
Research Article: Methods/New Tools | Novel Tools and Methods

ABLE: An Activity-Based Level Set Segmentation Algorithm for Two-Photon Calcium Imaging Data

ABLE: Activity-Based Level Set Segmentation Method

Stephanie Reynolds^{1,2}, Therese Abrahamsson³, Renaud Schuck^{2,4}, P. Jesper Sjöström³, Simon R. Schultz^{2,4} and Pier Luigi Dragotti¹

¹*Department of Electrical and Electronic Engineering, Imperial College London, South Kensington Campus, London, SW7 2AZ, UK*

²*Centre for Neurotechnology, Imperial College London, South Kensington Campus, London, SW7 2AZ, UK*

³*Centre for Research in Neuroscience, Brain Repair and Integrative Neuroscience Program, Department of Neurology and Neurosurgery, the Research Institute of the McGill University Health Centre, Montréal General Hospital, Montréal, Québec H3G 1A4, Canada*

⁴*Department of Bioengineering, Imperial College London, South Kensington Campus, London, SW7 2AZ, UK*

DOI: 10.1523/ENEURO.0012-17.2017

Received: 10 January 2017

Revised: 14 September 2017

Accepted: 15 September 2017

Published: 16 October 2017

Author Contributions: Designed research SR PLD SRS, Performed research SR TA RS PJS, Contributed unpublished reagents/ analytic tools, Analyzed data SR, Wrote the paper SR PLD SRS.

Funding: <http://doi.org/10.13039/501100000781>ERC | European Research Council (ERC): 277800. <http://doi.org/10.13039/501100000268>Biotechnology and Biological Sciences Research Council (BBSRC): BB/K001817/1. EU Marie Curie FP7 Initial Training Network: 289146. CIHR New Investigator Award: 288936. CIHR Operating Grant: 126137. <http://doi.org/10.13039/501100000038>Gouvernement du Canada | Natural Sciences and Engineering Research Council of Canada (NSERC): 418546-2.

Conflict of Interest: Authors report no conflict of interest.

Correspondence should be addressed to Pier Luigi Dragotti, Department of Electrical and Electronic Engineering, Imperial College London, South Kensington Campus, London SW7 2AZ, UK, E-mail address: p.dragotti@imperial.ac.uk or Simon R. Schultz, Address: Department of Bioengineering, Imperial College London, South Kensington Campus, London SW7 2AZ, UK, E-mail address: s.schultz@imperial.ac.uk

Cite as: eNeuro 2017; 10.1523/ENEURO.0012-17.2017

Alerts: Sign up at eneuro.org/alerts to receive customized email alerts when the fully formatted version of this article is published.

Accepted manuscripts are peer-reviewed but have not been through the copyediting, formatting, or proofreading process.

Copyright © 2017 Reynolds et al.

This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International license, which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37

TITLE PAGE

1. Manuscript Title (50 word maximum): ABLE: an Activity-Based Level Set Segmentation Algorithm for Two-Photon Calcium Imaging Data

2. Abbreviated Title (50 character maximum): ABLE: Activity-Based Level Set Segmentation Method

3. Author Names and Affiliations:

Stephanie Reynolds [1,2], Therese Abrahamsson [3], Renaud Schuck [2, 4], P. Jesper Sjöström [3], Simon R. Schultz [2, 4] and Pier Luigi Dragotti [1]

[1] Department of Electrical and Electronic Engineering, Imperial College London, South Kensington Campus, London SW7 2AZ, UK

[2] Centre for Neurotechnology, Imperial College London, South Kensington Campus, London SW7 2AZ, UK

[3] Centre for Research in Neuroscience, Brain Repair and Integrative Neuroscience Program, Department of Neurology and Neurosurgery, The Research Institute of the McGill University Health Centre, Montréal General Hospital, Montréal, Québec H3G 1A4, Canada

[4] Department of Bioengineering, Imperial College London, South Kensington Campus, London SW7 2AZ, UK

4. Author Contributions:

Designed research SR PLD SRS, Performed research SR TA RS PJS, Contributed unpublished reagents/analytic tools, Analyzed data SR, Wrote the paper SR PLD SRS.

5. Correspondence should be addressed to:

Name: Pier Luigi Dragotti

Address: Department of Electrical and Electronic Engineering, Imperial College London, South Kensington Campus, London SW7 2AZ, UK

Email address: p.dragotti@imperial.ac.uk

Name: Simon R. Schultz

Address: Department of Bioengineering, Imperial College London, South Kensington Campus, London SW7 2AZ, UK

Email address: s.schultz@imperial.ac.uk

6. Number of Figures: 6

7. Number of Tables: 0

8. Number of Multimedia: 0

9. Number of words for Abstract: 155

10. Number of words for Significance Statement: 113

11. Number of words for Introduction: 686

38 **12. Number of words for Discussion:** 980

39 **13. Acknowledgements:**

40 This work was supported by European Research Council starting investigator award [grant number
41 277800] (Pier Luigi Dragotti); Biotechnology and Biological Sciences Research Council [grant number
42 BB/K001817/1] (Simon R. Schultz); EU Marie Curie FP7 Initial Training Network [grant number
43 289146] (Simon R. Schultz); CIHR New Investigator Award [grant number 288936] (P. Jesper
44 Sjöström); CFI Leaders Opportunity Fund [grant number 28331] (P. Jesper Sjöström); CIHR Operating
45 Grant [grant number 126137] (P. Jesper Sjöström) and NSERC Discovery Grant [grant number
46 418546-2] (P. Jesper Sjöström).

47 **14. Conflict of Interest:** Authors report no conflict of interest. The authors declare no competing
48 financial interests.

49 **15. Funding sources:** See acknowledgements.

50

51

52

53

54

ABLE: an Activity-Based Level Set Segmentation Algorithm for Two-Photon Calcium Imaging Data

Abstract

1 We present an algorithm for detecting the location of cells from two-photon calcium imaging data. In our
2 framework, multiple coupled active contours evolve, guided by a model-based cost function, to identify cell
3 boundaries. An active contour seeks to partition a local region into two subregions, a cell interior and ex-
4 terior, in which all pixels have maximally 'similar' time courses. This simple, local model allows contours to
5 be evolved predominantly independently. When contours are sufficiently close, their evolution is coupled,
6 in a manner that permits overlap. We illustrate the ability of the proposed method to demix overlapping
7 cells on real data. The proposed framework is flexible, incorporating no prior information regarding a cell's
8 morphology or stereotypical temporal activity, which enables the detection of cells with diverse properties.
9 We demonstrate algorithm performance on a challenging mouse *in vitro* dataset, containing synchronously
10 spiking cells, and a manually labelled mouse *in vivo* dataset, on which ABLE achieves a 67.5% success
11 rate.

12

13 Significance statement

14 Two-photon calcium imaging enables the study of brain activity during learning and behaviour at single-
15 cell resolution. To decode neuronal spiking activity from the data, algorithms are first required to detect
16 the location of cells in the video. It is still common for scientists to perform this task manually, as the
17 heterogeneity in cell shape and frequency of cellular overlap impede automatic segmentation algorithms.
18 We developed a versatile algorithm based on a popular image segmentation approach (the Level Set
19 Method) and demonstrated its capability to overcome these challenges. We include no assumptions on
20 cell shape or stereotypical temporal activity. This lends our framework the flexibility to be applied to new
21 datasets with minimal adjustment.

22 1. Introduction

23 Two-photon calcium imaging has enabled the long-term study of neuronal population activity during
24 learning and behaviour (Peron et al., 2015b). State of the art genetically encoded calcium indicators have

1 sufficient signal-to-noise ratio (SNR) to resolve single action potentials (Chen et al., 2013). Furthermore,
2 recent developments in microscope design have extended the possible field-of-view in which individual neu-
3 rons can be resolved to 9.5mm^2 (Stirman et al., 2016), and enabled the simultaneous imaging of separate
4 brain areas (Lecoq et al., 2014). However, a comprehensive study of activity in even one brain area can
5 produce terabytes of imaging data (Peron et al., 2015a), which presents a considerable signal processing
6 problem.

7 To decode spiking activity from imaging data, one must first be able to accurately detect regions of
8 interest (ROIs), which may be cell bodies, neurites or combinations of the two. Heterogeneity in the ap-
9 pearance of ROIs in imaging datasets complicates the detection problem. The calcium indicator used to
10 generate the imaging video affects both a cell's resting fluorescence and its apparent shape. For exam-
11 ple, some genetically encoded indicators are excluded from the nucleus and therefore produce fluorescent
12 'donuts'. Moreover, imaging data is frequently contaminated with measurement noise and movement arte-
13 facts. These challenges necessitate flexible, robust detection algorithms with minimal assumptions on the
14 properties of ROIs.

15 Manual segmentation of calcium imaging datasets is still commonplace. While this allows the use
16 of complex selection criteria, it is neither reproducible nor scalable. To incorporate implicitly a human's
17 selection criteria, which can be hard to define mathematically, supervised learning from extensive human-
18 annotated data has been implemented (Valmianski et al., 2010; Apthorpe et al., 2016). Other approaches
19 rely on more general cellular properties, such as their expected size and shape (Ohki et al., 2005) and that
20 they represent regions of peak local correlation (Smith and Häusser, 2010; Kaifosh et al., 2014). The latter
21 approaches use lower-dimensional summary statistics of the data, which reduces computational complexity
22 but does not typically allow detection of overlapping regions.

23 To better discriminate between neighbouring cells, some methods make use of the temporal activity
24 profile of imaging data. The (2+1)-D imaging video, which consists of two spatial dimensions and one
25 temporal dimension, is often prohibitively large to work on directly. One family of approaches therefore
26 reshapes the (2+1)-D imaging video into a 2-D matrix. The resulting matrix admits a decomposition —
27 derived from a generative model of the imaging video — into two matrices, each encoding spatial and
28 temporal information. The spatial and temporal components are estimated using a variety of methods,
29 such as independent component analysis (Schultz et al., 2009; Mukamel et al., 2009) or non-negative
30 matrix factorization (Maruyama et al., 2014). Recent variants extend the video model to incorporate detail
31 on the structure of neuronal intracellular calcium dynamics (Pnevmatikakis et al., 2016) or the neuropil

1 contamination (Pachitariu et al., 2016). By expressing the (2+1)-D imaging video as a 2-D matrix, this
 2 type of approach can achieve high processing speeds. This does, however, come at the cost of discarded
 3 spatial information, which can necessitate post-processing with morphological filters (Pnevmatikakis et al.,
 4 2016; Pachitariu et al., 2016).

5 In this paper, we propose a method in which cell boundaries are detected by multiple coupled active
 6 contours. To evolve an active contour we use the level set method, which is a popular tool in bioimaging
 7 due to its topological flexibility (Delgado-Gonzalo et al., 2015). To each active contour, we associate a
 8 higher-dimensional function, referred to as the level set function, whose zero level set encodes the contour
 9 location. We implicitly evolve an active contour via the level set function. The evolution of the level set
 10 function is driven by a local model of the imaging data temporal activity. The data model includes no
 11 assumptions on a cell's morphology or stereotypical temporal activity. Our algorithm is therefore versatile, it
 12 can be applied to a variety of data types with minimal adjustment. For convenience, we refer to our method
 13 as ABLE (an Activity-Based LEvel set method). In the following, we describe the method and demonstrate
 14 its versatility and robustness on a range of *in vitro* and *in vivo* datasets.

15 2. Materials & Methods

16 2.1. Estimating the boundary of an isolated cell

17 Consider a small region of a video containing one cell (e.g. inside the dashed box, Fig. 1A). This region
 18 is composed of two subregions: the cell and the background. We want to partition the region into Ω^{in} and
 19 Ω^{out} , where Ω^{in} corresponds to the cell and Ω^{out} the background. We compute a feature of the respective
 20 subregions, \mathbf{f}^{in} and \mathbf{f}^{out} , with which to classify pixels into the cell interior or background. In particular, we
 21 define $\mathbf{f}^{\text{in}} \in \mathbb{R}^T$ and $\mathbf{f}^{\text{out}} \in \mathbb{R}^T$ as the average subregion time courses, where T is the number of frames in
 22 the video. We estimate the optimal partition as the one that minimizes discrepancies between a pixel's time
 23 course and the average time course of the subregion to which it belongs. To calculate this discrepancy, we
 24 employ a dissimilarity metric, D (see below), which is identically zero when the time courses are perfectly
 25 matched and positive otherwise. As such, we minimize the following cost function, which we refer to as the
 26 external energy,

$$\mathcal{E}_{\text{ext}}(\Omega^{\text{in}}, \Omega^{\text{out}}) = \int_{\Omega^{\text{in}}} D(\mathbf{I}(\mathbf{x}), \mathbf{f}^{\text{in}}) d\mathbf{x} + \int_{\Omega^{\text{out}}} D(\mathbf{I}(\mathbf{x}), \mathbf{f}^{\text{out}}) d\mathbf{x}, \quad (1)$$

27 where $\mathbf{I}(\mathbf{x}) \in \mathbb{R}^T$ is the time course of pixel \mathbf{x} .

1 The cell location estimate is iteratively updated by the algorithm. At each iteration, the cell exterior is
 2 defined as the set of pixels within a fixed distance of the current estimate of the cell interior, see Fig. 1B. The
 3 default distance is taken to be two times the expected radius of a cell. We refer to this exterior region as the
 4 narrowband to emphasise its proximity to the contour of interest. The boundary between the interior and
 5 the narrowband is the active contour. As an active contour is updated, so is the corresponding narrowband
 6 (Fig. 1F). The region of the video for which the optimal partition is sought is therefore not static; rather, it
 7 evolves as an active contour evolves.

8 2.2. Computing the dissimilarity metric

9 Due to the heterogeneity of calcium imaging data, we do not use a universal dissimilarity metric. When
 10 both the pattern and magnitude of a pixel's temporal activity are informative, as is typically the case for
 11 synthetic dyes, we use a measure based on the Euclidean distance, where

$$D^E(I(\mathbf{x}), \mathbf{f}) = \|I(\mathbf{x}) - \mathbf{f}\|^2, \quad (2)$$

12 for $\mathbf{f} \in \mathbb{R}^T$. When we have an image not a video (i.e. $I(\mathbf{x})$ and \mathbf{f} are one-dimensional) this dissimilarity met-
 13 ric reduces to the fitting term introduced by Chan and Vese (2001). For datasets in which the fluorescence
 14 expression level varies significantly throughout cells and, as a consequence, pixels in the same cell exhibit
 15 the same pattern of activity at different magnitudes, we use a measure based on the correlation, such that

$$D^C(I(\mathbf{x}), \mathbf{f}) = 1 - \text{corr}(I(\mathbf{x}), \mathbf{f}), \quad (3)$$

16 where corr represents the Pearson correlation coefficient. In this paper, as default, we use the Euclidean
 17 dissimilarity metric. Additionally, we present two notable examples in which the correlation-based metric is
 18 preferable.

19 2.3. External energy for neighbouring cells

20 We now extend the cost function presented in Eq. (1) to one suitable for partitioning a region into multiple
 21 cell interiors, $\{\Omega^{\text{in},1}, \Omega^{\text{in},2}, \dots, \Omega^{\text{in},M}\}$, and a global exterior, Ω^{out} , which encompasses the narrowbands of all
 22 the cells. We denote with $\mathbf{f}^{\text{in},i}$ the average time course of pixels exclusively in $\Omega^{\text{in},i}$. Due to the relatively low
 23 axial resolution of a two-photon microscope, fluorescence intensity at one pixel can originate from multiple
 24 cells in neighbouring z-planes. Accordingly, we allow cell interiors to overlap when this best fits the data.
 25 In particular, we assume that a pixel in multiple cells would have a time course well fit by the sum of the

1 interior time courses for each cell. The external energy in the case of multiple cells is thus

$$\mathcal{E}_{\text{ext}}(\Omega^{\text{in},1}, \dots, \Omega^{\text{in},M}, \Omega^{\text{out}}) = \int_{\Omega^{\text{out}}} D(I(\mathbf{x}), \mathbf{f}^{\text{out}}) d\mathbf{x} + \int_{\text{inside}} D(I(\mathbf{x}), \sum_{i \in C(\mathbf{x})} \mathbf{f}^{\text{in},i}) d\mathbf{x}, \quad (4)$$

2 where the area termed 'inside' denotes the union of all cell interiors and the function $C(\mathbf{x})$ identifies all
 3 cells whose interior contains pixel \mathbf{x} . When the region to be partitioned contains only one cell, the external
 4 energy in Eq. (4) reduces to that in Eq. (1).

5 2.4. Level Set Method

6 It is not possible to find an optimal cell boundary by minimizing the external energy directly (Chan and
 7 Vese, 2001). An alternative solution is to start from an initial estimate, see below, and evolve this estimate in
 8 terms of an evolution parameter τ . In this approach, the boundary is called an active contour. To update the
 9 active contour we use the Level Set Method of Osher and Sethian (1988). This method was first introduced
 10 to image processing by Caselles et al. (1993) and Malladi et al. (1995); it has since found widespread use
 11 in the field. We implicitly represent the evolving boundary estimate of the i^{th} cell — the i^{th} active contour —
 12 by a function ϕ_i , where ϕ_i is positive for all pixels in the cell interior, negative for those in the narrowband
 13 and zero for all pixels on the boundary (see Fig. 1C). We refer to ϕ_i as a level set function, as its zero
 14 level set identifies the contour of interest. We note that since the contour evolves with τ , ϕ_i itself depends
 15 upon τ . In the following, we present a set of M partial differential equations (PDEs) — one for each active
 16 contour — derived in part from Eq. (4), which dictate the evolution of the level set functions. The solution
 17 to the set of PDEs yields (as the zero level sets) the cell boundaries which minimize the external energy in
 18 Eq. (4).

19 From the external energy and a regularization term (Li et al., 2010), we define a new cost function

$$\mathcal{E}(\phi_1, \dots, \phi_M) = \lambda \mathcal{E}_{\text{ext}}(\phi_1, \dots, \phi_M) + \mu \mathcal{R}(\phi_1, \dots, \phi_M), \quad (5)$$

20 where the arguments to the external energy in Eq. (4) are replaced by the corresponding level set functions.
 21 The parameters λ and μ are real-valued scalars, which define the relative weight of the external energy
 22 and the regularizer. The regularizer is designed to ensure that a level set function varies smoothly in the
 23 vicinity of its active contour. The corresponding regularization energy is minimised when ϕ_i has gradient
 24 of magnitude one near the active contour and magnitude zero far from the contour. An example of such a
 25 function, a signed distance function (which is the shape of all level set functions upon initialization), can be
 26 seen in Fig. 1C.

1 A standard way to obtain the level set function that minimizes the cost function is to find the steady-state
 2 solution to the gradient flow equation (Aubert and Kornprobst, 2006), we do this for each ϕ_i :

$$\frac{\partial \phi_i}{\partial \tau} = -\frac{\partial \mathcal{E}}{\partial \phi_i}, \quad (6)$$

3 for $i \in \{1, 2, \dots, M\}$. From Eq. (5) we obtain

$$\frac{\partial \phi_i}{\partial \tau} = -\left(\lambda \frac{\partial \mathcal{E}_{\text{ext}}}{\partial \phi_i} + \mu \frac{\partial \mathcal{R}}{\partial \phi_i}\right). \quad (7)$$

4 We solve this PDE numerically, by discretizing the evolution parameter τ , such that

$$\phi_i(\tau + 1) = \phi_i(\tau) - \Delta\tau \left(\lambda \frac{\partial \mathcal{E}_{\text{ext}}}{\partial \phi_i} + \mu \frac{\partial \mathcal{R}}{\partial \phi_i}\right). \quad (8)$$

5 The regularization term, which encourages ϕ_i to vary smoothly in the image plane, helps to ensure the
 6 accurate computation of the numerical solution.

7 At every timestep τ , each level set function is consecutively updated until convergence. We must retain
 8 $\mu\Delta\tau < 0.25$ in order to satisfy the Courant-Friedrichs-Lewy condition (Li et al., 2010) — a necessary
 9 condition for the convergence of a numerically-solved PDE. This condition requires that the numerical
 10 waves propagate at least as fast as the physical waves (Osher and Fedkiw, 2003). We therefore set
 11 $\Delta\tau = 10$ and $\mu = 0.2/\Delta\tau$. For each dataset, we tune the value of λ based on the algorithm performance
 12 on a small section of the video. To attain segmentation results on the real datasets presented in this paper,
 13 we use $\lambda = 150$ (Section 3.1), $\lambda = 50$ (Section 3.3), $\lambda = 25$ (Section 3.4) and $\lambda = 10$ (Section 3.5).

14 2.5. External velocity

15 The movement of a level set function, ϕ_i , is driven by the derivatives in Eq. (8)— $\partial \mathcal{E}_{\text{ext}}/\partial \phi_i$ provides the
 16 impetus from the video data and $\partial \mathcal{R}/\partial \phi_i$ the impetus from the regulariser. In the following, we outline the
 17 calculation and interpretation of $\partial \mathcal{E}_{\text{ext}}/\partial \phi_i$; the regulariser is standard and its derivative is detailed in Li et
 18 al. (2010). As is typical in the level set literature (Zhao et al., 1996; Li et al., 2010), using an approximation
 19 of the Dirac delta function δ_ϵ , we obtain an approximation of the derivative: $\partial \mathcal{E}_{\text{ext}}/\partial \phi_i(\mathbf{x}) = \delta_\epsilon(\phi_i(\mathbf{x})) V_i(\mathbf{x})$,
 20 where

$$V_i(\mathbf{x}) = \begin{cases} D(I(\mathbf{x}), \mathbf{f}^{\text{in},i}) - D(I(\mathbf{x}), \mathbf{f}^{\text{out}}) & \text{if } \mathbf{x} \text{ is not in a neighbouring cell,} \\ D(I(\mathbf{x}), \mathbf{f}^{\text{in},i} + \sum_{j \in \mathcal{C}(\mathbf{x})} \mathbf{f}^{\text{in},j}) - D(I(\mathbf{x}), \sum_{j \in \mathcal{C}(\mathbf{x})} \mathbf{f}^{\text{in},j}) & \text{otherwise.} \end{cases} \quad (9)$$

21 We refer to $V_i(\mathbf{x})$ as the external velocity as it encapsulates the impetus to movement derived from the
 22 external energy in Eq. (4), see Fig. 1E for an illustrative example.

1 The term δ_ϵ , which is only non-zero at pixels on or near the cell boundary, acts as a localization operator,
2 ensuring that the velocity only impacts ϕ_i at pixels in the vicinity of the active contour. The parameter ϵ
3 defines the approximate radius, in pixels, of the non-zero band — here, we take $\epsilon = 2$. The product with
4 the localization operator means that, in practise, the external velocity must only be evaluated at pixels on
5 or near the cell boundary. As a consequence, although the external velocity contains contributions from all
6 cells in the video, the problem remains local — only neighbouring cells directly affect a cell's evolution.

7 Although Ω^{out} represents a global exterior, in practise, we calculate the corresponding time course in
8 Eq. (9), f^{out} , locally. To evaluate the external velocity of an active contour, we calculate f^{out} as the average
9 time course from pixels in the corresponding narrowband. This allows us to neglect components such as
10 intensity inhomogeneity and neuropil contamination, see Fig. 2, which we assume vary on a scale larger
11 than that of the narrowband.

12 The external velocity of a single active contour, Eq. (9), can be interpreted as follows: if a pixel, not
13 in another cell, has time course more similar to that of the contour interior than the narrowband, then the
14 contour moves to incorporate that pixel. If a pixel in another cell has time course better matched by the
15 sum of the interior time courses of cells containing that pixel plus the interior time course of the evolving
16 active contour, then the contour moves to incorporate it. Otherwise, the contour is repelled from that pixel.

17 2.6. Initialization

18 We devised an automatic initialization algorithm which selects connected areas of either peak local
19 correlation or peak mean intensity as initial ROI estimates. Initializing areas of peak mean intensity, which
20 may correspond to artefacts rather than active cells (see e.g. the electrode in Fig. 1A), is essential so that
21 these regions do not distort the narrowband signal of another ROI. We first compute the correlation image
22 of the video. For each pixel, this is the average correlation between that pixel's time course and those
23 of the pixels in its 8-connected neighbourhood. Local peaks in this image and the mean intensity image
24 are identified (by a built-in MATLAB function, 'imextendedmax') as candidate ROIs. The selectivity of the
25 initialization is set by a tuning parameter α , which defines the relative height with respect to neighbouring
26 pixels (in units of standard deviation of the input image) of the peaks that are suppressed. The higher the
27 value of α , the more conservative the initialization. We have found it best to use a low value for α (in the
28 range 0.2 - 0.8) so as to overestimate the number of ROIs; redundant estimates are automatically pruned
29 during the update phase of the algorithm. Moreover, smaller values of α produce smaller initializations,
30 which reduce errors due to initializations composed of multiple cells.

1 2.7. Convergence

2 We stop updating a contour estimate if a maximum number of iterations N_{\max} has been reached or the
3 active contour has converged — using one or both of these conditions is common in the active contour
4 literature, see, for example, Delgado-Gonzalo and Unser (2013); Li et al. (2010). A contour is deemed to
5 have converged if, in N_{con} consecutive iterations, the number of pixels that are added to or removed from
6 the interior is less than ρ . As default, we take $N_{\max} = 100$, $N_{\text{con}} = 40$ and $\rho = 2$.

7 The complexity of the level set method is intrinsically related to the dimensionality of the active contour;
8 the number of frames of the video is only relevant to the evaluation of the external velocity, Eq. (9), which
9 accounts for a small fraction of the computational cost. In Table 1, we demonstrate that increasing video
10 length by a factor of 10 has only a minor impact on processing time. As the framework includes no as-
11 sumptions on an ROI's stereotypical temporal activity, prior to segmentation a video can be downsampled
12 by averaging consecutive samples, thereby simultaneously enabling the processing of longer videos and
13 increasing signal-to-noise ratio.

14 Increasing cell density principally impacts the calculation of the external velocity and does, therefore,
15 not alter the computational complexity of the algorithm. On synthetic data, we observe that increasing cell
16 density only marginally affects the convergence rate, see Table 2. As emphasised in Section 2.5, updating
17 an active contour is a local problem — consequently, we observe that algorithm runtime increases linearly
18 with the total number of cells, see Table 1. Due to the independence of spatially separate ROIs in our
19 framework, further performance speed-ups are achievable by parallelizing the computation.

20 2.8. Merging and pruning ROIs

21 ABLE automatically merges two cells if they are sufficiently close and their interiors sufficiently corre-
22 lated — a strategy previously employed in the constrained matrix factorization algorithm of Pnevmatikakis
23 et al. (2016). When two contours are merged, their respective level set functions are replaced with a single
24 level set function, initialized as a signed distance function (Fig. 1C), with a zero level set that represents
25 the union of the contour interiors.

26 The required proximity for two cells to be merged is one cell radius (the expected cell radius is one of two
27 required user input parameters). To determine the correlation threshold we consider the correlation of two
28 noisy time courses corresponding to the average signals from two distinct sets of pixels belonging to the
29 same cell. We assume the underlying signal components — which correspond to the cellular signal plus
30 background contributions — have maximal correlation but that the additive noise reduces the correlation of

1 the noisy time courses. Assuming the noise processes are independent from the underlying cellular signal
 2 and each other, the correlation coefficient of the noisy time courses is

$$\frac{1}{1 + 10^{-\text{SNR}_{\text{dB}}/10}}, \quad (10)$$

3 where SNR_{dB} is the signal-to-noise-ratio (dB) of the noisy time courses. We thus merge components with
 4 correlation above this threshold. We select a default correlation threshold of 0.8, derived from a default
 5 expected SNR of 5 (dB). The user has the option to input an empirically measured SNR, which updates the
 6 correlation threshold using the formula in Eq. (10).

7 A contour is automatically removed ('pruned') during the update phase if its area is smaller or greater
 8 than adjustable minimum or maximum size thresholds, which, as default, are set at 3 and $3\pi r^2$ pixels,
 9 respectively, where r is the expected radius of a cell.

10 2.9. Metric definitions

11 The signal-to-noise ratio (SNR) is defined as the ratio of the power of a signal σ_x^2 and the power of the
 12 noise σ_ϵ^2 , such that $\text{SNR}_{\text{pow}} = \sigma_x^2 / \sigma_\epsilon^2$. We write the SNR in decibels (dB) as $\text{SNR}_{\text{dB}} = 10 \log_{10}(\text{SNR}_{\text{pow}})$.

13 Given two sets of objects, a ground truth set and a set of estimates, the precision is the percentage of
 14 estimates that are also in the ground truth set and the recall is the percentage of ground truth objects that
 15 are found in the set of estimates. As a complement of the precision we use the fall-out rate, the percentage
 16 of estimates not found in the ground truth set. The success rate (%) is

$$2 \frac{\text{precision} * \text{recall}}{\text{precision} + \text{recall}}. \quad (11)$$

17 When the objects are cells, an estimate is deemed to match a ground truth cell if their centres are within
 18 5 pixels of one another. When the objects are spikes, the required distance is 0.22s (3 sample widths).
 19 To quantify spike detection performance we also use the Root-Mean-Square-Error (RMSE), which is the
 20 square root of the average squared error between an estimated spike time ($\hat{t}_k \in \mathbb{R}$) and the ground truth
 21 spike time ($t_k \in \mathbb{R}$).

22 2.10. Simulations

23 To quantify segmentation performance, we simulated calcium imaging videos. In the following, we detail
 24 the method used to generate the videos. Cellular spike trains are generated from mutually independent
 25 Poisson processes. A cell's temporal activity is the sum of a stationary baseline component, the value
 26 of which is selected from a uniform distribution, and a spike train convolved with a stereotypical calcium

1 transient pulse shape. Cells are 'donut' (annulus) shaped to mimic videos generated by genetically encoded
2 calcium indicators, which are excluded from the nucleus. To achieve this, the temporal activity of a pixel in
3 a cell is generated by multiplying the cellular temporal activity vector by a factor in $[0, 1]$ that decreases as
4 pixels are further from the cell boundary. When two cells overlap in one pixel, we sum the contributions of
5 both cells at that pixel. Spatially and temporally varying background activity, generated independently from
6 the cellular spiking activity, is present in pixels that do not belong to a cell.

7 *2.11. Software accessibility*

8 The software described in the paper is freely available online at [redacted for double-blind review].

9 *2.12. Two-photon calcium imaging of quadruple whole-cell recordings*

10 All animal procedures were performed in accordance with the [Author University] animal care commit-
11 tee's regulations. P11-P15 mice of either sex were anaesthetised with isoflurane, decapitated, and the
12 brain was rapidly dissected in 4°C external solution consisting of 125 mM NaCl, 2.5 mM KCl, 1 mM MgCl₂,
13 1.25 mM NaH₂PO₄, 2 mM CaCl₂, 26 mM NaHCO₃, and 25 mM dextrose, bubbled with 95% O₂/5% CO₂
14 for oxygenation and pH. Quadruple whole-cell recordings in acute visual cortex slices were carried out at
15 32°C-34°C with internal solution consisting of 5 mM KCl, 115 mM K-gluconate, 10 mM K-HEPES, 4 mM
16 MgATP, 0.3 mM NaGTP, 10 mM Na-phosphocreatine, and 0.1% w/v biocytin, adjusted with KOH to pH
17 7.2-7.4. On the day of the experiment, 20 μM Alexa Fluor 594 and 180 μM Fluo-5F pentapotassium salt
18 (Life Technologies) were added to the internal solution. Electrophysiology amplifier (Dagan Corporation
19 BVC-700A) signals were recorded with a National Instruments PCI-6229 board, using in-house software
20 running in Igor Pro 6 (WaveMetrics). Two-photon excitation was achieved by raster-scanning a Spectra-
21 physics MaiTai BB Ti:Sa laser tuned to 820 nm across the sample using an Olympus 40x objective and
22 galvanometric mirrors (Cambridge Technologies 6215H, 3 mm, 1 ms/line, 256 lines). Substage photomul-
23 tiplier tube signals (R3896, Hamatsu) were acquired with a National Instruments PCI-6110 board using
24 ScanImage 3.7 running in MATLAB (MathWorks). Layer-5 pyramidal cells were identified by their promi-
25 nent apical dendrites using infrared video Dodt contrast. Unless otherwise stated, all drugs were obtained
26 from Sigma-Aldrich.

27 *2.13. Two-photon calcium imaging of bulk loaded hippocampal slices*

28 All animal procedures were performed in accordance with the [Author University] animal care com-
29 mittee's regulations. Juvenile wild-type mice of either sex (C57Bl6, P13-P21) were anaesthetised using

1 isoflurane prior to decapitation procedure. Brain slices (400 μm thick) were horizontally cut in 1-4°C venti-
2 lated (95% O₂, 5% CO₂) slicing Artificial Cerebro-Spinal Fluid (sACSF: 0.5 mM CaCl₂, 3.0 mM KCl, 26 mM
3 NaHCO₃, 1 mM NaH₂PO₄, 3.5 mM MgSO₄, 123 mM Sucrose, 10 mM D-Glucose). Hippocampal slices
4 containing Dentate Gyrus, CA3 and CA1 were taken and resting in ventilated recovery ACSF (rACSF: 2
5 mM CaCl₂, 123 mM NaCl, 3.0 mM KCl, 26 mM NaHCO₃, 1mM NaH₂PO₄, 2mM MgSO₄, 10mM D-Glucose)
6 for 30min at 37°C. After this the slices were placed in an incubation chamber containing 2.5 mL of venti-
7 lated rACSF and 'painted' with 10 μL of the following solution: 50 μg of Cal-520 AM (AAT Bioquest), 2 μL
8 of Pluronic-F127 20% in DMSO (Life Technologies) and 48 μL of DMSO (Sigma Aldrich) where they were
9 left for 30 min at 37°C in the dark. Slices were then washed in rACSF at room temperature for 30 min
10 before imaging. Dentate Gyrus granular cells were identified using oblique illumination prior to being im-
11 aged using a standard commercial galvanometric scanner based two-photon microscope (Scientifica Ltd)
12 coupled to a mode-locked Mai Tai HP Ti Sapphire (Spectra-Physics) laser system operating at 810 nm.
13 Functional calcium images of granular cells were acquired with a 40X objective (Olympus) by raster scan-
14 ning a $180 \times 180 \mu\text{m}^2$ square Field of View at 10 Hz. Electrical stimulation was accomplished with a tungsten
15 bipolar concentric microelectrode (WPI) where the tip of the electrode was placed into the molecular layer
16 of the Dentate Gyrus (20 pulses with a pulse-width of 400 μs and a 60 μA amplitude were delivered into
17 the tissue with a pulse repetition rate of 10 Hz, repeated every 40 sec). Unless otherwise stated, all drugs
18 were obtained from Sigma-Aldrich.

19 **3. Results**

20 *3.1. ABLE is robust to heterogeneity in cell shape and baseline intensity*

21 ABLE detected 236 ROIs with diverse properties from the publicly available mouse *in vivo* imaging
22 dataset of Peron et al. (2015c), see Fig. 2. Automatic initialization on this dataset produced 253 ROIs with
23 17 automatically removed during the update phase of the algorithm after merging with another region.

24 To maintain a versatile framework we included no priors on cellular morphology in the cost function
25 that drives the evolution of an active contour. This allowed ABLE to detect ROIs with varied shapes (Fig.
26 2A) and sizes (Fig. 2D). The smaller detected ROIs correspond to cross-sections of dendrites (Fig. 2E),
27 whereas the majority correspond to cell bodies. The topological flexibility of the level set method allows
28 cell bodies and neurites to be segmented as separate (Fig. 2G) or connected (Fig. 2A) objects, depending
29 on the correlation between their time courses. ABLE automatically merges neighbouring regions that are
30 sufficiently correlated (Fig. 2F). Cell bodies and dendrites that are initialised separately and exhibit distinct

1 temporal activity, however, are not merged. For example, the cell body and neurite in Fig. 2G were not
2 merged as the cell body's saturating fluorescence time course was not sufficiently highly correlated with
3 that of the neurite.

4 Evaluating the external velocity, which drives an active contour's evolution, requires only data from
5 pixels in close proximity to the contour (see Section 2.5). This region has radius of the same order as
6 that of a cell body. Background intensity inhomogeneity, caused by uneven loading of synthetic dyes or
7 uneven expression of virally inserted genetically encoded indicators, tend to occur on a scale larger than
8 this. On this dataset we show that, as a result of this local approach, ABLE is robust to background intensity
9 inhomogeneity. This is illustrated by the wide range of baseline intensities of the detected ROIs (Fig. 2C),
10 some of which are even lower than the video median.

11 No prior information on stereotypical neuronal temporal activity is included in our framework. Cells
12 detected by ABLE exhibit both stereotypical calcium transient activity (Fig. 2B:1-9) and non-stereotypical
13 activity (Fig. 2B:10-12), perhaps corresponding to saturating fluorescence, higher firing cell types such as
14 interneurons, or non-neuronal cells.

15 The scattering of photons when imaging at depth can result in leakage of neuropil signal into cellular
16 signal. To obtain decontaminated cellular time courses it is thus important to perform neuropil correction in a
17 subsequent stage, once cells have been located. This involves computation of the decontaminated cellular
18 signal by subtracting the weighted local neuropil signal from the raw cellular signal. As illustrated in Fig. 2H,
19 the proposed method naturally facilitates neuropil correction, as it computes the required components as a
20 by-product of the segmentation process (see Section 2.5). The appropriate value of the weight parameter
21 varies depending on the imaging set-up (Peron et al. 2015a; Chen et al. 2013; Kerlin et al. 2010). We
22 therefore do not include neuropil-correction as a stage of the algorithm, preferring instead to allow users
23 the flexibility to choose the appropriate parameter in post-processing.

24 *3.2. ABLE demixes overlapping cells*

25 When imaging through scattering tissue, a two-photon microscope can have relatively low axial resolu-
26 tion (on the order of ten microns) in comparison to its excellent lateral resolution. As a consequence, the
27 photons collected at one pixel can in some cases originate from multiple cells in a range of z-planes. For
28 this reason, cells can appear to overlap in an imaging video (for an example, see Fig. 3E). It is crucial
29 that segmentation algorithms can delineate the true boundary of 'overlapping' cells, which we refer to as
30 'demixing', so that the functional activity of each cell can be correctly extracted and analysed. In a set of
31 experiments on real and simulated data, we demonstrated that ABLE can demix overlapping cells.

1 On synthetic data containing 25 cells, 17 of which had some overlap with another cell, we measured
 2 the success rate of ABLE's segmentation compared to the ground truth cell locations (Fig. 3A-C), when the
 3 algorithm was initialised on a fixed grid (Fig. 3D). For full description of the performance metric used, see
 4 Section 2.9. Performance was measured over 10 realizations of noise at each noise level. On average, over
 5 all cells and noise realizations, ABLE achieved success rate greater than 99% when the noise standard
 6 deviation was less than 90 (Fig. 3B). Cells were simulated with uneven brightness to mimic the 'donut' cells
 7 generated by some genetically encoded indicators that are excluded from the nucleus. Consequently, the
 8 correlation-based dissimilarity metric was used on this data. As a result, pixels with significantly different
 9 resting fluorescence, but identical temporal activity pattern, were segmented in the same cell (Fig. 3A).

10 On the publicly available mouse *in vivo* imaging dataset of Peron et al. (2015c), ABLE demixed over-
 11 lapping cells (Fig. 3E-F). In this dataset, the vibrissal cortex was imaged at various depths, from layer 1 to
 12 deep layer 3, whilst the mouse performed a pole localization task (Peron et al., 2015a; Guo et al., 2014).
 13 Some cells appear to overlap, due to the relatively low axial resolution when imaging at depth through
 14 tissue. When an ROI was initialised in each separate neuron, ABLE accurately detected the overlapping
 15 cell boundaries using the Euclidean distance dissimilarity metric, Eq. (2). On the Neurofinder Challenge
 16 dataset presented in Section 3.4, ABLE demixed overlapping cells when performing segmentation with the
 17 correlation-based dissimilarity metric, Eq. (3), see Fig. 3G.

18 3.3. ABLE detects synchronously spiking, densely packed cells

19 ABLE detected 207 ROIs from mouse *in vitro* imaging data (Fig. 4). Cells in this dataset exhibit activity
 20 that is highly correlated with other cells and the background as the brain slice was electrically stimulated
 21 (at rate 10Hz for 2s every 40s) during imaging. When the cell interior and narrowband time courses are
 22 highly correlated, the external velocity of the active contour, Eq. (9), derived from the Euclidean distance
 23 dissimilarity metric, Eq. (2), is driven by the discrepancy between the baseline intensities of the subregions.
 24 This is evident when we consider the average time course of the cell interior (f^{in}) and exterior (f^{out}) as a
 25 sum of a stationary baseline component — the resting fluorescence — and an activity component that is
 26 zero when a neuron is inactive, such that $f^{\text{in}} = b^{\text{in}} + a^{\text{in}}$ and $f^{\text{out}} = b^{\text{out}} + a^{\text{out}}$. The time course of a pixel
 27 \mathbf{x} is $I(\mathbf{x}) = b^{\mathbf{x}} + a^{\mathbf{x}}$. Substituting these expressions into Eq. (9), for pixels not in another cell, we obtain
 28 the external velocity $V(\mathbf{x}) = \|b^{\mathbf{x}} - b^{\text{in}}\|^2 - \|b^{\mathbf{x}} - b^{\text{out}}\|^2 + R$, where the residual, R , encompasses all terms
 29 with contributions from the activity components. When the cell and the background are highly correlated
 30 — meaning that the discrepancy between activity components is low and, consequently, the contribution
 31 from R is comparatively small — the external velocity will drive the contour to include pixels with baselines

1 more similar to the interior than the background. As a result of this, ABLE detected ROIs despite their high
2 correlation with the background (Fig. 4C). Furthermore, inactive ROIs were detected (Fig. 4 H-J), when
3 their baseline fluorescence allowed them to be identified from the background (Fig. 4I).

4 The algorithm was automatically initialised on this dataset with 250 ROIs, initializations in the bar (an
5 artefact that can be seen in the top right of Fig. 4A) were prohibited. Of the initialised ROIs, 19 were pruned
6 automatically during the update phase of the algorithm as (i) their interior time course was not sufficiently
7 different from that of the narrowband (3 ROIs), (ii) they merged with another region (2 ROIs) or (iii) they
8 crossed the minimum and maximum size thresholds (14 ROIs).

9 3.4. Algorithm comparison on manually labelled dataset

10 We compared the performance of ABLE with two state of the art calcium imaging segmentation al-
11 gorithms — CNMF (Pnevmatikakis et al., 2016) and Suite2p (Pachitariu et al., 2016) — on a manually
12 labelled dataset from the Neurofinder Challenge, see Fig. 5. The dataset, which can be accessed at the
13 Neurofinder Challenge website (see references), was recorded at 8Hz and generated using the genetically
14 encoded calcium indicator GCaMP6s. Consequently, we apply ABLE with the correlation-based dissimilar-
15 ity metric, Eq. (3), which is well suited to neurons with low baseline fluorescence and uneven brightness.
16 As the dataset is large enough (512x512x8000 pixels) to present memory issues on a standard laptop,
17 we run the patch-based implementation of CNMF, which processes spatially-overlapping patches of the
18 dataset in parallel. We optimise the performance of each algorithm by selecting a range of values for each
19 of a set of tuning parameters and generating segmentation results for all combinations of the parameter
20 set. The results are visualised on the correlation image and the parameter set that presents the best match
21 to the correlation image is selected. This process is representative of what a user may do in practise when
22 applying an algorithm to a new dataset.

23 ABLE achieved the highest success rate (67.5%) when compared to the manual labels, see Table 3.
24 For a definition of the success rate and other performance metrics used, see Section 2.9. ABLE achieved
25 a lower fall-out rate than Suite2p and CNMF (Fig. 5C) — 67.5% of the ROIs it detected matched with
26 the manually labelled cells. Some of the ‘false positives’ were consistent among algorithms (Fig. 5C) and
27 corresponded to local peaks in the correlation image (Fig. 5D), whose extracted time courses displayed
28 stereotypical calcium transient activity (Fig. 5E). A subset of these ROIs may thus correspond to cells
29 omitted by the manual operator. The highest proportion of the manually labelled cells were detected by
30 Suite2p, which detected the greatest number of cells not detected by any other algorithm (Fig. 5B). A small
31 proportion (13.2%) of cells were detected by none of the algorithms. As can be seen from Fig. 5A, these

1 do not correspond to peaks in the correlation image, and may reflect inactive cells detected by the manual
2 operator.

3 3.5. Spikes are detected from ABLE-extracted time courses with high temporal precision

4 Typically, after cells have been identified in calcium imaging data, spiking activity is detected from the
5 extracted cellular time courses and the relationship between cellular activity (and, if measured, external
6 stimuli) is analysed. On a mouse *in vitro* dataset (21 videos, each 30s long), we demonstrated that time
7 courses from cells automatically segmented by ABLE allow spikes to be detected accurately and with high
8 temporal precision (Fig. 6). The dataset has simultaneous electrophysiological recordings from four cells
9 (the electrodes can be seen in the mean image Fig. 6A), which enabled us to compare inferred spike times
10 from the imaging data with the ground truth. We performed spike detection automatically with an existing
11 algorithm (Anonymous, —). On average, over all cells and recordings, 78% of ground truth spikes are
12 detected with a precision of 88% (Fig. 6D). The error in the location of detected spikes is less than one
13 sample width — the average absolute error was 0.053 (s).

14 4. Discussion

15 In this paper, we present a novel approach to the problem of detecting cells from calcium imaging data.
16 Our approach uses multiple coupled active contours to identify cell boundaries. The core assumption is
17 that the local region around a single cell (e.g. inside the dashed box Fig. 1A) can be well-approximated by
18 two subregions, the cell interior and exterior. The average time course of the respective subregions is used
19 as a feature with which to classify pixels into either subregion. We assume that pixels in which multiple cells
20 overlap have time courses that are well-approximated by the sum of each cell's time course. We form a cost
21 function based on these assumptions that is minimised when the active contours are located at the true cell
22 boundaries. Our results on real and simulated data indicate that this is a versatile and robust framework for
23 segmenting calcium imaging data.

24 The cost function in our framework (Eq. 4) penalises discrepancies between the time course of a pixel
25 and the average time course of the subregion to which it belongs. To calculate this discrepancy we use
26 one of two dissimilarity metrics: one based on the correlation, which compares only patterns of temporal
27 activity, the other based on the Euclidean distance, which implicitly takes into account both pattern and
28 magnitude of temporal activity. When the latter metric is used, our cost function is closely related to that
29 of Chan and Vese (2001). If we were to take as an input one frame of a video (or a 2D summary statistic

1 such as the mean image), the external energy in our cost function for an isolated cell would be identical to
2 the fitting term of Chan and Vese (2001). The lower-dimensional approach is, however, not sufficient for
3 segmenting cells with neighbours that have similar baseline intensities. By incorporating temporal activity
4 we can accurately delineate the boundaries of neighbouring cells (Fig. 3A).

5 We evolve one active contour for each cell identified in the initialization. Contours are evolved predomi-
6 nantly independently, with the exception of those within a few pixels of another active contour (see Section
7 2.5). In contrast to previous approaches to coupling active contours (Zimmer and Olivo-Marin, 2005; Dufour
8 et al., 2005), we do not penalise overlap of contour interiors. This is because low axial resolution when
9 imaging through scattering tissue can result in the signals of multiple cells being expressed in one pixel. We
10 therefore permit interiors to overlap when the data is best fit by the sum of average interior time courses.
11 Using this method we can accurately demix the contribution of multiple cells from single pixels in real and
12 simulated data (Fig. 3).

13 ABLE is a flexible method: we include no priors on a region's morphology or stereotypical temporal
14 activity. Due to this versatility, ABLE segmented cells with varying size, shape, baseline intensity and cell
15 type from a mouse *in vivo* dataset (Fig. 2). Moreover, only 2 parameters need to be set by a user for a
16 new dataset. These are the expected radius of a cell and λ , the relative strength of the external velocity
17 compared to the regulariser, see Eq. (5). In order to permit ABLE to segment irregular shapes such as cell
18 bodies attached to dendritic branches (Fig. 2A), the weighting parameter, λ , must be set sufficiently high
19 to counter the regulariser's implicit bias towards smooth contours.

20 Unlike matrix factorization (Maruyama et al., 2014; Pnevmatikakis et al., 2016) and dictionary learning
21 (Diego Andilla and Hamprecht, 2014), which fit a global model to an imaging video, our approach requires
22 only local information to evolve a contour. To evolve an active contour, ABLE uses temporal activity from
23 an area around that contour with size on the order of the radius of a cell. This allows us to omit from our
24 model the spatial variation of the neuropil signal and baseline intensity inhomogeneities, which we assume
25 to be constant on our scale. Our local approach means that the algorithm is readily parallelizable and, in
26 the current implementation, runtime is virtually unaffected by video length (Table 1) and increases linearly
27 with the number of cells.

28 Like any level set method, the performance of ABLE is bounded by the quality of the initialization —
29 if no seed is placed in a neuron it will not be detected, if a seed is spread across multiple neurons they
30 may be jointly segmented. In this work, we developed an automatic initialization algorithm that selects local
31 peaks in the correlation and mean images as candidate ROIs. This approach, however, can lead to false

1 negatives in dense clusters of cells in which the correlation image can appear smooth. In future work, an
2 initialization based on temporal activity, rather than a 2D summary statistic, could overcome this issue. Our
3 algorithm included minimal assumptions about the objects to be detected. To tailor ABLE to a specific data
4 type (e.g. somas versus neurites), it is possible to incorporate terms relating to a region's morphology or
5 stereotypical temporal activity into the cost function. Furthermore, the level set method is straightforward to
6 extend to higher dimensions (Dufour et al., 2005), which means our framework could be adapted to detect
7 cells in light-sheet imaging data (Ahrens et al., 2013).

8 Here we have presented a framework in which multiple coupled active contours detect the boundaries
9 of cells from calcium imaging data. We have demonstrated the versatility of our framework, which includes
10 no priors on a cell's morphology or stereotypical temporal activity, on real *in vivo* imaging data. In this data,
11 we are able to detect cells of various shapes, sizes, and types. We couple the active contours in a way that
12 permits overlap when this best fits the data. This allows us to demix overlapping cells on real and simulated
13 data, even in high noise scenarios. Our results on a diverse array of real datasets indicate that ours is a
14 flexible and robust framework for segmenting calcium imaging data.

15 References

- 16 Ahrens MB, Orger MB, Robson DN, Li JM, Keller PJ (2013) Whole-brain functional imaging at cellular resolution using light-sheet
17 microscopy. *Nature methods* 10:413–420.
- 18 Apthorpe N, Riordan A, Aguilar R, Homann J, Gu Y, Tank D, Seung HS (2016) Automatic neuron detection in calcium imaging data
19 using convolutional networks In: *Advances in Neural Information Processing Systems 29*, pp3270–3278. Curran Associates, Inc.
- 20 Aubert G, Kornprobst P (2006) *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of*
21 *Variations*. New York: Springer-Verlag.
- 22 Caselles V, Catté F, Coll T, Dibos F (1993) A geometric model for active contours in image processing. *Numerische Mathe-*
23 *matik* 66:1–31.
- 24 Chan TF, Vese LA (2001) Active contours without edges. *IEEE Transactions on Image Processing* 10:266–277.
- 25 Chen TW, Wardill TJ, Sun Y, Pulver SR, Renninger SL, Baohan A, Schreiter ER, Kerr RA, Orger MB, Jayaraman V, Looger LL,
26 Svoboda K, Kim DS (2013) Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499:295–300.
- 27 Delgado-Gonzalo R, Uhlmann V, Schmitter D, Unser M (2015) Snakes on a plane: A perfect snap for bioimage analysis. *IEEE Signal*
28 *Processing Magazine* 32:41–48.
- 29 Delgado-Gonzalo R, Unser M (2013) Spline-based framework for interactive segmentation in biomedical imaging. *IRBM* 34:235 – 243.
- 30 Diego Andilla F, Hamprecht FA (2014) Sparse space-time deconvolution for calcium image analysis In: *Advances in Neural Information*
31 *Processing Systems 27*, pp64–72. Curran Associates, Inc.
- 32 Dufour A, Shinin V, Tajbakhsh S, Guillen-Aghion N, Olivo-Marin JC, Zimmer C (2005) Segmenting and tracking fluorescent cells in
33 dynamic 3-D microscopy with coupled active surfaces. *IEEE Transactions on Image Processing* 14:1396–1410.

- 1 Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K (2014) Flow of cortical activity underlying a tactile decision
2 in mice. *Neuron* 81:179 – 194.
- 3 Kaifosh P, Zaremba JD, Danielson NB, Losonczy A (2014) SIMA: Python software for analysis of dynamic fluorescence imaging data.
4 *Frontiers in Neuroinformatics* 8:80.
- 5 Kerlin AM, Andermann ML, Berezovskii VK, Reid RC (2010) Broadly tuned response properties of diverse inhibitory neuron subtypes
6 in mouse visual cortex. *Neuron* 67:858 – 871.
- 7 Lecoq J, Savall J, Vučinić D, Grewe BF, Kim H, Li JZ, Kitch LJ, Schnitzer MJ (2014) Visualizing mammalian brain area interactions by
8 dual-axis two-photon calcium imaging. *Nature Neuroscience* 17:1825–1829.
- 9 Li C, Xu C, Gui C, Fox MD (2010) Distance regularized level set evolution and its application to image segmentation. *IEEE Transactions*
10 *on Image Processing* 19:3243–3254.
- 11 Malladi R, Sethian JA, Vemuri BC (1995) Shape modeling with front propagation: a level set approach. *IEEE Transactions on Pattern*
12 *Analysis and Machine Intelligence* 17:158–175.
- 13 Maruyama R, Maeda K, Moroda H, Kato I, Inoue M, Miyakawa H, Aonishi T (2014) Detecting cells using non-negative matrix factor-
14 ization on calcium imaging data. *Neural Networks* 55:11 –19.
- 15 Mukamel EA, Nimmerjahn A, Schnitzer MJ (2009) Automated analysis of cellular signals from large-scale calcium imaging data.
16 *Neuron* 63:747–760.
- 17 Neurofinder Challenge <http://neurofinder.codeneuro.org/> Accessed: 2017-05-29.
- 18 Ohki K, Chung S, Ch'ng YH, Kara P, Reid RC (2005) Functional imaging with cellular resolution reveals precise micro-architecture in
19 visual cortex. *Nature* 433:597–603.
- 20 Osher S, Fedkiw R (2003) *Level set methods and dynamic implicit surfaces*. New York: Springer-Verlag.
- 21 Osher S, Sethian JA (1988) Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations.
22 *Journal of computational physics* 79:12–49.
- 23 Pachitariu M, Stringer C, Schröder S, Dipoppa M, Rossi LF, Carandini M, Harris KD (2016) Suite2p: beyond 10,000 neurons with
24 standard two-photon microscopy. *bioRxiv* .
- 25 Peron S, Chen TW, Svoboda K (2015b) Comprehensive imaging of cortical networks. *Current Opinion in Neurobiology* 32:115 – 123.
- 26 Peron S, Freeman J, Iyer V, Guo C, Svoboda K (2015c) Volumetric calcium imaging data recorded during performance of a single
27 whisker object localization task, sampling activity in the majority of the relevant superficial barrel cortex neurons (75%, 12,000
28 neurons per mouse) CRCNS.org.
- 29 Peron SP, Freeman J, Iyer V, Guo C, Svoboda K (2015a) A cellular resolution map of barrel cortex activity during tactile behavior.
30 *Neuron* 86:783 – 799.
- 31 Pnevmatikakis EA, Soudry D, Gao Y, Machado TA, Merel J, Pfau D, Reardon T, Mu Y, Lacefield C, Yang W, Ahrens M, Bruno R,
32 Jessell TM, Peterka DS, Yuste R, Paninski L (2016) Simultaneous denoising, deconvolution, and demixing of calcium imaging
33 data. *Neuron* 89:285–299.
- 34 Schultz SR, Kitamura K, Post-Uiterweer A, Krupic J, Häusser M (2009) Spatial pattern coding of sensory information by climbing
35 fiber-evoked calcium signals in networks of neighboring cerebellar purkinje cells. *Journal of Neuroscience* 29:8005–8015.
- 36 Smith SL, Häusser M (2010) Parallel processing of visual space by neighboring neurons in mouse visual cortex. *Nature Neuro-*
37 *science* 13:1144–1149.
- 38 Stirman JN, Smith IT, Kudenov MW, Smith SL (2016) Wide field-of-view, multi-region, two-photon imaging of neuronal activity in the
39 mammalian brain. *Nature Biotechnology* 34:857–862.

- 1 Valmianski I, Shih AY, Driscoll JD, Matthews DW, Freund Y, Kleinfeld D (2010) Automatic identification of fluorescently labeled brain
 2 cells for rapid functional imaging. *Journal of Neurophysiology* 104:1803–1811.
- 3 Zhao HK, Chan T, Merriman B and Osher S (1996) A variational level set approach to multiphase motion. *Journal of Computational*
 4 *Physics* 127:179–195.
- 5 Zimmer C, Olivo-Marin JC (2005) Coupled parametric active contours. *IEEE Transactions on Pattern Analysis and Machine Intelli-*
 6 *gence* 27:1838–1842.

7 Legends

8 Table 1: On synthetic data with dimensions $512 \times 512 \times T$, the runtime of ABLE (minutes) increases linearly
 9 with the number of cells and is not significantly affected by increasing number of frames, T . Runtime was
 10 measured on a PC with 3.4GHz Intel Core i7 CPU.

11

12 Table 2: On synthetic data the average number of iterations to convergence, over 100 realizations of noisy
 13 data, marginally increases as the number of cells in a given cell's narrowband ('neighbouring cells') in-
 14 creases.

15

16 Table 3: On a manually labelled dataset from the Neurofinder Challenge, we compare the performance of
 17 three segmentation algorithms: ABLE, CNMF (Pnevmatikakis et al., 2016) and Suite2p (Pachitariu et al.,
 18 2016), using the manual labels as ground truth.

19

20 Figure 1: A flow diagram of the main steps of the proposed segmentation algorithm: initialization (**A-C**),
 21 iterative updates of the estimate (**D-G**) and convergence (**H-J**). When cells are sufficiently far apart we can
 22 segment them independently — in this example we focus on the isolated cell in the dashed box on the
 23 maximum intensity image in **A**. We make an initial estimate of the cell interior, from which we form the
 24 corresponding narrowband (**B**) and level set function ϕ (**C**). Based on the discrepancy between a pixel's
 25 time course and the time courses of the interior and narrowband regions (**D**), we calculate the velocity of
 26 ϕ at each pixel (**E**). ϕ evolves according to this velocity (**G**), which updates the location of the interior and
 27 narrowband (**F**). Final results for: one cell (**H**), the average signals from the corresponding interior and
 28 narrowband (**I**) and segmentation of all four cells (**J**).

29

30 Figure 2: ABLE detects cells with varying size, shape and baseline intensity from mouse *in vivo* imaging
 31 data. The 236 detected ROIs are superimposed on the mean image of the imaging video (**A**). Extracted

1 neuropil-corrected time series and corresponding ROIs are displayed for a subset of the detected regions
2 (**B**). Cells with both stereotypical calcium transient activity (**B**, 1-9) and saturating fluorescence (**B**, 10-12)
3 are detected. The performance of ABLE does not deteriorate due to intensity inhomogeneity: ROIs with
4 baseline fluorescence from beneath the video median to just below saturation are detected (**C**). The area of
5 detected regions varies (**D**) with the smallest ROIs corresponding to cross-sections of dendrites (**E**). Neigh-
6 bouring regions with sufficiently high correlation are merged (**F**), those with lower correlation are not merged
7 (**G**). In **F** we plot the ROIs prior to and after merging along with the corresponding neuropil-corrected time
8 courses. In **G** we plot the separate ROIs and the neuropil-corrected time courses. The proposed method
9 naturally facilitates neuropil-correction — the removal of the weighted, local neuropil time course from the
10 raw cellular time course (**H**).

11
12 Figure 3: ABLE demixes overlapping cells in real and simulated data. With high accuracy, we detect the
13 true boundaries of overlapping cells from noisy simulated data, the detected contours for one realization
14 of noise with standard deviation (σ) 60 are plotted on the correlation image in **A**. Given an initialisation on
15 a fixed grid, displayed on the mean image in **D**, we detect the true cell boundaries with success rate of at
16 least 99% for $\sigma < 90$ (**B**). The central marker and box edges in **B** indicate the median and the 25th and
17 75th percentiles, respectively. For noise level reference, we plot the average time course from inside the
18 green contour in **A** at various levels (**C**). ABLE demixes overlapping cells in real GCaMP6s mouse *in vivo*
19 data — detected boundaries are superimposed on the mean image (**E** and **F**) and correlation image (**G**),
20 respectively.

21
22 Figure 4: ABLE detects synchronously spiking, densely packed cells from mouse *in vitro* imaging data.
23 The boundaries of the 207 detected ROIs are superimposed on the thresholded maximum intensity image
24 (**A**) and the correlation image (**D**). For all correlation data we use Pearson's correlation coefficient. ABLE
25 detects ROIs that exhibit high correlation with the background **C** and neighbouring synchronously spiking
26 ROIs (**B**). Panel **B** displays the neuropil-corrected extracted time courses of the 207 ROIs (each plotted as
27 a row of the matrix) along with the video mean raw activity and the time points of the electrical stimulations.
28 Panel **C** displays the histogram of the correlation coefficient between the mean raw activity of the video and
29 the extracted time series of each ROI. ABLE detected both active (**E-G**) and inactive ROIs (**H-I**). We display
30 the contours of the two detected ROIs on the correlation image (**E** and **H**), the mean image (**F** and **I**) and
31 the corresponding extracted time courses (**G** and **J**).

1

2 Figure 5: We compare the segmentation results of ABLE, CNMF (Pnevmatikakis et al., 2016) and Suite2p
 3 (Pachitariu et al., 2016) on a manually labelled dataset from the Neurofinder Challenge. On the correlation
 4 image we plot the boundaries of the manually labelled cells colour-coded by the combination of algorithms
 5 that detected them (**A**), undetected cells are indicated by a white contour. Suite2p detected the highest
 6 proportion of manually labelled cells (**B**), whereas ABLE had the lowest fall-out rate (**C**), which is the per-
 7 centage of detected regions not present in the manual labels. Some algorithm-detected ROIs that were not
 8 present in the manual labels are detected by multiple algorithms (**D**) and have time courses which exhibit
 9 stereotypical calcium transient activity (**E**). The correlation image in **D** is thresholded to enhance visibility
 10 of local peaks in correlation. In **E**, we plot the extracted time courses of the ROIs in **D**.

11

12 Figure 6: Spikes are detected from ABLE-extracted time courses with high accuracy. On an *in vitro* dataset
 13 (21 imaging videos, each 30s long) we demonstrate spike detection performance compared to electrophys-
 14 iological ground truth on time courses extracted from cells segmented by ABLE. We plot the labelled cells
 15 (**A**) and corresponding boundaries detected by ABLE (**B**) on the mean image of one imaging video. The
 16 extracted cellular time courses and detected spikes are plotted in **C**. Spike detection was performed with
 17 an existing algorithm (Anonymous, —). On average over all videos, 78% of spikes are detected with a
 18 precision of 88% **D**.

19 Tables, Figures and Multimedia

Table 1: Runtime (minutes) on synthetic data of size $512 \times 512 \times T$.

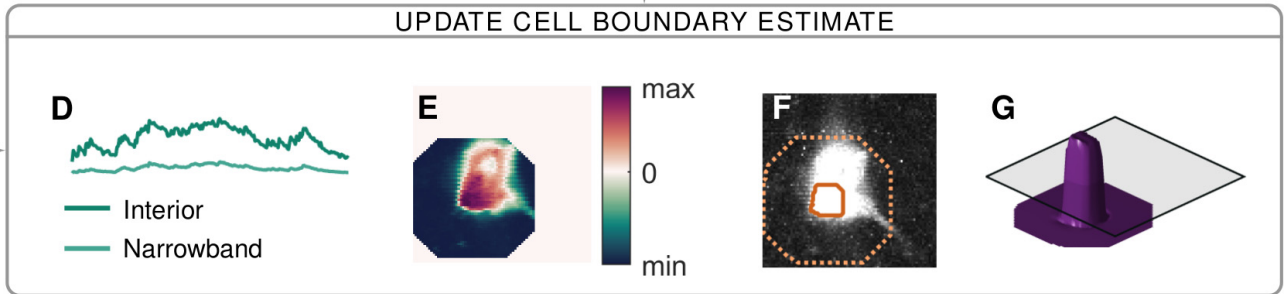

		Num. cells		
		25	125	225
Num. frames (T)	100	1.1	6.5	11.2
	1000	1.3	6.5	12.7

Table 2: Number of iterations to convergence as cell density increases.

Num. neighbours	0	1	2	3	4
Num. iterations	33	33	35	35	36

Table 3: Algorithm success rate (1dp) on manually labelled dataset.


	Success rate (%)	Precision (%)	Recall (%)
ABLE	67.5	67.5	67.5
CNMF	63.4	60.7	66.5
Suite2p	63.7	56.5	73.1




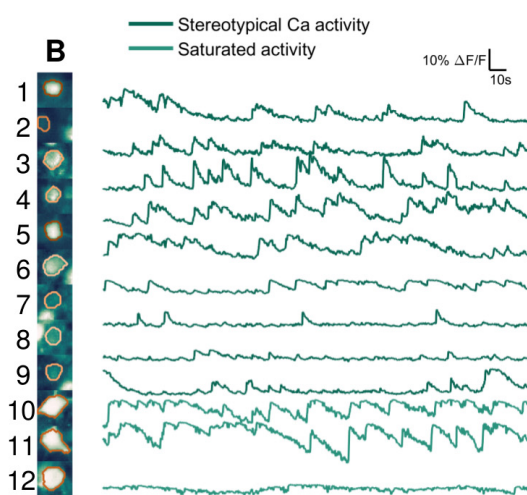



No


Has boundary converged?

Yes






F Automatically merged ROIs and final ROI.



G Not merged, correlation coefficient = 0.72.



H

