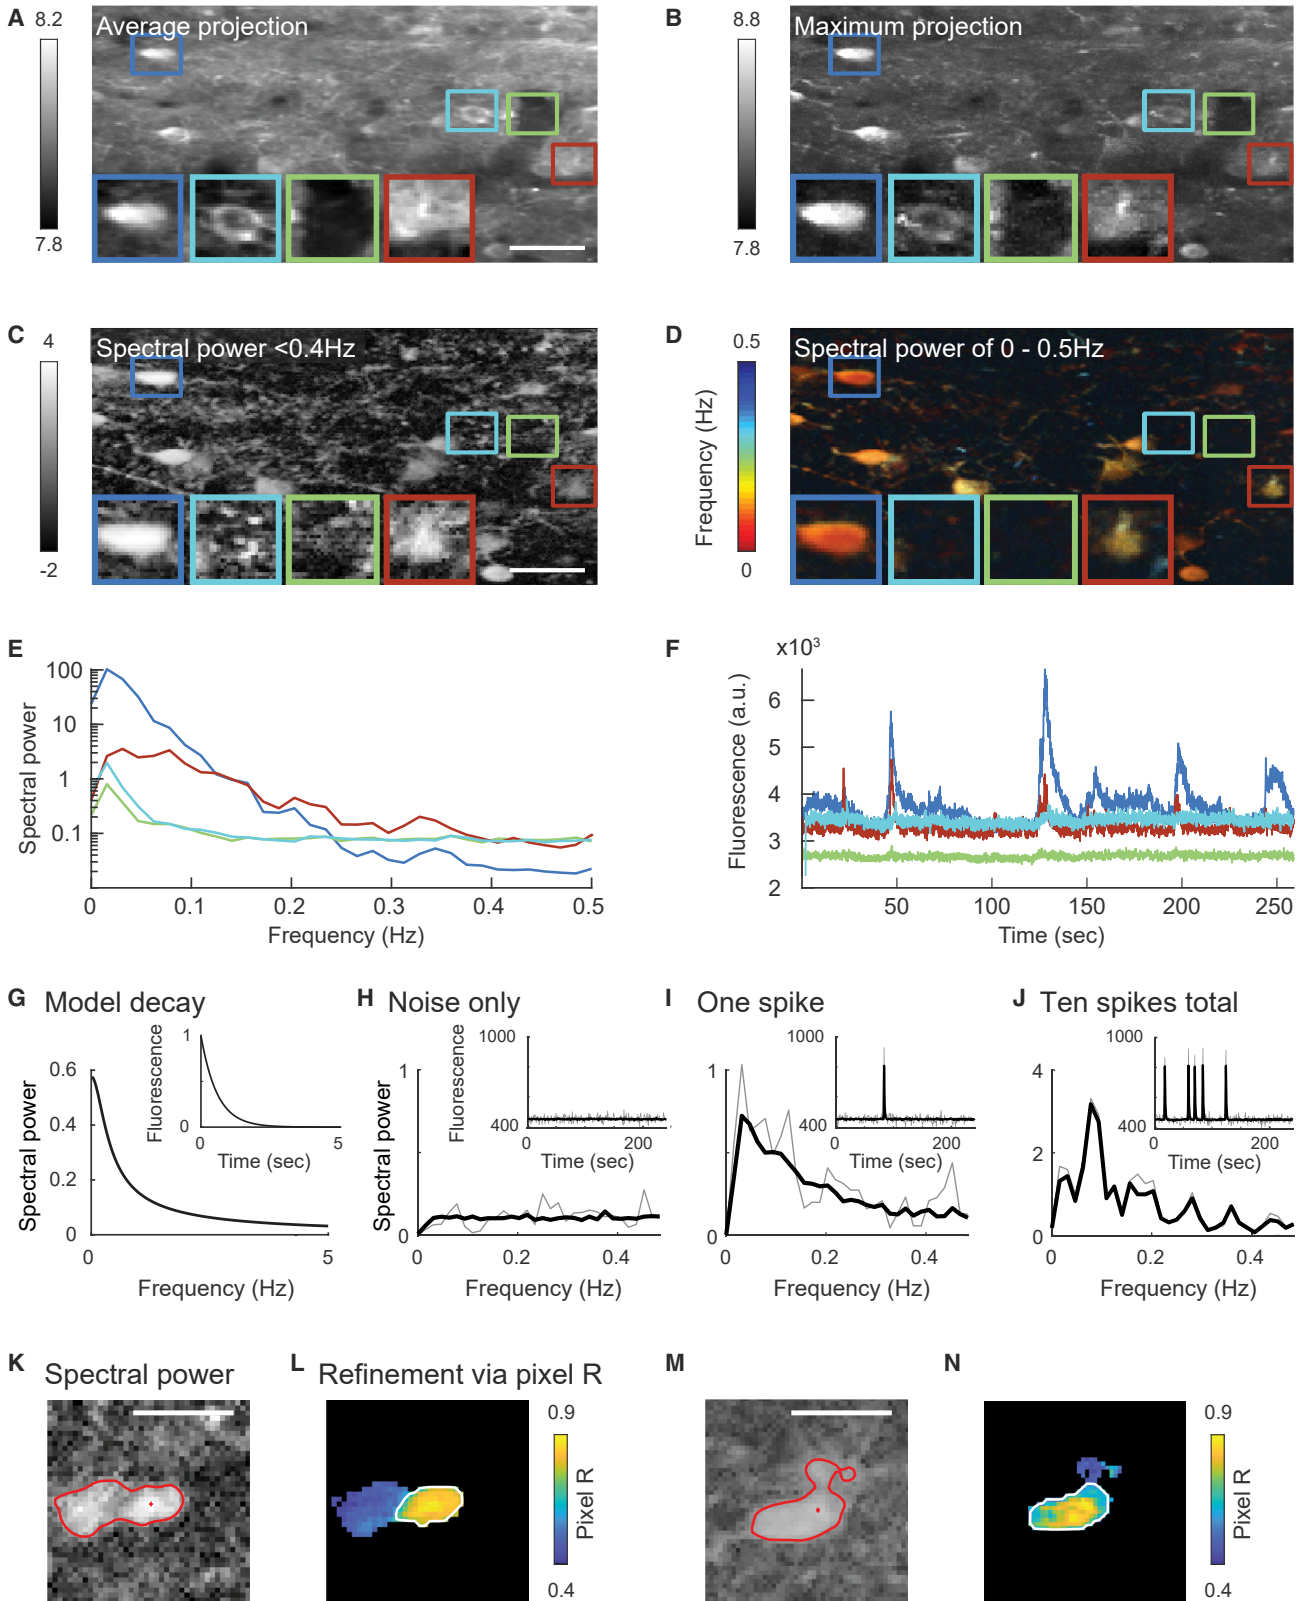
Figure 1. Overview of the pipeline for automated region of interest (ROI) selection and signal extraction

After collection, calcium imaging data need to be motion corrected using NoRMCorre. ROI selection involves data reorganization, extraction of frequency components of pixel traces, and the drawing and refinement of ROIs around the peaks identified in the images based on the frequency components. A user interface enables the manual rejection, splitting, or adding of ROIs. Next, the neuropil is subtracted and the signal extracted and deconvolved. A toolbox for matching of ROIs in repeated recordings is also provided.

use of calcium imaging data acquired repeatedly using two-photon microscopy in mice expressing the genetically encoded calcium indicator GCaMP6f in the primary visual cortex (V1).

The first step (`stacktranspose.m`) is to transpose the image sequences to place time in the first dimension and width x height in the second, and to down-sample the data in time to approximately 1 Hz (`DecimateTrans.m`), using the MATLAB `decimate` function (Mathworks). This way, each pixel trace is organized in sequential order in memory and can be accessed rapidly, speeding up the next steps in the process.

Next is the cross-spectral power calculation (`spectral.m`). Each pixel trace is cut into overlapping 1-min segments, and a discrete Fourier transformation is applied to each segment to extract frequency components between 0.013 and 0.5 Hz, with a bin width of 0.017 Hz. Then we calculate the cross-spectral density function of each pixel with its eight neighbors, and average these over all segments. Finally, the average cross-spectral density functions are normalized with the variances of each pixel and its neighbors, and the average cross-spectral power for each frequency component at each pixel is calculated



(legend on next page)

(this estimate is a measure for how well pixels correlate with their direct neighbors at each frequency component). This results in a series of 30 images, each representing a different frequency component.

These images of spectral components provide a better basis for ROI selection around active neurons than the fluorescent signal (Figures 2A and 2B). We find that spectral components below 0.4 Hz are the most indicative for active neural elements (Figures 2C–2F) (Chen et al., 2013). When we compose images by translating the cross-spectral power of a particular low-frequency component into brightness, active elements become clearly separated from the background (Figures 2C and 2D), even when their level of fluorescence is low (see the neuron in the red box).

As a result, neurons may look different in cross-spectral images and often have a less pronounced dark central nuclear region. Low-frequency pixel correlations tend to show a central maximum that declines to the border of a neuron (Figures 2K–2N). In contrast, non-active neurons with high levels of fluorescence, which are visible in the average fluorescence projection (see the neuron in the cyan box, Figures 2A–2F) and background fluorescence in the neuropil signal disappear in the spectral images. ROIs can thus be selected by searching for the largest local maxima in these spectral images and drawing contours around them (see STAR Methods for details) (Figures 2K and 2M). Because the images of different frequency components reveal different neurons and/or neuronal compartments, a complete set of ROIs is created by adding up all ROIs detected in the spectral images of all low-frequency components (Figure 2D).

The low-frequency components in the fluorescent signal are due to the exponential decay of GCaMP6f, which, like various other indicators, has a time constant in the order of 400 ms (Chen et al., 2013). A theoretical model of the Fourier transform and spectral power of calcium transients (Figure 2G; see STAR Methods) with a time constant based on this decay half time make clear that the greatest power is expected to be found for frequency components below 1 Hz. To illustrate this in more detail, we simulated video data at a 1-Hz sampling rate. Simple neurons were simulated with fluorescence spikes that had a half time of around 500 ms. Noise was added to each pixel in the simulated video. Power spectra for this video were calcu-

lated using SpecSeg. ROIs with spikes show low-frequency components (Figures 2H–2J). Noise in the fluorescence data (Figure 2H) adds higher frequencies to the power spectra. The low-frequency components are therefore a signature of spiking neurons.

Automated ROI refinement

Some contours selected by SpecSeg may contain overlapping or closely juxtaposed neuronal elements (Figures 2K and 2M), especially if the density of cell bodies or neurites is high. To prevent this problem, contours are further constrained based on whether the fluorescence pixel traces (not the Fourier transforms) within the contour are tightly correlated as would be expected if they are from the same neuron. We calculate this by taking the median fluorescence trace of the pixel from the local maximum (seedpoint) and its eight neighboring pixels and correlating this trace with each pixel trace in the preselected contour. This results in correlation values for every pixel in the ROI (pixel R). Next, using a threshold set halfway the maximum correlation strength and the correlation strength in a band of tissue surrounding the ROI (this threshold can be adapted), a new contour is selected around the best-correlated pixels. To save an estimate of how well the fluorescence pixel traces are correlated inside an ROI, we square the pixel correlation values of an ROI and average them, resulting in the mean R^2 (Figures 2L and 2N). The mean R^2 indicates how much of the pixel variance is explained by a shared signal within an ROI. We assume that this signal is the actual activity trace of a neuron, possibly including a general neuropil signal. The remaining variance is due to sources surrounding a neuron influencing individual pixel traces separately.

Versatility of ROI selection approach

Restrictions can be applied to select ROIs around cell bodies, such as the minimal and maximal surface area of the ROI or its roundedness. Without these restrictions, thin elongated contours can be easily selected, making it also possible to define ROIs on dendrites and axons (Figures 3A–3C). We tested the ROI selection approach on various datasets, obtained by two-photon calcium imaging using a GRIN-lens in the visual thalamus (Figure 3D) as well as single-photon miniscope imaging in the pre-motor cortex (M2) and striatum (Figures 3E and 3F).

Figure 2. Comparison of fluorescence and spectral power images

- (A) Average fluorescence projection. The colored squares indicate regions containing an easily recognizable active neuron (dark blue); a silent neuron (cyan); no neuron (green); a hardly recognizable but active neuron (red). Scale bar, 40 μm .
- (B) Maximum fluorescence projection. Note that the neuron in the cyan region can be easily detected, while the neuron in the red region is hardly visible.
- (C and D) Cross-spectral images. Note that the silent neuron in the cyan region is hardly visible in the cross-spectral images, while the active cell in the red region can be easily identified.
- (E) The cross-spectral power of the four example areas. Note the logarithmic y axis.
- (F) Pixel traces of the example areas denoted in (A–D).
- (G) A simple exponentially decaying signal over time (inset) with a half time of 400 ms, produces the highest values in low frequency.
- (H) Simulated data are used in the SpecSeg pipeline. The thick black line belongs to an ROI consisting of 61 pixels. The thin gray line is the signal belonging to one example pixel. A trace containing only noise and no calcium transients does not contain low-frequency components.
- (I) A trace with one calcium transient, with a half time around 600 ms, and noise produces high values in the low-frequency spectral components in the SpecSeg pipeline.
- (J) More transients produce higher spectral power values, with a different spectral profile.
- (K and M) Automatic ROI creation via the spectral image. The contour that is found around these putative neurons is displayed in red. For these neurons, the threshold for pixel inclusion was too low, resulting in ROIs that are too large. These ROIs are automatically refined based on pixel correlation. Scale bar, 20 μm .
- (L and N) The pixel correlations of the ROIs displayed in (K and M). The color displays the pixel correlation of the signal from each pixel with the signal from the ROI local maximum. The pixel correlation threshold facilitates ROI refinement. The new contour, shown in white, is determined by the pixel correlation values.




Figure 3. Examples of fluorescent signals, cross-spectral power, and ROI selection in datasets derived from different brain regions and/or imaging techniques

Left column: maximum fluorescence projections. Middle column: cross-spectral power, with different frequency components in different colors. Right column: cross-spectral power maximum projection, with contours of ROIs in red.

(A) Two-photon microscopy of layer 5 apical tufts in mouse V1.

(B) Two-photon microscopy of layer 5 apical trunks in mouse V1.

(C) Two-photon microscopy of visual lateromedial area axon projections to V1.

(D) Two-photon microscopy of visual thalamus using a GRIN-lens.

(E) Single-photon miniscope imaging in mouse premotor cortex (M2).

(F) Single-photon miniscope imaging in mouse striatum.

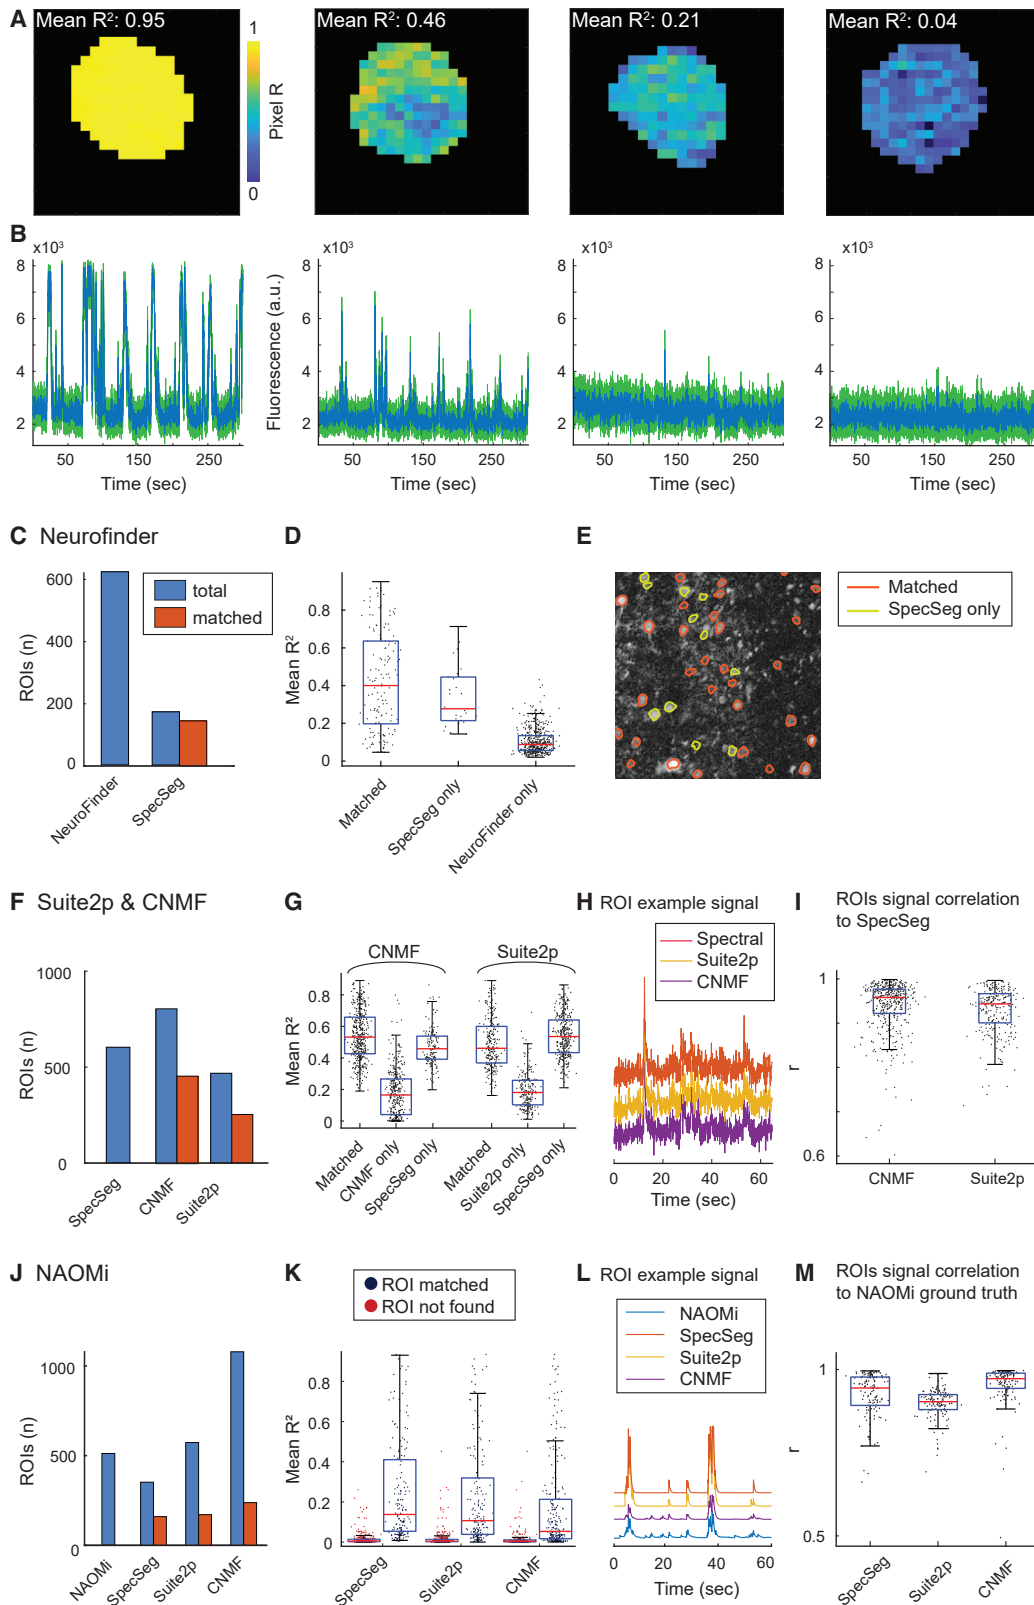
We also tested how ROI selection is influenced by the resolution of the recordings by imaging the same part of V1 at different magnifications (Figure S1). We found that ROI selection is quite consistent at different resolutions. Only when ROI sizes become very small (~ 25 pixels total, so $\sim 5 \times 5$ pixels) does the number of ROIs start to decrease significantly.

Comparing SpecSeg ROI selection in two-photon imaging data with other packages

To compare the efficiency of our method with other packages, we first made use of Neurofinder datasets (<http://neurofinder.codeneuro.org/>). Our results (SegSpect0, by Spectral Segmentation) scored in the mid-range of all tested methods. Because this seemed a relatively low score considering the good results we obtained with SpecSeg on our own datasets, we assessed the Neurofinder results in more detail. We noticed that, in several Neurofinder datasets, SpecSeg missed over 50% of Neurofinder “ground-truth” ROIs and found additional ROIs that were not included as ground-truth ROIs, both contributing to the lower score. To understand why SpecSeg selected different ROIs than those considered ground truth, we calculated the mean R^2 of the pixels within the latter ROIs (Figures 4A and 4B)

(see STAR Methods). In the Neurofinder datasets, most ground-truth ROIs had a mean R^2 below 0.15, suggesting that >85% of their signal variance originated from external sources or noise. ROIs selected by SpecSeg (Figures 4C–4E) had much

Single-photon miniscope signals are generally more correlated and contaminated by out-of-focus signals. Therefore, we included an additional background correction tool for miniscope data (BackgroundSubtractSbx.m) (Figures 3E and 3F).



(legend on next page)

higher mean R^2 values (Figure 4D). This resulted in only part of the ROIs defined as ground truth in the Neurofinder datasets to correspond with those drawn using SpecSeg. Importantly, SpecSeg missed less than 1% of the ground-truth ROIs with a mean R^2 over 0.3. SpecSeg also selected additional ROIs representing active neurons that were not identified among the Neurofinder ground-truth ROIs. These ROIs had high pixel correlations and showed clearly defined responses (Figure 4B). This shows that it is crucial to select ROIs based on activity and not solely on cell morphology visualized by the baseline fluorescence of the calcium indicator, as some responsive neurons may be invisible, especially when using calcium indicators with very low calcium-free fluorescence, such as jRCaMP7c (Dana et al., 2019).

We also compared ROIs selected by SpecSeg with those obtained using two frequently used software packages used for ROI selection: Suite2p (Pachitariu et al., 2017) and CNMF/CalmAn (Giovannucci et al., 2019). We analyzed the ROIs selected by Suite2p, CNMF, and SpecSeg in a dataset of a densely labeled population of cortical neurons in a transgenic mouse line (Ai93 x G35-3-cre) (Madisen et al., 2015; Sawtell et al., 2003). In the same dataset, CNMF detected 808 ROIs, SpecSeg detected 693, and Suite2p detected 468 (Figure 4F). SpecSeg ROIs that overlapped with CNMF or Suite2p had high mean R^2 values, while those that were selected by Suite2p and CNMF, but missed by SpecSeg, had low mean R^2 values (Figure 4G). We also compared the raw signal traces of the matching ROIs (Figures 4H and 4I) that we extracted using all three methods and calculated the correlation between the raw signals using SpecSeg with those from CNMF and Suite2p (Figure 4I). This shows that the signal traces SpecSeg extracts are highly similar to those extracted with the other methods.

The main limitation in the above approaches to compare the different methods is that reliable ground-truth data with which the selected ROIs or extracted signals can be compared is lacking. Therefore, we simulated a calcium imaging dataset using NAOMi (Song et al., 2021) and tested how the three ROI selection toolboxes performed at selecting ROIs and extracting the signal. CNMF detected the highest number of total ROIs (1,079) and found 238 of all 514 ground-truth somas. Suite2p detected 573 ROIs and identified 171 of the ground-truth somas. SpecSeg

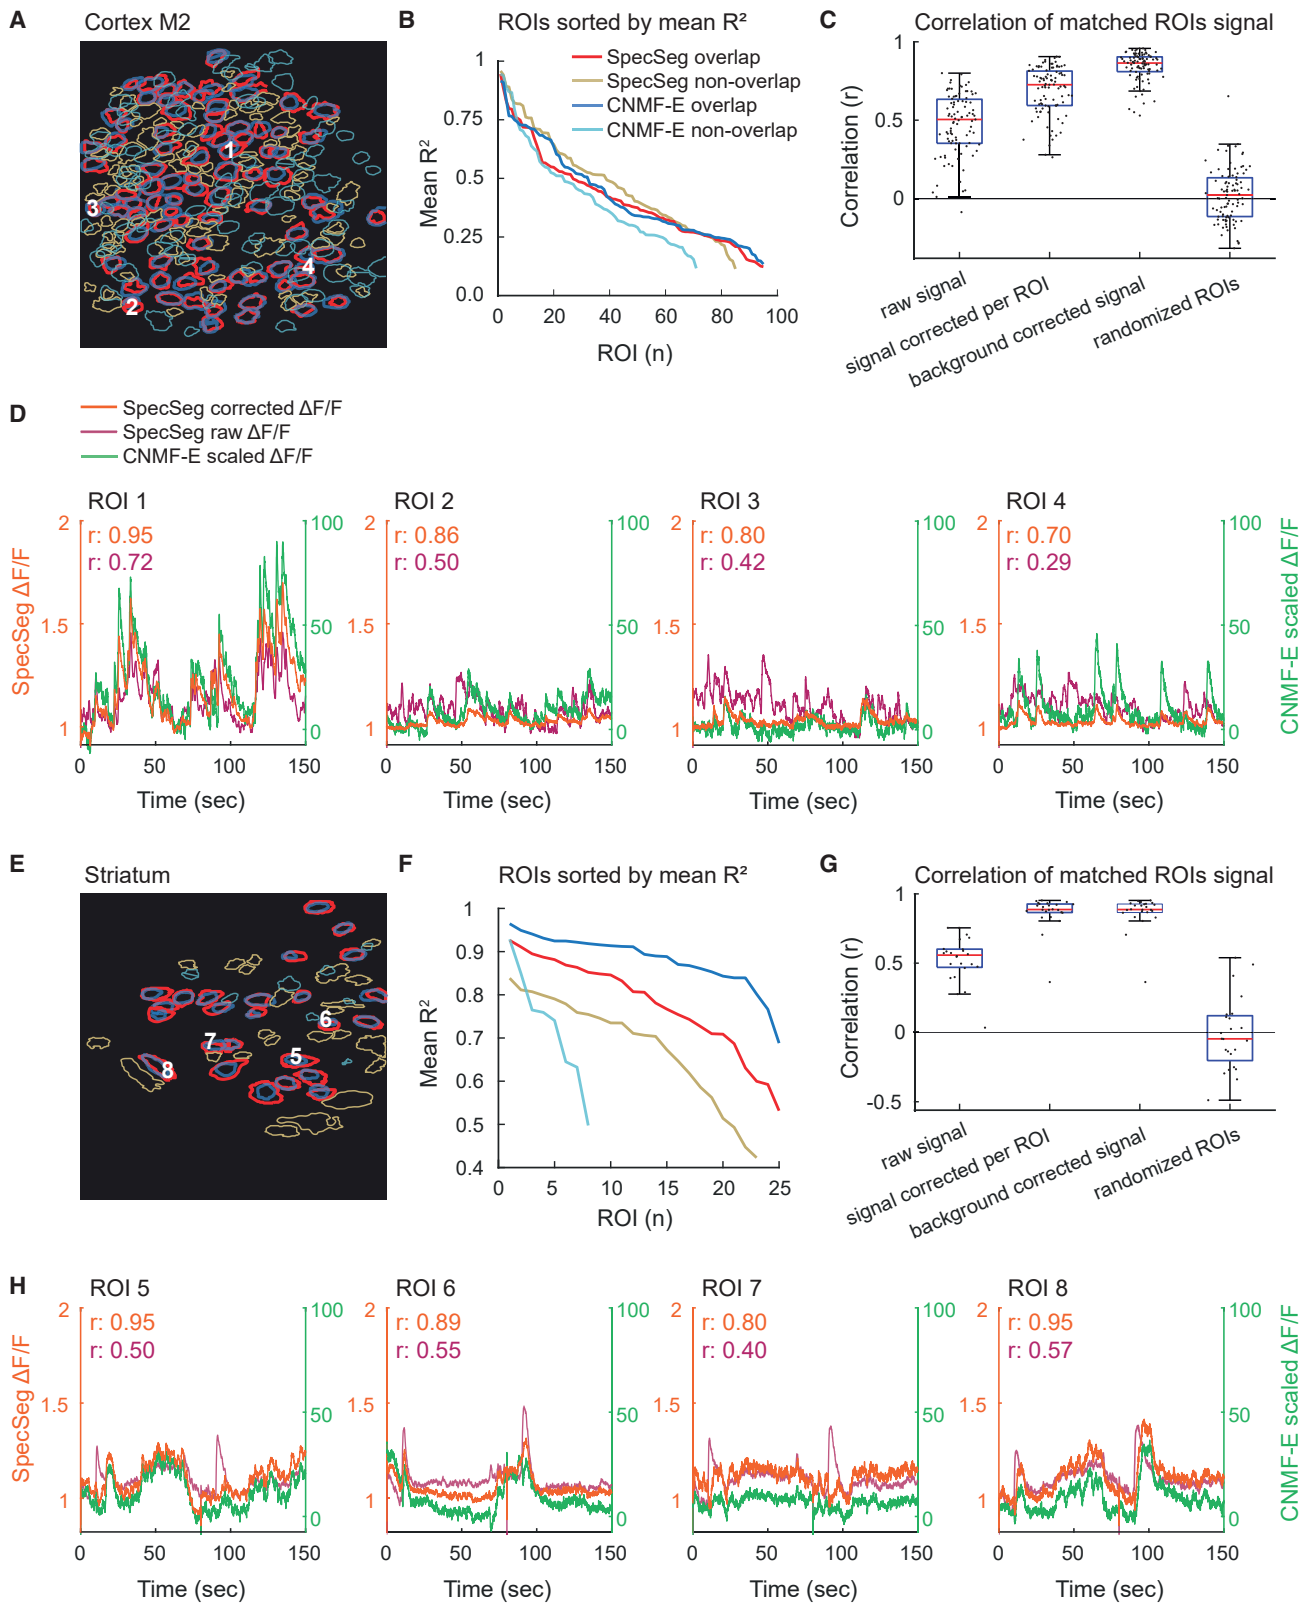
found 352 ROIs, of which 160 were ground-truth somas. This means that, while CNMF detected most total and ground-truth ROIs, precision was on the lower side as only 22% of all ROIs matched ground-truth. Precision of Suite2p was higher, with 30% of ROIs matching ground-truth. SpecSeg detected the lowest number of ROIs, but of those found 46% matched ground-truth (Figure 4J). Also, on average, ROIs that SpecSeg found had higher mean R^2 values than those detected by Suite2p and CNMF. The ground-truth cells that were missed by all methods had very low mean R^2 values (Figure 4K). We conclude that, in this analysis, CNMF was the most sensitive approach, while SpecSeg was the most precise.

When we compared the (denoised) signal traces of ROIs detected by all three methods with the ground-truth signal traces (Figure 4L), we observed that they were all highly correlated (Figure 4M). This implies that the extracted signals are not strongly affected by bleed-through or neuropil contamination. The signal traces from CNMF matched ground-truth data best, probably because CNMF denoises the signal traces during the same step as it identifies the ROIs. SpecSeg signal traces correlated somewhat better than those from Suite2p, although both used MLspike (Deneux et al., 2016) to denoise the traces. Together, these results show that SpecSeg is very effective at selecting ROIs from which a signal can be reliably extracted. This increases the quality of the extracted signal and circumvents the need for signal demixing, as all pixels included in the ROI are highly correlated.

A possible caveat may be that decreasing the size of the ROI to exclude uncorrelated pixels, as SpecSeg does (as opposed to signal demixing), could make the signal noisier or less reliable. Therefore, we tested whether extracting the signals from a smaller part of a cell would result in a lower signal quality. We calculated the mean R^2 as a proxy for signal quality from a large set of ROIs from simulated NAOMi data, of which we cut an increasing number of pixels (Figures S2A–S2C). We found that signal quality only started to decrease once more than 50% of the ROI was removed. In this analysis, 50% removal usually resulted in the loss of the original seed point, causing the decline in signal quality. This will never occur when SpecSeg removes non-correlated pixels. A similar result was obtained using actual

Figure 4. Comparing SpecSeg

- (A) Correlation of each pixel trace within an ROI with the median pixel trace at the local maximum of cross-spectral power and its eight neighboring pixels, of four example ROIs with decreasing mean R^2 from the 03.00 Neurofinder dataset.
- (B) Fluorescence traces (blue) plotted against their SEM \times 3 (green) for the four example ROIs. Note that around a mean R^2 of 0.2 or lower, the signal becomes non-significant.
- (C) The amount of ROIs in a Neurofinder dataset, according to Neurofinder, SpecSeg, and how many ROIs of the two were matched together.
- (D) Mean R^2 for ground truth ROIs in (C).
- (E) Spectral image of the Neurofinder dataset, with SpecSeg ROIs encircled by red (overlapping) and yellow (non-overlapping) contours.
- (F) A very densely labeled population of cortical neurons in a transgenic mouse line is analyzed with SpecSeg, Suite2p, and CNMF. We compared the results from SpecSeg against those of Suite2p and CNMF. The horizontal line shows the number of ROIs found with SpecSeg.
- (G) The mean R^2 of all the ROIs in the dataset. ROIs that were not found by SpecSeg had a low mean R^2 , while there were also ROIs that were not found by Suite2p or CNMF, which did have a higher mean R^2 .
- (H) Raw fluorescence traces from an ROI found with all three methods, with average mean R^2 . The signal correlates well between methods.
- (I) Pearson correlation of the raw fluorescence signals between all matched ROIs.
- (J) Using simulated calcium imaging data with NAOMi, we compared that ground-truth dataset with SpecSeg, Suite2p, and CNMF. The horizontal line shows there were 514 neural somas near the focal plane.
- (K) Mean R^2 NAOMi ROIs. All methods find somas with high mean R^2 . The NAOMi ROIs that were not found had a very low mean R^2 .
- (L) Denoised signals from one example ROI that was found by all three methods.
- (M) The denoised signal correlates to the NAOMi ground truth very well for all three methods.



(legend on next page)

calcium imaging data. We quantified this for large numbers of simulated or real neurons (Figures S2D and S2E). We conclude that extracting the signal traces from part of an ROI to exclude contaminated signals does not significantly decrease the quality of the signal.

Single-photon miniscope data analysis comparison with CNMF-E

Single-photon miniscope imaging data present additional problems for ROI selection and signal extraction. Single-photon data suffer heavily from out-of-focus fluorescence. Background fluorescence can contain fluctuations that are stronger than the signal from individual neurons themselves. The images also contain vignetting, meaning they are darker near the edges. These problems require additional processing. A widely used single-photon analysis is CNMF-E (Zhou et al., 2018). With this method, ROIs are initialized based on a heavily filtered image, which is not suited for signal extraction. Therefore, the ROIs are refined and the signal estimated by constrained matrix factorization. SpecSeg uses a different approach, and first removes the out-of-focus fluorescence and vignetting from the video data itself via background subtraction. The background subtraction estimates background fluorescence for every frame individually by filtering each frame with a large disk (~200 μm). The resulting background fluorescence image is subtracted from the original frame (see STAR Methods). This approach decreases both out-of-focus fluorescence and vignetting, while small signal sources remain present. The data from Figures 3E and 3F were analyzed with both CNMF-E and SpecSeg. We found that 60% of the ROIs had spatial footprints that were more than 30% overlapping between the CNMF-E and SpecSeg analysis. Matched ROIs are shown in Figures 5A and 5E. Unlike in the two-photon Neurofinder dataset comparison, the mean R^2 was not significantly higher in the SpecSeg ROIs (Figures 5B and 5F), probably because CNMF-E also selects ROIs based on activity, and not predominantly on anatomical features. The extracted signal traces from ROIs correlated well with those from corresponding ROIs from CNMF-E after background correction. Without background correction the signal traces correlate little with CNMF-E or the SpecSeg corrected signals (Figures 5C, 2D, 2G, and 2H).

Speed

To test the speed of SpecSeg, we automatically timed the different components of the pipeline on two different computers analyzing 16 different datasets, varying in size from 6 to 51 gigabytes. Depending on the hardware used and the size of the dataset, the different components of the analysis pipeline varied in the time required for completion (Figure S3A). This makes it difficult to provide hard numbers about the speed at which the SpecSeg pipeline processes the data. However, a useful indicator is that 1 h of acquired data imaged at 15 Hz was processed by automated analysis in approximately 5 h, from NoRMCorre motion correction until signal extraction. More than half of the time was used by the NoRMCorre motion correction. We also assessed how the dimensions of the video influenced the processing time. We varied the size (from 70 \times 70 to 300 \times 300 pixels) and the durations (from 166 to 5,500 s). The time for processing was approximately linear with the size of the recording (Figures S3B and S3C). We conclude that the speed of the SpecSeg pipeline is not a bottleneck, which in most cases performs the analysis overnight.

User interface for rejecting, splitting, or manually defining ROIs

In the SpecSeg pipeline, we included a graphical user interface to add, split, or delete ROIs. ROIs are not drawn around neurons that are silent in the imaging session. Such neurons may need to be included in certain analyses, for example in chronic experiments in which these neurons are active in other recording sessions. The user interface therefore enables adding ROIs of such neurons (Figure S4A). Sometimes ROIs are created that envelop multiple neurons (Figure S4B). To visually guide the splitting of these ROIs, pixel traces within the ROI are correlated with four reference points and correlation values are color coded. In the case in which the ROI contains signals from more than one neuron, this becomes clearly visible (Figure S4C). The correlation values are then used in k-means clustering to split the ROI into multiple ROIs, or to delete parts of the ROI, on the user's request (Figure S4D). ROIs can also be deleted entirely or marked to be kept when more stringent criteria for ROI selection are applied. Finally, ROIs can be drawn in manually, or added by the computer, based on the fluorescence or spectral image (Figure S4A). Any image with the same size as the

Figure 5. Comparison of CNMF-E and SpecSeg for the analysis of single-photon miniscope data

- (A) Miniscope calcium-imaging data from premotor cortex (M2) are analyzed using both SpecSeg and CNMF-E. ROIs identified by SpecSeg and CNMF-E are shown. Colors indicate whether or not the ROIs identified by the two methods overlapped (see (B) for color coding). White numbers, identify the overlapping ROIs whose signals are shown in (D).
- (B) Mean R^2 of ROIs sorted by magnitude. Overlapping CNMF-E and SpecSeg ROIs represent the same neurons. However, their spatial footprints are slightly different, also resulting in slightly different mean R^2 . Non-overlapping ROIs do not have a significantly lower mean R^2 than overlapping ROIs.
- (C) Signals extracted by SpecSeg are strongly correlated with the signals extracted by CNMF-E for overlapping ROIs. In contrast, correlating the raw signal with the CNMF-E signal results in lower r values. Background correction of the ROI signal with a surrounding donut ROI, as used for two-photon data, increases the correlation. Pixel- and frame-wise background correction increases the correlation even more, illustrating the importance of background correction. The fourth condition "randomized ROIs" is a control and correlates the background corrected signal from the SpecSeg ROIs with random ROIs from CNMF-E.
- (D) Signals from four example neurons, marked in (A), are shown. The r values represent the Spearman correlations between the ROI SpecSeg signals and the CNMF-E signals. The correlation between the CNMF-E signal and the background-corrected SpecSeg signal (background-corrected signal in (C)) is indicated in orange, while correlation with the raw SpecSeg signal is indicated in magenta.
- (E) Same as (A) but for calcium imaging data from striatum.
- (F) Mean R^2 of ROIs sorted by magnitude.
- (G) For the striatum data, ROI-wise signal correction is more effective than in the M2 data, because the striatum has a sparser population of neurons. This allows for more data to be used in the subtraction, making it more similar to the pixel-wise background correction.
- (H) The signal from the four marked ROIs marked in (E).

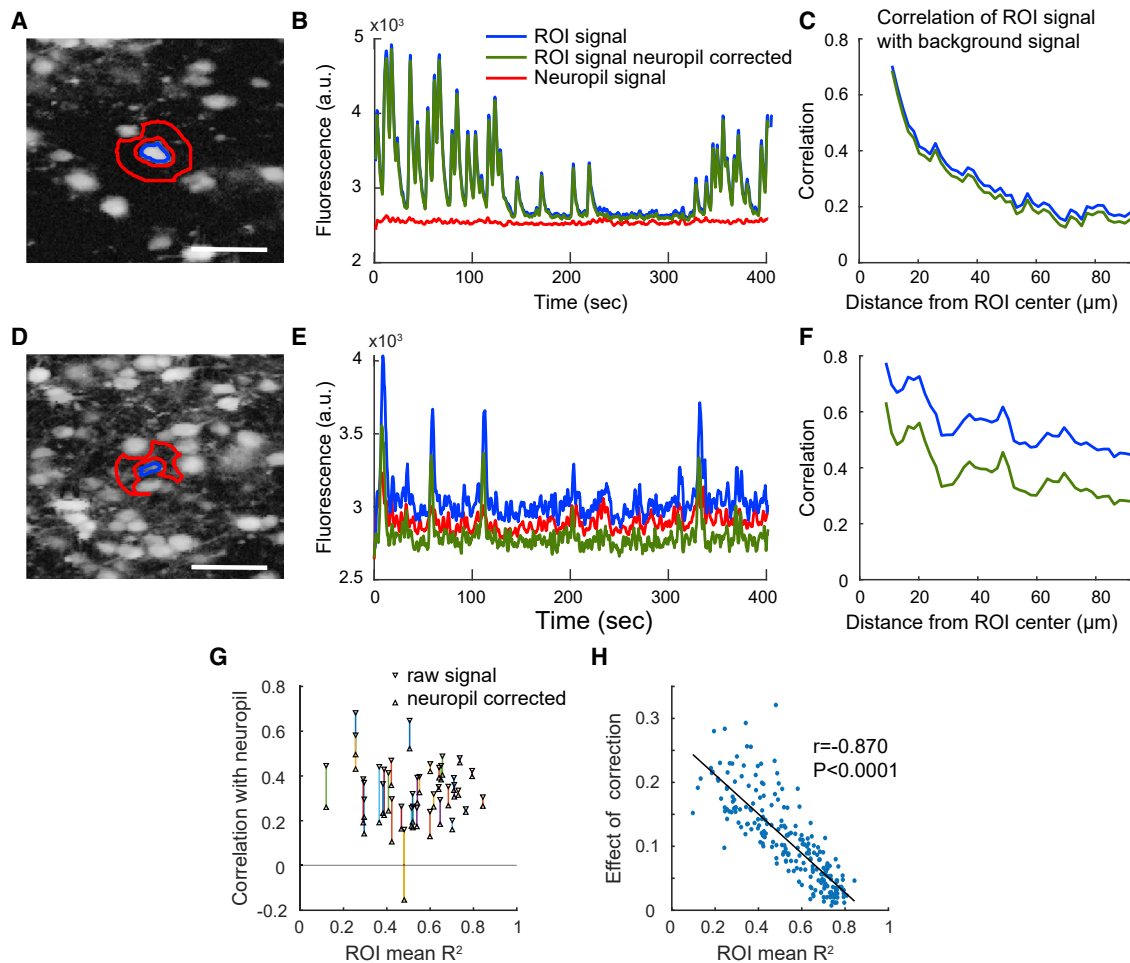


Figure 6. Neuropil estimation from surrounding pixels

- (A) Spectral image with an example ROI (blue contour) and surrounding background (red contour). Scale bar, 40 μm .
 (B) Fluorescence signal of the example ROI of (A). The neuropil signal is almost non-existent and neuropil correction does not change the ROI signal much.
 (C) Correlation of neuropil signal with ROI signal depends on the distance from the ROI. The fluorescence signal was extracted from rings at increasing distances around the example ROI, excluding other ROIs. Those background signals were correlated with both the raw signal of the ROI (blue line), and with the ROI neuropil corrected signal (green line). Neuropil correction slightly decreases the correlation of the background signal.
 (D) Spectral image with a different example ROI (blue contour) and surrounding background (red contour).
 (E) In this ROI, the neuropil signal caused more contamination than in the example in (A). Neuropil signal was larger because more dendrites are present, and the ROI itself has a less strong signal. The neuropil correction seems necessary for this ROI.
 (F) The neuropil correction decreases the correlation with the surrounding background signal.
 (G) Correlation with neuropil signal of multiple ROIs, before and after neuropil correction. The distance between these shows the “neuropil correction effect.”
 (H) The neuropil correction effect is significantly smaller when an ROI has a stronger mean R^2 .

recording can be loaded from the workspace into the user interface and used as a source for ROI creation. During any of these procedures, the user interface supports the visualization and exploration of the data in various ways. The calcium trace from any location can be plotted and immediately compared with calcium traces from other chosen locations. Together, these additional tools provide the means to adjust the ROI selection process in an easy, insightful, and user-friendly way.

Neuropil correction and signal extraction

Once the ROIs are selected, the calcium signals from the cells can be extracted by averaging all pixel traces within the ROI con-

tours. However, because calcium signals from the neuropil can contaminate those of the cells, the neuropil signal is first subtracted from the ROIs (Zhou et al., 2018). To achieve this, we select an area surrounding each ROI, taking care to avoid other ROIs (Figure 6, see STAR Methods for details). The average of the pixel traces in this area is used as an estimate of the local neuropil activity. After multiplying it by 0.7 it is subtracted from the averaged fluorescence trace of the selected ROI. This method of neuropil subtraction is widely used (Chen et al., 2013; Keemink et al., 2018; Khan et al., 2018; Tegtmeier et al., 2018). To test how neuropil subtraction alters the signal, we correlated the pixel trace of ROIs with the pixel traces from the

surrounding area (Figures 6A–6F). We found that, in most cases, neuropil correction barely changed the ROI signal (Figures 6A–6C). In some cases, however, the signal improved after neuropil subtraction (Figures 6D–6F). There may be multiple reasons why neuropil subtraction has less effect in our hands than in studies by others (Pachitariu et al., 2017). One plausible reason could be that ROIs with a relatively high mean R^2 , preferentially selected by SpecSeg, are less influenced by neuropil contamination. To test this, we plotted the effect of neuropil correction on the signal extracted from ROIs against their mean R^2 (Figures 6G and 6H). This revealed that, indeed, ROIs with higher mean R^2 are barely affected by neuropil correction while those with low mean R^2 are strongly affected. Thus, selection of ROIs with high mean R^2 has the added advantage that it reduces the need for neuropil subtraction.

Automated retrieval of the same ROIs in chronically recorded datasets

One major strength of multi-photon microscopy of calcium responses is that changes in neuronal activity can be followed at the single-cell level over prolonged periods of time. For this, it is essential to identify the same ROIs in chronic recordings. In large datasets containing many recordings, this is a daunting task to do by hand, even after automated ROI selection. We therefore developed a toolbox for automated matching of ROIs in complete series of chronically recorded datasets. It involves the registration of the sequential recordings with each other, through translation and rotation of the spectral images. Next, overlapping ROIs are searched for in all possible pairs of recordings in the series (i.e., each recording is compared with all other recordings). By thresholding for a (user-defined) minimal amount of overlap between each pair of ROIs in both directions, only ROIs with similar shapes and sizes are matched. Finally, merging all the matched pairs of ROIs results in a “matching matrix” containing the matched ROIs that are found in the series of recordings. This multi-session ROI matching tool can also be used with ROIs selected using other tools (instructions can be found in the PDF manual).

An example of ROI matching in three sequential recordings of V1 neurons, performed with 2–4 weeks in between sessions, is shown in Figure 7. In these color-coded images, many cells can be retrieved in all three recordings. However, some neurons are not retrieved in one or two of the recordings, as indicated by the triangles. This may be caused by various factors, such as the precise angle and depth of the recording, changes in viral expression of the calcium indicator, or the lack of activity of a neuron during one of the recordings.

Comparing ROIs between recordings

For a quantitative assessment of the number of ROIs that can be matched between recording sessions we analyzed data from 5 mice that were imaged 7, 8, 12, 13, and 12 times, respectively, over a period of 2–5 months. This resulted in 21, 28, 66, 78, and 66 possible pairs of recordings, respectively, in which we investigated the number of matched ROIs. With all data of the five mice pooled together, on average 38% of the ROIs were matched between recordings pairs, corresponding to 94 ± 57 ROIs. The percentage of matched ROIs was calculated as the percentage of matched ROIs in the recording with the least

ROIs per recording-pair. The number of matches was influenced by the time between two recordings. Figure 7B shows that there was a significant correlation between the number of matched neurons and the time between recordings in two of the three mice (Spearman correlation: for mice 1 and 2, $p < 0.005$; for mice 3, 4, and 5, $p < 0.0005$; $r = -0.6, -0.35, -0.66, -0.93$, and -0.68 for mice 1 to 5, respectively). The fits shown in the figure are exponential fits ($y = a \times e^{b \times x}$) with the initial value (a) at 75%, 71%, 53%, 56%, and 61%, and decay rate (b) at $-0.0084, -0.0036, -0.0171, -0.0160$, and -0.0123 days for mice 1 to 5, respectively). This suggests that it is not only variations in the angle or exact location of the recordings that determines whether ROIs can be matched, but probably also slower, biological processes, such as viral expression of the calcium indicator, learning- or age-induced changes in neural activity, cell death, or anatomical changes of the brain. Indeed, Figure 7C shows that correlations between the registered spectral images decreases over time (Spearman correlation: $p = 0.0006$ for mouse 1; for all other mice, $p < 0.0005$, $r = -0.68, -0.58, -0.61, -0.93$, and -0.82 for mice 1 to 5, respectively), suggesting that the structure of the imaged location and activity patterns and GCaMP6f labeling of the neurons alter over a period of weeks to months. Changes in GCaMP6f labeling and/or neuronal activity is supported by the observation that more ROIs were detected on later imaging sessions for some mice (Spearman correlation: $p = 0.2, p = 0.008, p = 0.51, p = 0.97$, and $p = 0.015$, and $r = 0.56, 0.83, -0.21, -0.01$, and 0.68 for mice 1 to 5, respectively) (Figure 7E).

We also asked through how many sessions individual ROIs could be tracked. For this analysis, we pooled the data from the five imaged mice. We limited the number of sessions to seven (constrained by the mouse with the fewest recordings), choosing the sessions of the mice in such a way that the number of days between the first and last recording was close to 65. We then calculated in how many imaging sessions the ROIs from recording 1 were found back in the later sessions (Figure 7D). These did not need to be consecutive sessions as ROIs could disappear and reappear (Figure 7). We found that 22% of all ROIs could be matched in all 7 recordings, averaged over the 5 mice, which corresponded to 213 neurons, or $213 \times 7 = 1,491$ separate ROIs in total. Twenty-eight percent of ROIs (354 ROIs) were found only in the first recording.

To understand to what extent ROIs are missed by SpecSeg in subsequent recordings due to differences in activity or to specific workings of SpecSeg, it is necessary to exclude changes caused by human error or biological processes. To achieve this, we split up an imaging session in two parts, and tested whether the same ROIs were found (Figure S5). Of 320 ROIs, 247 were found in both the first and the second half, while 13 were only found in the full recording, but not in the first or second half separately. Twenty-eight ROIs were only found in the first half, and 32 only in the second. The average mean R^2 was much higher in ROIs found in both halves of the recordings. ROIs that were missed in one half of the recording had lower mean R^2 values that were even lower in the half in which they were not detected. Often, this was caused by the absence of spiking activity in that period. We concluded that the main cause of not detecting an ROI is a weak signal.

A chronic tracking example

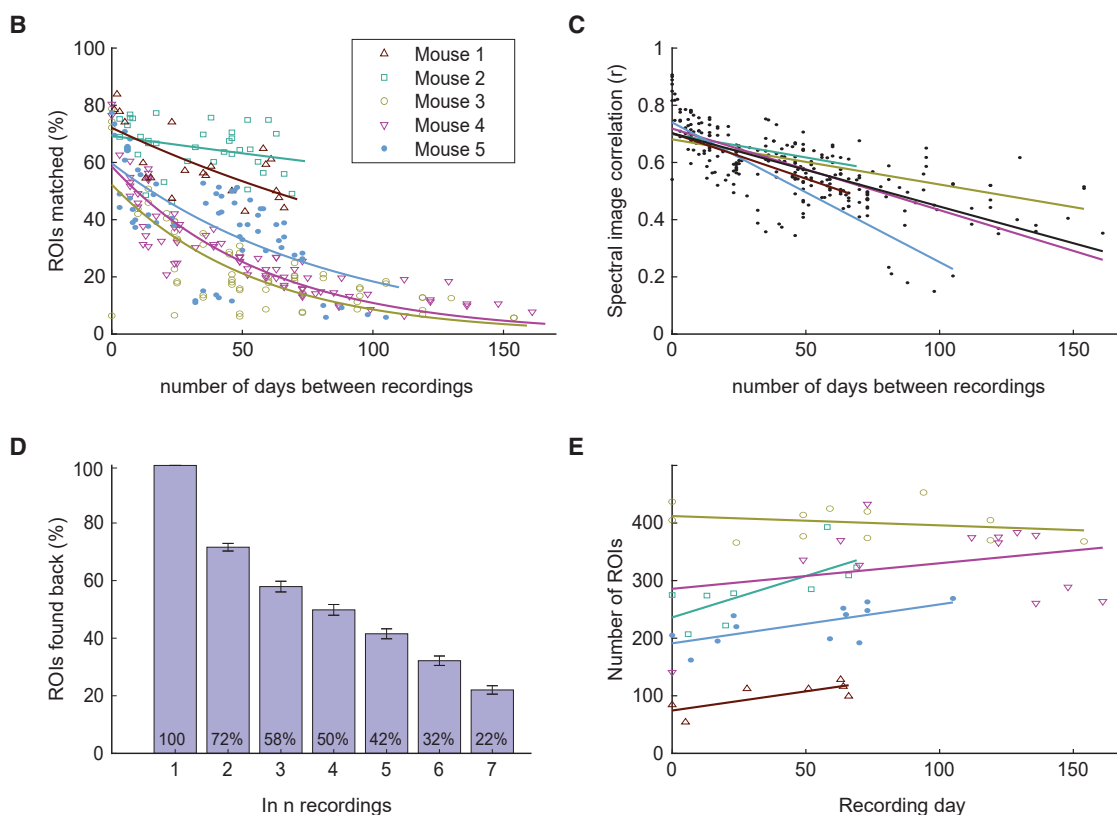



Figure 7. Chronic ROI matching

(A) An example of matched ROIs in three sequential recordings. An ROI's center of mass (CoM) is shown with a dot. The dot's color shows in which recordings ("rec") it was retrieved. The triangles denote the position of some CoMs of ROIs that were present in other recordings, but missed in the recording in which the triangle is shown. The color denotes in which recording the ROI was present. The right bottom figure shows the three spectral images overlaid with red, green, and blue, resulting in the same color mixing as the ROI CoM dots. Scale bar, 40 μ m.

(B) The number of matches between recordings decreases if there is more time between the recording pairs.

(C) Spectral correlation for each recording pair. Recordings that are recorded closer together in time have better correlation coefficients. This shows that the imaged brain changes over time.

(D) Checking in how many recordings the ROIs from recording one were found back in all seven recordings. To be able to pool all the mice together, seven recordings were analyzed per mouse. Recordings were chosen so that the time between the first and last recording was as close to 65 days as possible. The number of putative neurons in each condition were: 1,146, 792, 606, 516, 421, 319, and 213 for conditions 1 to 7, respectively. Error bar shows standard error of the mean for the five mice.

(E) The number of ROIs found in the mice recordings did increase significantly for mice 2 and 5 (Spearman correlation: $p < 0.05$; for mice 1, 3, and 4 the p values were 0.20, 0.51, and 0.97, respectively).

DISCUSSION

Here, we describe SpecSeg, an open-source toolbox for automated ROI selection, ROI editing, neuropil correction, signal extraction, and chronic matching of ROIs. Like several other approaches (Kaifosh et al., 2014; Smith and Häusser, 2010; Spaen et al., 2019), SpecSeg ROI selection is based on correlation of co-active pixel traces. In SpecSeg, however, the first step of ROI selection involves calculating the low-frequency (0.013–0.5 Hz) cross-spectral power components. These spectral components can be visualized, yielding exceptionally clear images of the local functional anatomy that include neuronal cell bodies and neurites. The low-frequency cross-spectral power results from the exponential decay of calcium transients. In contrast to the high-frequency components, the low-frequency components are less sensitive to noise. Moreover, activity in neuropil is less correlated, resulting in reduced low-frequency activity, thus enabling the efficient separation of neuronal structures from the background.

This first step of ROI selection requires few initial constraints, permitting inclusion of ROIs with irregular shapes or sizes. Next, a threshold is set on the pixel correlations within the ROI, leading to robust separation of ROIs containing highly correlated pixels and rejection of overlapping neural elements. Additional filters for size or roundedness of the ROIs can be included by the user, allowing ROI selection around cell bodies, dendrites, or axons. These morphological constraints are intuitive, making ROI selection versatile and user friendly.

The pipeline also includes a graphical user interface, facilitating quality control of the selected ROIs. This interface includes an alternative method for visualizing pixel correlations, aimed at detecting whether multiple cells are present within one ROI. This makes it easy to split or reject ROIs if needed. Also, ROIs can be added if the experimenter believes that cells are missed, and the quality of their signal can be assessed immediately.

We employed various methods to compare the efficiency of SpecSeg with other ROI segmentation packages, such as Suite2p (Pachitariu et al., 2017) and CNMF/Calman (Giovannucci et al., 2019). The first approach was to make use of NeuroFinder (<http://neurofinder.codeneuro.org/>), which is based on calcium imaging datasets with ground-truth ROIs selected by human annotators. However, we noticed that, within many Neurofinder ground-truth ROIs, pixel traces showed little correlation (low high mean R^2), implying that the signal extracted from such ROIs was predominantly derived from the surrounding neuropil. In contrast to human annotators, SpecSeg predominantly selected ROIs with a relatively high mean R^2 , even if they were dim, but missed silent or very noisy ROIs even if they were bright. We think it is debatable whether silent or noisy ROIs should be included in the data analysis, as they may contaminate the actual neural code of the imaged neurons. Only in some instances it is useful to include a neuron that does not show any activity in a dataset, for example when chronically tracking the activity of individual neurons. For such cases, the SpecSeg graphical interface provides the option to manually include neurons when they are not detected by the activity-based approach.

We also used two approaches that were independent of human annotators to compare the performance of SpecSeg with

Suite2p and CNMF. The first was to select ROIs using SpecSeg, CNMF, and Suite2p on a densely labeled dataset from mouse V1 and compare the number of selected ROIs detected by the different methods and their mean R^2 . In this analysis, SpecSeg retrieved slightly fewer ROIs than CNMF but considerably more than Suite2p. Those found by SpecSeg had the highest mean R^2 . Interestingly, SpecSeg ROIs that were missed by the other approaches had high mean R^2 values, while those found by the other packages but missed by SpecSeg had low mean R^2 values.

The second approach was to simulate calcium imaging data using NAOMi and test how well SpecSeg, Suite2p, and CNMF retrieved the ground-truth ROIs. In contrast to the Neurofinder approach, NAOMi ground-truth ROIs are well defined and the actual spiking activity of the modelled neurons is known. In this analysis, SpecSeg found fewer ground-truth ROIs than Suite2p and CNMF. However, Suite2p and CNMF found considerably more ROIs that did not match ground-truth ROIs. Moreover, the mean R^2 values of SpecSeg ROIs were higher than those detected by the other approaches. Signals extracted from corresponding ROIs detected by Suite2p, CNMF, and SpecSeg were highly correlated. CNMF performed slightly better at denoising the signal than MLspike used with Suite2p or SpecSeg, probably because CNMF denoises during the same step as it identifies the ROIs. The differences were small, however, and future denoising approaches may further improve signal quality. Taken together, when comparing SpecSeg with other ROI-selection approaches, it is highly efficient at selecting ROIs from which a reliable signal can be extracted and selects few spurious ROIs.

We included a neuropil subtraction approach in SpecSeg that works by subtracting the signal from a donut-shaped area around the ROI from the ROI-derived signal, similarly to methods described previously (Chen et al., 2013; Khan et al., 2018; Tegtmeyer et al., 2018). We noticed that, for most SpecSeg ROIs, neuropil correction did not significantly improve the neuronal signal. More thorough analysis revealed that this was mainly due to the relatively high mean R^2 of SpecSeg ROIs: ROIs with high mean R^2 were much less contaminated by the neuropil signal than those with low mean R^2 . This observation also confirms that mean R^2 is a good measure for ROI quality. It should be noted that in ROIs with low mean R^2 , neuropil subtraction may actually introduce an artefact in the signal, again illustrating the importance of removing such ROIs from the analysis.

We also optimized the pipeline for the analysis of calcium-imaging data obtained using single-photon miniscopes. We found that SpecSeg identified many ROIs overlapping with those detected by CNMF-E, and that signals extracted from the ROIs selected by both methods were strongly correlated. However, different ROIs were also detected by both approaches, and many of these ROIs did actually show partial overlap and identified (parts of) the same neurons. This is probably caused by the fact that CNMF-E can create overlapping ROIs around neurons. The signal extraction procedure then separates these overlapping signals. In contrast, SpecSeg selects ROIs around those parts of the neurons that do not overlap and then extracts the signals from these ROIs. This approach does not deteriorate the quality of the extracted signal unless more than 50% of the ROI is removed. Removing noisy pixels from the ROI thus

appears to be highly effective and a good alternative to signal demixing, also for two-photon calcium imaging data. We expect that, in most cases, results using either method will be very similar. The main advantage of SpecSeg is its graphical user interface, which provides extensive insight into the activity profiles of the ROIs and facilitates quality control and post-hoc ROI selection and editing.

It is difficult to compare the speed of our pipeline with other available software packages. However, we found that motion correction using NoRMCorre was the slowest step in the process, taking up approximately half the time of the ROI selection process. This implies that the actual ROI selection process is not a bottleneck. Importantly, the speed of ROI selection was more than sufficient for all practical purposes except on-line ROI selection.

One strength of calcium imaging approaches is that one can follow responses of individual neurons over prolonged periods of time. To facilitate this approach, we developed a tool for automated matching of ROIs in series of sequentially recorded calcium imaging sessions. The percentage of ROIs that could be matched between two randomly selected pairs of recordings was quite variable, depending on various factors. In some recordings of a series, ROIs were missed due to a low mean R^2 during the session. ROIs could also be missed due to slight misalignments of the imaged brain region. Interestingly, we found that, when the time between two recordings was longer, the number of matched ROIs was reduced. Most likely, biological factors cause this reduction. The brain slightly changes its shape over time, causing loss of some neurons from the field of view. Moreover, neurons die, or lose or gain expression of the viral vector. The most interesting cause of missing ROIs in some recordings is that neurons change activity patterns over time, in line with the finding that, when mice perform the same task over prolonged periods of time, different sets of neurons are recruited (Driscoll et al., 2017).

In conclusion, SpecSeg is a powerful, complete, and open-source pipeline for ROI selection, signal extraction, and chronic ROI matching that can be used on a variety of single- and multi-photon calcium imaging data. Its main advantages over several existing calcium imaging toolboxes are the ease of use and simplicity, the intuitive way ROIs are selected and constrained, the selection of ROIs that represent neurons whose responses can be well separated from background noise, the possibility to select ROIs of various shapes and sizes, its graphical interface for ROI editing, and its use for analyzing both miniscope and multi-photon microscope calcium imaging data.

Limitations of the study

While SpecSeg processes datasets without a maximum file size, recordings do have to be longer than 2 min to calculate all low-frequency components. Also, SpecSeg cannot be used online during real-time data acquisition. In SpecSeg, ROI detection is based on correlated changes in fluorescence. If data contain a highly correlated fluctuating background signal it can be difficult to separate neurons from background and thus to extract useful ROIs. Such a highly correlated background signal may be diminished by background subtraction, which we provide for miniscope imaging. Similarly, it is difficult to detect inactive neurons

that nonetheless have a higher baseline fluorescence than background. The user can include such ROIs by selecting them semi-automatically using maximum or average fluorescence projections in the ROI manager user interface.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead contact
 - Materials availability
 - Data and code availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
 - Mouse experiments
- METHOD DETAILS
 - Viral injections for two-photon microscopy experiments
 - Handling and habituation
 - Cranial window surgery for two-photon microscopy
 - Two-photon imaging
 - Visual stimulation during two-photon microscopy
 - Gradient-index lens surgery
 - Baseplating and miniscope imaging
 - Miniscope recording and preprocessing
 - NAOMI simulated data
 - SpecSeg data analysis code
 - ROI selection based on cross-spectral power
 - Miniscope background correction
 - User interface for ROI refinement
 - Signal extraction and neuropil subtraction
 - Spike estimation
 - Identification of the same ROIs in chronically recorded datasets
 - Quantification and statistical analysis

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.crmeth.2022.100299>.

ACKNOWLEDGMENTS

The authors thank Emma Ruimschotel for technical assistance and the Mechatronics department and Animal facilities of the NIN for their services. This project received funding from the European Union's Horizon 2020 Research and Innovation Programme under grant agreement nos. 785907 (HBP SGA2) and 945539 (HBP SGA3). We thank the Genetically Encoded Neuronal Indicator and Effector (GENIE) Project and the Janelia Research Campus of the Howard Hughes Medical Institute for developing and sharing GCaMP6s and Dr. Karl Deisseroth for providing AAVdj.CaMKIIa.GCaMP6s and AAVdj.hSyn.GCaMP6s vectors.

AUTHOR CONTRIBUTIONS

L.d.K. wrote code, analyzed data, and wrote the paper. K.S., P.T., and B.J.G.v.d.B. performed the experiments, analyzed data, and wrote the paper. I.F.P. analyzed data. I.W. provided materials and data for analysis. C.N.L.

supervised the project and wrote the paper. C.v.d.T. wrote code, supervised the project, and wrote the paper.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: April 9, 2021
Revised: February 11, 2022
Accepted: August 29, 2022
Published: September 20, 2022

REFERENCES

- Aharoni, D., Khakh, B.S., Silva, A.J., and Golshani, P. (2019). All the light that we can see: a new era in miniaturized microscopy. *Nat. Methods* 16, 11–13.
- Apthorpe, N.J., Riordan, A.J., Aguilár, R.E., Homann, J., Gu, Y., Tank, D.W., and Seung, H.S. (2016). Automatic neuron detection in calcium imaging data using convolutional networks. *Adv. Neural Inf. Process. Syst.*, 3278–3286.
- Cai, D.J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Song, W., Wei, B., Veshkini, M., La-Vu, M., Lou, J., et al. (2016). A shared neural ensemble links distinct contextual memories encoded close in time. *Nature* 534, 115–118.
- Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295–300.
- Cichon, J., and Gan, W.B. (2015). Branch-specific dendritic Ca²⁺ spikes cause persistent synaptic plasticity. *Nature* 520, 180–185.
- Dana, H., Sun, Y., Mohar, B., Hulse, B.K., Kerlin, A.M., Hasseman, J.P., Tsegaye, G., Tsang, A., Wong, A., Patel, R., et al. (2019). High-performance calcium sensors for imaging activity in neuronal populations and microcompartments. *Nat. Methods* 16, 649–657.
- de Groot, A., van den Boom, B.J., van Genderen, R.M., Coppens, J., van Veldhuijzen, J., Bos, J., Hoedemaker, H., Negrello, M., Willuhn, I., De Zeeuw, C.I., and Hoogland, T.M. (2020). Ninscope, a versatile miniscope for multi-region circuit investigations. *Elife* 9, e49987.
- Deneux, T., Kaszas, A., Szalay, G., Katona, G., Lakner, T., Grinvald, A., Rózsa, B., and Vanzetta, I. (2016). Accurate spike estimation from noisy calcium signals for ultrafast three-dimensional imaging of large neuronal populations *in vivo*. *Nat. Commun.* 7, 12190.
- Driscoll, L.N., Pettit, N.L., Minderer, M., Chettih, S.N., and Harvey, C.D. (2017). Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell* 170, 986–999.e16.
- Engert, F., and Bonhoeffer, T. (1999). Dendritic spine changes associated with hippocampal long-term synaptic plasticity. *Nature* 399, 66–70.
- Gambino, F., Pagès, S., Kehayas, V., Baptista, D., Tatti, R., Carleton, A., and Holtmaat, A. (2014). Sensory-evoked LTP driven by dendritic plateau potentials *in vivo*. *Nature* 515, 116–119.
- Ghosh, K.K., Burns, L.D., Cocker, E.D., Nimmerjahn, A., Ziv, Y., Gamal, A.E., and Schnitzer, M.J. (2011). Miniaturized integration of a fluorescence microscope. *Nat. Methods* 8, 871–878.
- Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B.L., Koay, S.A., Taxis, J., Najafi, F., Gauthier, J.L., Zhou, P., et al. (2019). CalmAn: an open source tool for scalable calcium imaging data analysis. *Elife* 8, e38173.
- Iacuruso, M.F., Gasler, I.T., and Hofer, S.B. (2017). Synaptic organization of visual space in primary visual cortex. *Nature* 547, 449–452.
- Jaepel, J., Hübener, M., Bonhoeffer, T., and Rose, T. (2017). Lateral geniculate neurons projecting to primary visual cortex show ocular dominance plasticity in adult mice. *Nat. Neurosci.* 20, 1708–1714.
- Jia, H., Rochefort, N.L., Chen, X., and Konnerth, A. (2010). Dendritic organization of sensory input to cortical neurons *in vivo*. *Nature* 464, 1307–1312.
- Jo, J.A., Fang, Q., Papaioannou, T., and Marcu, L. (2004). Fast Model-free Deconvolution of Fluorescence Decay for Analysis of Biological Systems 9, 743–752.
- Kaifosh, P., Zaremba, J.D., Danielson, N.B., and Losonczy, A. (2014). SIMA: Python software for analysis of dynamic fluorescence imaging data. *Front. Neuroinform.* 8, 80.
- Keemink, S.W., Lowe, S.C., Pákan, J.M.P., Dylida, E., Van Rossum, M.C.W., and Rochefort, N.L. (2018). FISSA: a neuropil decontamination toolbox for calcium imaging signals. *Sci. Rep.* 8, 3493.
- Khan, A.G., Poort, J., Chadwick, A., Blot, A., Sahani, M., Mrsic-Flogel, T.D., and Hofer, S.B. (2018). Distinct learning-induced changes in stimulus selectivity and interactions of GABAergic interneuron classes in visual cortex. *Nat. Neurosci.* 21, 851–859.
- Klibisz, A., Rose, D., Eicholtz, M., Blundon, J., and Zakharenko, S. (2017). Fast, simple calcium imaging segmentation with fully convolutional networks. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support* (Springer Verlag), pp. 285–293.
- Liberti, W.A., Perkins, L.N., Leman, D.P., and Gardner, T.J. (2017). An open source, wireless capable miniature microscope system. *J. Neural. Eng.* 14, 045001.
- Madisen, L., Garner, A.R., Shimaoka, D., Chuong, A.S., Klapoetke, N.C., Li, L., van der Bourg, A., Niino, Y., Egolf, L., Monetti, C., et al. (2015). Transgenic mice for intersectional targeting of neural sensors and effectors with high specificity and performance. *Neuron* 85, 942–958.
- Mank, M., Santos, A.F., Drenberger, S., Mrsic-Flogel, T.D., Hofer, S.B., Stein, V., Hendel, T., Reiff, D.F., Levelt, C., Borst, A., et al. (2008). A genetically encoded calcium indicator for chronic *in vivo* two-photon imaging. *Nat. Methods* 5, 805–811.
- Maruyama, R., Maeda, K., Moroda, H., Kato, I., Inoue, M., Miyakawa, H., and Aonishi, T. (2014). Detecting cells using non-negative matrix factorization on calcium imaging data. *Neural Netw.* 55, 11–19.
- Masullo, A. (2020). `xcorr2_fft(ab)` (MathWorks).
- Mishne, G., Coifman, R.R., Lavzin, M., and Schiller, J. (2018). Automated cellular structure extraction in biological images with applications to calcium imaging data. Preprint at bioRxiv, 313981.
- Mukamel, E.A., Nimmerjahn, A., and Schnitzer, M.J. (2009). Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron* 63, 747–760.
- Pachitariu, M., Stringer, C., Dipoppa, M., Schröder, S., Rossi, L.F., Dalgleish, H., Carandini, M., and Harris, K.D. (2017). Suite2p: beyond 10,000 neurons with standard two-photon microscopy. Preprint at bioRxiv, 061507.
- Petersen, A., Simon, N., and Witten, D. (2018). Scalpel: extracting neurons from calcium imaging data. *Ann. Appl. Stat.* 12, 2430–2456.
- Petreanu, L., Mao, T., Sternson, S.M., and Svoboda, K. (2009). The subcellular organization of neocortical excitatory connections. *Nature* 457, 1142–1145.
- Pnevmatikakis, E.A., and Giovannucci, A. (2017). NoRMCorr: an online algorithm for piecewise rigid motion correction of calcium imaging data. *J. Neurosci. Methods* 291, 83–94.
- Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., et al. (2016). Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron* 89, 285–299.
- Resendez, S.L., Jennings, J.H., Ung, R.L., Namboodiri, V.M.K., Zhou, Z.C., Otis, J.M., Nomura, H., Mchenry, J.A., Kosyk, O., and Stuber, G.D. (2016). Visualization of cortical, subcortical and deep brain neural circuit dynamics during naturalistic mammalian behavior with head-mounted microscopes and chronically implanted lenses. *Nat. Protoc.* 11, 566–597.
- Reynolds, S., Abrahamsson, T., Schuck, R., Sjöström, P.J., Schultz, S.R., and Dragotti, P.L. (2017). ABLE: An activity-based level set segmentation algorithm for two-photon calcium imaging data. *ENeuro* 4, ENEURO.0012-0017.2017.
- Sawtell, N.B., Frenkel, M.Y., Philpot, B.D., Nakazawa, K., Tonegawa, S., and Bear, M.F. (2003). NMDA receptor-dependent ocular dominance plasticity in adult visual cortex. *Neuron* 38, 977–985.
- Smith, S.L., and Häusser, M. (2010). Parallel processing of visual space by neighboring neurons in mouse visual cortex. *Nat. Neurosci.* 13, 1144–1149.

Song, A., Gauthier, J.L., Pillow, J.W., Tank, D.W., and Charles, A.S. (2021). Neural anatomy and optical microscopy (NAOMi) simulation for evaluating calcium imaging methods. *J. Neurosci. Methods* 358, 109173.

Spaen, Q., Asín-Achá, R., Chettih, S.N., Minderer, M., Harvey, C., and Hochbaum, D.S. (2019). HNCcorr: a novel combinatorial approach for cell identification in calcium-imaging movies. *ENeuro* 6. ENEURO.0304-18.2019.

Svoboda, K., Tank, D.W., and Denk, W. (1996). Direct measurement of coupling between dendritic spines and shafts. *Science* 272, 716–719.

Szalay, G., Judák, L., Katona, G., Ócsai, K., Juhász, G., Veress, M., Szadai, Z., Fehér, A., Tompa, T., Chiovini, B., et al. (2016). Fast 3D imaging of spine, dendritic, and neuronal assemblies in behaving animals. *Neuron* 92, 723–738.

Tegtmeier, J., Brosch, M., Janitzky, K., Heinze, H.-J., Ohl, F.W., and Lippert, M.T. (2018). CAVE: an open-source tool for combined analysis of head-mounted calcium imaging and behavior in MATLAB. *Front. Neurosci.* 12, 958.

Wilson, D.E., Whitney, D.E., Scholl, B., and Fitzpatrick, D. (2016). Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. *Nat. Neurosci.* 19, 1003–1009.

Winnubst, J., Cheyne, J.E., Niculescu, D., and Lohmann, C. (2015). Spontaneous activity drives local synaptic plasticity *in vivo*. *Neuron* 87, 399–410.

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

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and virus strains		
Viral vector AAV2/9.syn.GCaMP6f	UPenn Vector Core	N/A
Viral vector AAVdj.CaMKIIa.GCaMP6s	Stanford Neuroscience Gene Vector and Virus Core	N/A
Chemicals, peptides, and recombinant proteins		
Isoflurane	Zoetis Isoflo	REG NL 10416 UDD
Metacam	Boehringer Ingelheim	EU/2/97/004/003-005,029
Dexamethasone	Alfasan	REG NL 1578 UDD
Cavasan eye ointment	AST Farma	REG NL 4006
Experimental models: Organisms/strains		
Mice C57BL/6	Janvier Labs	N/A
Ai14 reporter mice	Jackson Labs	#007908
VIP-cre mice	Jackson Labs	#010908
SOM-cre mice	Jackson Labs	#013044
G35-3 Cre mice	Sawtell et al., 2003	N/A
Ai93 GCaMP6f mice	Jackson Labs	#024103
Software and algorithms		
MATLAB	Mathworks™	N/A
NoRMCorre	Pnevmatikakis and Giovannucci, 2017	https://github.com/flatironinstitute/NoRMCorre
NAOMi	Song et al. 2021	https://bitbucket.org/adamshch/naomi_sim/src/master/
CNMF-E	Zhou et al. 2018	https://github.com/zhoup/cnmf_e
SpecSeg	This paper	https://doi.org/10.5281/zenodo.6993003
Suite2P	Pachitariu et al., 2017	https://github.com/MouseLand/suite2p
Other		
NeuroLabware standard two-photon microscope	NeuroLabware	N/A
GRIN lens (1 mm diameter, 0.5 mm working distance)	Grintech	N/A
Miniscope v3.2	Cai et al., 2016 ; Ghosh et al., 2011	N/A

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, Dr. Christiaan N. Levelt (c.levelt@nin.knaw.nl).

Materials availability

- This study did not generate new unique reagents.

Data and code availability

- Calcium imaging data reported in this paper will be shared by the [lead contact](#) upon request.
- All original code has been deposited at Github and is publicly available as of the date of publication and can be downloaded at <https://github.com/Leveltlab/SpectralSegmentation>. An archival version is also available at <https://doi.org/10.5281/zenodo.6993003>
- Any additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Mouse experiments

To develop the ROI selection tools described in this study, we made use of two-photon imaging data from male and female mice that were repeatedly imaged in as yet unpublished behavioral studies. Here we describe the methodology developed for the analysis of such data. All animal experiments were approved by the institutional animal care and use committees of the Royal Netherlands Academy of Arts and Sciences. We used male and female mice that were 2–7 months of age. The mice were C57Bl/6 mice, or offspring of Ai14 mice (Cre-dependent tdTomato reporter mice, strain 007908) crossed with mice expressing Cre in vasoactive intestinal polypeptide (VIP)-expressing or somatostatin (SOM)-expressing interneurons (Jackson Laboratories, www.jaxmice.jax.org, strains 010908 and 013044 respectively). However, we did not make use of the calcium signals selectively derived from these interneuron subsets for this study. We also made use of two-photon imaging data acquired in V1 of Ai93D (cre-dependent GCaMP6f-expressing strain 024103, Jackson Laboratories) (Madisen et al., 2015) x G35-cre (Sawtell et al., 2003) mice, expressing GCaMP6f in most excitatory neurons of the cortex. These strains were kept on a mixed C57Bl/6 x 129S background. All animals were kept in a 12 h reverse day/night cycle with access to food and water *ad libitum*. Experiments were carried out during the dark cycle.

METHOD DETAILS

Viral injections for two-photon microscopy experiments

Mice were injected with a viral vector driving expression of the genetically encoded calcium sensor GCaMP6f in neurons (AAV2/9.syn.GCaMP6f, UPenn Vector Core facility). Anesthesia was induced with 5% isoflurane and maintained at 1.6% isoflurane in Oxygen (0.8 L/min flow rate). Mice were administered Metacam (1 mg/kg subcutaneously (s.c.), for analgesia) and dexamethasone (8 mg/kg s.c.) to prevent cerebral edema/inflammation after induction of anesthesia. Mice were head-fixed on a stereotax, scalp and soft tissue overlying the visual cortex were incised and the skull exposed. A small hole was drilled in the skull overlying the center of primary visual cortex (V1). A pulled capillary with AAV2/9.syn.GCaMP6f or AAVdj.CaMKIIa.GCaMP6s was inserted vertically through this hole to a depth of 200–400 μ m from the brain surface. Approximately 20 to 100 nL of virus (titer $\sim 10^{12}$ viral genomes per mL) was injected slowly using a Nanoject and the hole was covered with bone wax. During the surgery, the temperature was maintained with euthermic pads. Respiration was monitored to adjust depth of anesthesia. Eyes were protected from light and from drying using Cavasan eye ointment. Once the window was made the exposed dura was continuously kept moist with artificial aCSF, consisting of a solution of 125 NaCl, 10 HEPES, 5 KCl, 2 MgSO₄, 2 CaCl₂, and 10 Glucose, in mM. Later the scalp was sutured, and the animal let to recover from anesthesia.

Handling and habituation

Once an animal recovered from the viral injection, and before window implantation, animals were handled daily for 5 min (or until they started to groom while being handheld) to reduce handling stress during later training. Next, animals were trained and habituated for 3 days with head-restraint in the training setup with a running wheel. After habituation, animals were placed in a two-photon microscopy setup. Once the mice were comfortable with the setup, they performed a visual detection task.

Cranial window surgery for two-photon microscopy

One month after viral vector injection, mice were anesthetized again as described above. Mice were head-fixed on a stereotax and scalp and soft tissue overlying the visual cortex were incised and the skull exposed. A metal ring (5 mm inner diameter) was fixed on the skull centered on V1, with dental cement. A cranial window was made inside the ring and the dura was exposed. The cranial window was then covered with a double coverslip (to reduce brain movements under the microscope) and fixed to the metal ring using dental cement. Animals were allowed to recover after the dental cement dried. After a minimum of 2 weeks of recovery, mice were submitted to further handling and training. During the training and recording periods described below, animals were typically in the setup 5 days per week.

Two-photon imaging

For imaging we used a NeuroLabware standard microscope (CA, USA) equipped with a Ti-sapphire laser (Mai-Tai, Spectra-physics, CA, USA). A black cloth was used to cover the objective in order to prevent light coming from the monitor to the objective. Two-photon laser scanning microscopy was performed at 920 nm and neurons were imaged with 16x water-immersion objective (0.8NA) with computer-optimized optics of 1.6x magnification.

Two-photon calcium-imaging sequences were recorded in awake behaving mice and saved in a continuous binary format (sbx, NeuroLabware). The dimensions of the images were 812 by 512 pixels, with 16 bit unsigned integer pixel depth. These files are associated with a metadata file (mat, MATLAB) that defines, among other parameters, pixel dimensions, number of channels, number of sections, and number of frames recorded. For two-photon GRIN lens imaging in the visual thalamus, we used a 4x objective (0.2NA) with computer-optimized optics with 3.2 magnification. All further processing was done with MATLAB (MathworksTM).

Visual stimulation during two-photon microscopy

Stimuli were presented on a gamma-corrected Dell-P2314H 23" full HD LED monitor, placed 15 cm in front of the mouse. Stimuli were made with custom-made MATLAB scripts. Receptive field size of each region of interest (ROI) was estimated by reverse correlation after presenting 3 black and 3 white squares (7.5° - size degrees) simultaneously at pseudo random locations on the screen. The stimuli were repeated for 10–15 times at each location. The duration was 0.5 s, the interstimulus was an isoluminant gray screen, duration 1.5 s, contrast: 1.0 and maximum luminance was kept to 20% of screen max (max luminance). For visual stimulation, various stimuli were used, depending on the behavioral experiment.

Gradient-index lens surgery

For single-photon miniscope experiments, animals were anaesthetized with isoflurane (1.5% in O₂ and air) and placed in a stereotaxic apparatus (David Kopf Instruments) on a heated pad (37°C), injected with Metacam (1 mg/kg subcutaneously) as analgesic, dexamethasone (8 mg/kg subcutaneously) to prevent inflammation, and saline (80 mL/kg subcutaneously) to prevent dehydration. A craniotomy and durotomy were made above the region of interest. To image calcium dynamics in premotor cortex, a total of 500 nL AAVdj.CaMKIIa.GCaMP6s (Titer max. 10E12 genomes per mL, Stanford Neuroscience Gene Vector and Virus Core) was injected at four locations 100 μm off center of premotor cortex (AP:2.3 ML:0.35 DV: -0.3) using a Hamilton-syringe (100 nL/min). A GRIN objective lens (1.8 mm diameter, Edmund Optics Ltd.) was placed on top of the brain. To image striatum, a custom-built robotized stereotaxic arm was used to slowly lower (300 nm/min) a 25G needle to ease lens implantation and reduce tissue damage above the ventral striatum (AP:1.1 ML:1.1 DV:-4.5). A total of 500 nL AAVdj.hSyn.GCaMP6s was injected at two different location 100 μm off center of ventral striatum (AP:1.1 ML:1.1 DV:-4.5), followed by lowering (100 nm/min) of a GRIN relay lens (0.6 mm diameter, Inscopix Inc.). The gap between lens and skull was covered with cyanoacrylate glue and secured to the skull using Superbond dental cement (Sun Medical Company Ltd.). Using dental acrylic cement, a head bar was placed further caudal and skin was glued to the cement headcap. The lens was protected by a layer of Twinsil Speed (Picodent GmbH) silicon.

For two-photon imaging in the visual thalamus, animals were anesthetized as described above and injected with a total of 100 nL AAVdj-hSyn.GCaMP6s (Titer ~10E12 genomes per mL, Stanford Neuroscience Gene Vector and Virus Core) at two locations (AP:2.2 & 2.4, ML:2.1) and three depths per location (DV:2.35, 10 nL; 2.5; 20 nL, 2.7; 20 nL). A doublet GRIN lens (1 mm diameter, 0.5 mm working distance, GRINTECH) was implanted right above the visual thalamus as described above for striatum, centered in between the two injection locations (AP:2.3, ML:2.1, DV:2.1) and fixed on the skull using Superbond dental cement. A custom-made head restriction ring was placed around the GRIN lens and fixed on the skull using Superbond and Tetric EvoFlow® (A1). The GRIN lens was protected by a custom-made 3D-printed cap that could be screwed on the head restriction ring.

Baseplating and miniscope imaging

Four to eight weeks after virus injection and GRIN lens placement, mice were head restricted using a custom-built device (that included a running belt to allow for locomotion) and fluorescent signal was assessed. The miniscope (Miniscope v3.2) (Cai et al., 2016; Ghosh et al., 2011) was lowered above the implanted GRIN lens and, once individual cells were detectable in the field of view, a baseplate was secured to the cement head cap using dental acrylic cement (coated with black nail polish). Recording imaging sessions were performed in an open-field arena (30 cm × 30 cm). Prior to each imaging session, mice were head restricted, and the miniscope was attached to the baseplate. Calcium dynamics were recorded with a frame-rate acquisition of 15 Hz, and data was stored in videos files with a resolution of 752 (width) × 480 (height) pixels (AVI format).

Miniscope recording and preprocessing

Raw AVI files were spatially down sampled by a factor four (reducing frames to 376 × 240 pixels) and stored as TIFF files. Rigid and non-rigid motion in imaging data was corrected using NoRMCorre (Pnevmatikakis and Giovannucci, 2017), followed by neuron and signal extraction using either SpecSeg or CNMF-E (Zhou et al., 2018). Every cell included in the analyses was verified by visual inspection. CNMF-E output was manually cleaned using a custom-written GUI in MATLAB (2019b). Components with non-circular contours, spatial contours smaller than neuron size, or artifacts in temporal traces were excluded. Duplicate copies of individual cells were identified by overlapping contours and temporal activity was subsequently merged.

NAOMi simulated data

For comparing SpecSeg, CNMF and Suite2p we made use of previously published data produced by NAOMi (Song et al., 2021). The NAOMi dataset consisted of 20,000 frames simulated from a 500 × 500 × 100 μm volume with 1 μm sampling at 30 Hz scanning using a 0.6-NA Gaussian excitation numerical aperture (NA) at 40 mW average power. The ROI sets for Suite2p and CNMF were also published (Song et al., 2021). The ROI sets were created in the MATLAB 2017 versions of the software. Parameters were adjusted manually for each algorithm to optimize the output.

For CNMF the following parameters were used: fr = 30, tsub = 5, patch_size = [40, 40], overlap = [8,8], K = 7, tau = 6, p = 0, and num_bg = 1. For Suite2P the following parameters were used: diameter = 12, DeleteBin = 1, sig = 0.5, nSVDforROI = 1000, NavgFramesSVD = 5,000, signalExtraction = 'surround', innerNeuropil = 1, outerNeuropil = Inf, minNeuropilPixels = 400, ratioNeuropil = 5, imageRate = 30, sensorTau = 0.5, maxNeuropil = 1, sensorTau = 0.5, and redmax = 1. For SpecSeg ROIs we used cutOffHz-Min = 0, cutOffHzMax = 0.44, areasz = [45, 300], roundedness = 0.666, voxel = 50, cutOffCor = 0.05, useFluorescenceImg = true.

To calculate which ROIs matched with NAOMi, the ROI masks of each method were flattened into a 2D mask, like SpecSeg uses. The ROIs were matched to the NAOMi ground truth somas that were present in the volume present 8 μm around the focal plane. To match to a NAOMi ROI, the pixel overlap between ROIs had to be more than 45%.

SpecSeg data analysis code

The code for all data analysis procedures described below is available at Github: (<https://github.com/Leveltlab/SpectralSegmentation>). A flowchart of the analysis pipeline is shown in Figure 1 and is also present in the Github repository.

ROI selection based on cross-spectral power

We developed a pipeline for ROI extraction that works as follows: First, the images are cropped to remove border artifacts before being aligned using rigid registration with NoRMCorre (<https://github.com/flaticoninstitute/NoRMCorre>), a toolbox provided by the Simons Foundation (Pnevmatikakis and Giovannucci, 2017). We adapted the entry function of this toolbox in order for it to work with files in .sbn format, produced by the NeuroLabware microscope that we used, and integrated it in the pipeline. After visual inspection to ascertain that the images are well aligned, the image sequences are transposed to place time in the first dimension and width * height in the second dimension (StackTranspose.m). This makes the processing of pixel traces much more efficient in the time domain. This data is saved in binary format for later use with “_Trans.dat” as extension. In addition, we also save a temporally downsampled version of this file with “_DecTrans.dat” as extension. Temporal downsampling is performed with “decimate” (MATLAB) to a sampling rate of ± 1 Hz.

Next, the cross-spectral power of the fluorescent signal between neighboring pixels over time is calculated. To achieve this, the data is first detrended and for each half-overlapping time window (60 s) the data is convolved with a hamming window. Cross-spectral density functions are calculated from the discrete Fourier transforms of these pixel-trace segments for each pixel with its eight neighbors and averaged over all time windows. Additionally, the total variance is calculated for each pixel from their average auto-spectral density function. Then, the cross-spectral power functions are normalized with their respective variances according to the formula:

$$P_{xy}(f) = \frac{|G_{xy}(f)|^2}{V_x V_y}$$

Where $P_{xy}(f)$ is the average normalized cross-spectral power, $G_{xy}(f)$ is the cross spectral density between x and y for frequency component f, and V_x and V_y the average variance of x and y respectively.

Finally, the normalized cross-spectral power functions of a pixel with its eight neighbors are averaged. The result is a 2D matrix representing cross-spectral power at each frequency component for all pixels. We used this to generate an image for each spectral component. This data is saved in a separate file with “SPSIG.mat” as extension.

Cross-spectral images are generated for spectral components between 0.017 Hz to 0.5 Hz, from image sequences recorded at a sampling rate between 10 and 30 Hz (Figure 2).

Given a general model of calcium responses to spikes characterized as a convolution between the fluorescence impulse response function and spikes as a series of onset times (Jo et al., 2004):

$$\begin{aligned} c(t) &= h(t) * s(t) + e(t) \\ &= \sum_{n=1}^N h(t - \tau_n) + e(t) \end{aligned}$$

where $c(t)$ is the calcium trace, $s(t)$ is the spiking events, $h(t)$ is the calcium response and $e(t)$ is the independent noise, and τ_n is the spike time, the Fourier transform is:

$$\begin{aligned} \mathcal{F}[c(t)](\omega) &= \sum_{n=1}^N \mathcal{F}[h(t - \tau_n)](\omega) + E(\omega) \\ &= \sum_{n=1}^N e^{-j\omega\tau_n} H(\omega) + E(\omega) \\ &= H(\omega) \left[\sum_{n=1}^N e^{-j\omega\tau_n} \right] + E(\omega) \end{aligned}$$

where $E(\omega)$ is independent noise in frequency ω , and $H(\omega)$ is the Fourier transform of the calcium response. Since the calcium response: $h(t) = e^{-at}$, approximated a one-sided decaying exponential function, it follows that the Fourier transform $H(\omega)$ is $1/(j\omega + a)$,

Magnitude, $|H(\omega)| = \frac{1}{\sqrt{a^2 + \omega^2}}$; for all ω

The number of spikes and their position in time is encoded in the height and phase of the Fourier transform.

For all calcium indicators the decay half time has been shown to be in the order of 400 ms (Chen et al., 2013), this approximates to a time constant of $a = -\ln(0.5)/0.4 \approx 1.73$ 1/s, since $0.5 = e^{-at}$, at $t = 0.4$ s. The power spectrum of an exponential function based on this time constant can be shown to contain the highest power in frequency components below 1 Hz (Figure 2G).

This explains the observation that pixels within active neurons display strong cross-spectral power below 0.4 Hz. Background neuropil is usually non correlated making neurons clearly visible in these images (Figures 2 and 3). The images for cross-spectral components between 0.017 and 0.4 Hz are therefore used to extract an initial estimation of the ROIs (spectral.m).

For each cross-spectral image in the selected frequency range, a series of morphological constraints are applied to find as many non-overlapping ROIs as possible (getSpectrois.m). First, all maxima (cross-spectral peaks) in the image are detected and sorted according to magnitude. A fraction (0.25–0.4) of these peaks with the highest values are selected and sorted in decreasing order. Based on this selection, a square area of pixels (Voxel; 50x50 pixels) is sampled for each peak, centered on the peak and contours are detected using contourc (MATLAB) (Figures 2L and 2N). Closed contours containing the selected peak are selected and constrained by a set of criteria:

1. To avoid that multiple peaks from the same neuron are selected, the minimal distance between peaks in each ROI should be greater than a defined threshold (20 pixels).
2. To ensure that the magnitude of the peak is well above background level, the peak should be greater than (95%) of the pixel range within a voxel.
3. Given the expected size range of cell bodies, the minimum and maximum area (number of pixels) should be between 40 and 400 pixels (a pixel is $\pm 1.5 \mu\text{m}^2$ within our field of view: $\pm 1 \text{ mm}$).
4. Roundedness; The relationship between the area of a contour and the length of its circumference ($4 \cdot \text{Pi} \cdot \text{Area} / \text{Circumference}^2$) should be greater than 0.6 (1 = completely round, 0 = any shape), These parameters can be adapted depending on the type of ROI that a user is interested in and the magnification of the images. To select cell bodies, for example, the roundedness should be greater than 0.6 and the number of pixels should be within a constrained range based on the expected size of the cell bodies. Because the criteria are simple and self-evident, optimal values can easily be established with some experimentation.

Depending on the density of active neurons and their processes in an image sequence, contours around correlated pixels may still represent overlapping or adjacent neural elements. Therefore, the ROIs are further restrained based on the assumption that pixels in adjacent cell bodies are not highly correlated. To achieve this, all pixel traces within the original ROI are correlated with the averaged down-sampled trace of eight neighboring pixels at the original local maximum. Based on this calculation, the ROI is constrained to an area containing pixels with correlations over a threshold set halfway the maximum correlation strength and the correlation strength in a band of tissue surrounding the ROI (this threshold can be adapted) (Figures 2L and 2 N). In principle, this is a computationally expensive approach, but it is feasible because the numbers of pixel traces in a preselected ROI are limited in number and the traces are decimated to 1 Hz. For analysis and selection purposes we also saved the mean R^2 of an ROI. To calculate this we averaged the squared pixel correlations within the final ROI.

Miniscope background correction

We included a background correction tool (BackgroundSubtractSbx.m) for miniscope calcium imaging data, which is applied after the motion correction step. A unique background estimation image is created for every frame by filtering a padded version of that frame. This background estimation is subtracted from the frame. Before subtraction, the values of the background estimation are lowered by a fixed amount to prevent underexposure of the data in the final result. There are three different filter kernels to choose from: a Gaussian, a disk, or a donut. A disk of 111x111 pixels ($\sim 200\mu\text{m}$) was chosen in the example data (Figures 3E, 3F and 5). The gap in the donut filter prevents the background estimation from containing signal of a neuron at the position of the filter. We found, however, that the difference in results compared to a circular filter was very small. A Gaussian filter provides very similar results to a smaller disk filter, and much of the miniscope data is already out of focus in a way that resembles a Gaussian kernel. This is why the circular kernel is preferred over the Gaussian. The size of the filter can be edited by the user, based on the data's scale. When selecting the donut-shaped filter, both the gap size and outer radius of the filter can be adjusted. Background subtraction takes more processing time to calculate with larger filters or larger datasets. To diminish processing time, the background can be calculated using a downscaled version of each frame. An optional feature in the background subtraction process is data smoothing. To decrease sensor banding noise and increase correlation between neighboring pixels, data can be smoothed with a Gaussian in either the vertical, horizontal or both directions. When the correlation between neighboring pixels is increased, it becomes easier to detect ROIs with low mean R^2 . This smoothing was not used in the example data. All filtering parameters can first be tested on a small subset of the data with the BackgroundSubtractSbxExampleFrame.m script.

User interface for ROI refinement

To give users control over the selected ROIs, we developed a graphical user interface to reject ROIs outside a preferred range of properties (RoiManagerGUI.m). Additionally, the user interface allows the user to manually delete or keep ROIs. In rare cases, some ROIs have areas with pixels that have low or negative correlations. This may indicate that the ROI contains signals from multiple neuronal sources, or that registration was suboptimal and the neuron was not located at the ROI during the entire recording. We therefore developed the option to split these ROIs (Figures S4B–S4E). The user interface creates four reference points at the distal edges of the ROI. The fluorescence signal from these reference points is then correlated with the signal of each pixel in the ROI. This creates four correlation values per pixel of the ROI. These are used to create a color-coded image in the user interface, indicating the

location of different signal sources in the ROI. The pixels in the ROI can then be subjected to k-means clustering. After setting the number of clusters the user can decide to split such an ROI in two or more ROIs or to delete part of the ROI. The user can also add more reference points if clustering needs to be improved.

We also included the option to create new ROIs in the user interface, for example when neurons were missed by the automated analysis. These new ROIs can be added in two different ways. One option is to manually draw the contour of the ROI. The other option is to let the computer draw the ROI, based on the cross-spectral image (or any other user-defined background image, such as the fluorescent image) and a threshold set by the user (Figure S4A).

Signal extraction and neuropil subtraction

To extract the calcium signals for each ROI (`retrievesignals.m`), all pixel traces within the ROI contour are averaged. The signals are neuropil corrected by subtracting the averaged signal of a neuropil area from the ROI signal. A donut-shaped neuropil area is created for each ROI. To achieve this, a small buffer area with a width of 2 pixels ($\sim 2 \mu\text{m}$) is first added around each ROI. The neuropil area is defined by first enlarging the ROI by 20 pixels ($\sim 20 \mu\text{m}$) using a circular filter, after which the central and adjacent ROIs and their buffer areas are excluded from this area. To prevent overcorrection, the neuropil signal is multiplied by 0.7, in accordance with previous studies (Chen et al., 2013; Khan et al., 2018; Tegtmeier et al., 2018).

Spike estimation

The extracted ROI signal can be converted into an estimate of a spike train (`DeconvolveSignals.m`) with the `MLspike` toolbox (Deneux et al., 2016), which has to be downloaded via their github link (<https://github.com/MLspike>). The spike train is then saved in the same format as the regular calcium signals.

Identification of the same ROIs in chronically recorded datasets

Finally, we include a toolbox to identify the same ROIs in chronic recordings of the same brain region (`ChronicMatching.m`). It first registers the spectral images of the recordings to ensure that each ROI will be located at exactly the same position in all recordings. To achieve this, the spectral images are normalized so that all values are in the range from -1 to 1 . Next, the registration of the spectral images is done by repeated translation and rotation. Each recording is registered to a reference recording chosen by the user. Preferentially, one of the middle recordings is used as a reference in order to minimize the difference between the reference and the recordings to be registered.

The distance over which recordings need to be translated is calculated by doing a 2D cross-correlation between the reference image and the spectral images of the other recordings. The cross-correlation is calculated with `xcorr2_fft` (Masullo, 2020), which is much faster than the MATLAB built-in cross-correlation function. After the translations are applied, the images are padded with zeros to maintain the same dimensions between recordings.

To correct the rotation, each recording is rotated using nearest-neighbor interpolation over a range of different angles, from 1° clockwise to 1° counterclockwise in steps of 0.05° . Each rotation is compared to the reference image with a 2D correlation and the best fit is applied. Several rounds of translation and rotation can be applied until no significant further improvement is achieved.

After registration, ROIs between all chronic recordings within an experiment are matched. A minimum percentage of overlap threshold is set. In the example presented here, we chose 67.5% based on the best balance between false positives and false negatives as determined by inspection with the chronic viewer user interface. The ROI matching is represented in a 2D matrix (the “match matrix”), consisting of a column for each recording and rows with the ROI numbers in each recording that are matched to each other. If no match for a particular ROI is found in a recording, the cell is kept blank.

To create the match matrix, the following three steps are taken:

1. For each ROI in each recording, the overlapping ROIs from all other recordings are identified. The percentage of the reference ROI that is covered by each overlapping ROI of other recordings is calculated (“overlap”) and saved as a putative match if it is above the overlap threshold.
2. Next, the overlap for each pair of putatively matched ROIs is calculated in the opposite direction (i.e. the percentage of overlapping ROI that is covered by the reference ROI). The two overlap values are averaged, and putative matches whose average overlap is below the overlap threshold (in our example 67.5%) are discarded. Calculating the overlap in these two simple steps automatically takes many parameters of the neurons into account: if the sizes, shapes or positions of two ROIs differ significantly the overlap value will always be small.
3. To obtain the final matching matrix, the data needs to be merged in order to match the same ROIs in all recordings, not just in pairs of recordings. The merging process includes ‘mutual linking’: if an ROI from recording A is matched only to recording B, but that ROI from recording B has matches to A and C, it is implied that the ROI from recording A is also matched to the ROI in recording C. This mutual linking increases the number of matched ROIs.

The results can be checked in a user interface (`ChronicViewer.m`). The matches can be verified by selecting a cell in the match matrix, which will light up the ROI contours of the match. Clicking on ROIs in the main image will show the ROI numbers and match information in a table. Matches can be added, edited or removed if necessary.

Quantification and statistical analysis

All bar graphs show the mean of the data, with error bars showing the standard error of the mean (Figures 7A and S3A). All statistical tests were done using MATLAB built-in functions. Data was tested for normality using the Shapiro-Wilk test (swtest). Correlations were tested using corr with Spearman's correlation.