A deep learning algorithm for 3D cell detection in whole mouse brain image datasets

Adam L Tyson, Charly V Rousseau, Christian J Niedworok, Sepiedeh
Keshavarzi, Chryssanthi Tsitoura, Lee Cossell, Molly Strom, Troy W Margrie

RESEARCH ARTICLE

A deep learning algorithm for 3D cell detection in whole mouse brain image datasets
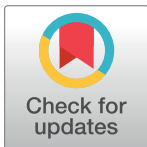
Adam L. Tyson❶◉*, Charly V. Rousseau◉¤, Christian J. Niedworok◉, Sepiedeh Keshavarzi❶, Chryssanthi Tsitoura❶, Lee Cossell❶, Molly Strom, Troy W. Margrie*

Sainsbury Wellcome Centre, University College London, London, United Kingdom

◉ These authors contributed equally to this work.
¤ Current address: Sorbonne Université, Institut du Cerveau - Paris Brain Institute - ICM, Inserm, CNRS, Paris, France
* adam@adamltyson.com (ALT); t.margrie@ucl.ac.uk (TWM)

Abstract

Understanding the function of the nervous system necessitates mapping the spatial distributions of its constituent cells defined by function, anatomy or gene expression. Recently, developments in tissue preparation and microscopy allow cellular populations to be imaged throughout the entire rodent brain. However, mapping these neurons manually is prone to bias and is often impractically time consuming. Here we present an open-source algorithm for fully automated 3D detection of neuronal somata in mouse whole-brain microscopy images using standard desktop computer hardware. We demonstrate the applicability and power of our approach by mapping the brain-wide locations of large populations of cells labeled with cytoplasmic fluorescent proteins expressed via retrograde trans-synaptic viral infection.

Author summary

Mapping cells in the brain is a key method in neuroscience, and was traditionally carried out on manually prepared thin sections. Today, modern microscopy approaches allow the entire mouse brain to be imaged in 3D at high resolution. Due to their often complex somatic morphology, detecting cytoplasmically labelled neurons in these large image datasets is highly challenging compared, for example, to detecting spherical cell nuclei. Additionally, a neuron can often be mistakenly detected multiple times, or two cells can be interpreted as a single cell. Here we have developed a freely available algorithm for detecting cytoplasmically labelled neuronal somata in these images which can be run faster than the data can be acquired, and without the bias of manual analysis. The ability to quickly map cellular distributions throughout the mouse brain will lead to a greater understanding of both its structure and function. As with flies, nematodes and fish, detecting and mapping cells in 3D throughout the entire mammalian brain will allow for new experiments designed to understand the structural basis of its myriad complex functions.

This is a *PLOS Computational Biology* Methods paper.

Introduction

To understand the circuits underlying computations in the brain, it is necessary to map cell types, connections and activity across the entire structure. Advances in labelling [1–3], tissue clearing [4–6] and imaging [7–12] now allow for the meso- and microscopic study of brain structure and function across the rodent brain. Analysis of these whole-brain images has lagged behind the developments in imaging [13]. Although there are many relevant commercial and open-source bio-image analysis packages available [14–17], these have traditionally been developed for 2D images or for 3D volumes much smaller than a rodent brain.

In rodent studies, an increasingly common whole-brain image analysis task is the identification of individual, labelled cells across the entire brain. Traditionally, this was carried out manually [18–21], but this approach does not scale to all biological questions, particularly when many cells are labelled. Considering that a mouse brain has around 100 million neurons [22], even if only 0.01% of cells in the brain are labelled, a manual approach becomes impractical for any kind of routine analysis.

There are methods that work well for identifying labelled cells in serial 2D sections, [23–27]. However, 2D analysis can be subject to bias as detected cell numbers can be under, or overestimated depending on sampling in the third dimension. Methods have been developed for 3D cell detection in whole-brain microscopy images, but many of these were only developed for nuclear labels [28, 29]. Although nuclear labels are much simpler to detect than membrane or cytoplasmic markers (as they have a simple shape and can be approximated as spheres and are far less likely to be overlapping in the image) there are many applications in which a nuclear label is not practical or even useful, as in the case of in vivo functional imaging. Cytoplasmic cell detection methods have also been developed, but these have some key disadvantages, such as requiring multiple methods for each brain area [30], or only showing the application or validation in very small regions of the brain [24, 31]. Any approach must be applicable and validated throughout the brain, as the signal to noise (SNR) characteristics can vary between brain regions (e.g. sources of noise could be falsely detected as a cell).

As 3D whole-brain microscopy is becoming commonplace, and many neuroscientists do not have extensive computational training, it is important that any software can be installed easily, and the method can be applied quickly on standard desktop computing hardware without any programming knowledge. There does not yet exist a quick, easily applied method for 3D detection of cells with cytosolic labels that has been validated throughout an entire brain. Meeting this need is a highly desired goal within systems neuroscience.

To overcome the limitations of traditional computer vision, machine learning—and particularly deep learning [32]—has revolutionised the analysis of biomedical and cellular imaging [33]. Deep neural networks (DNNs) now represent the state of the art for the majority of image analysis, and have been applied to analyse whole-brain images, to detect cells in 2D [23, 26] or to segment axons [34]. However, they have two main disadvantages when it comes to 3D whole brain analysis. Firstly, they require large amounts of manually-annotated training data (e.g. for cell segmentation, this would potentially require the painstaking annotation of hundreds or thousands of cell borders in 3D). Secondly, the complex architecture of DNNs means that for big data (e.g. whole-brain images at cellular resolution), large amounts of computing infrastructure is required to train these networks, and then process the images in a reasonable time frame.

To harness the power of deep learning for 3D identification of labelled cells in whole-brain images, we developed a computational pipeline which uses classical image analysis approaches to detect potentially labelled cells with high sensitivity (cell candidates), at the expense of detecting false positives (i.e. geometrically similar objects). This is then followed by application of a DNN to classify cell candidates as either true cells, or artefacts to be rejected. Harnessing the power of deep learning for object classification rather than cell segmentation at a voxel level speeds up analysis (since there are billions of voxels, but many fewer cell candidates) and simplifies the generation of training data. Rather than annotating cell borders in 3D, cell candidates from the initial step can be further classified by the addition of a single (cell or artefact) label.

Results

To illustrate the problem and to demonstrate the software, whole mouse brain images were acquired following retrograde rabies virus labelling. Viral injections were performed into visual or retrosplenial cortex, causing thousands of cells to be cytoplasmically labelled throughout the brain. Data was acquired using serial two-photon microscopy as previously described [19] (Fig 1). Briefly, coronal sections are imaged at high-resolution (2 μm x 2 μm x 5 μm voxel size) and stitched to provide a complete coronal section. This is carried out for ten imaging planes after which a microtome removes the most superficial 50 μm of tissue and the process is repeated until the entire brain data set is collected. Light emitted from the specimen is filtered

Fig 1. Simplified schematic diagram of the serial two-photon microscope and data acquisition process. A: The tissue is excited using a femtosecond Ti-sapphire laser (emission wavelength = 800 nm). For data collection, 50 μm of tissue (at approximately 40 μm to 90 μm below the tissue surface) is imaged in ten, 5 μm thick planes. An in-built microtome then physically removes a 50 μm thick section from the optical face. This process is repeated to generate a complete 3D dataset of the specimen. B: For signal collection, the emitted lightpath is split into two channels whereby the primary channel detects the fluorescence signal of interest from labelled cells (e.g. mCherry at 610 nm) and the second channel (e.g. at 450 nm) detects the tissue autofluorescence signal that reveals gross anatomical structure.
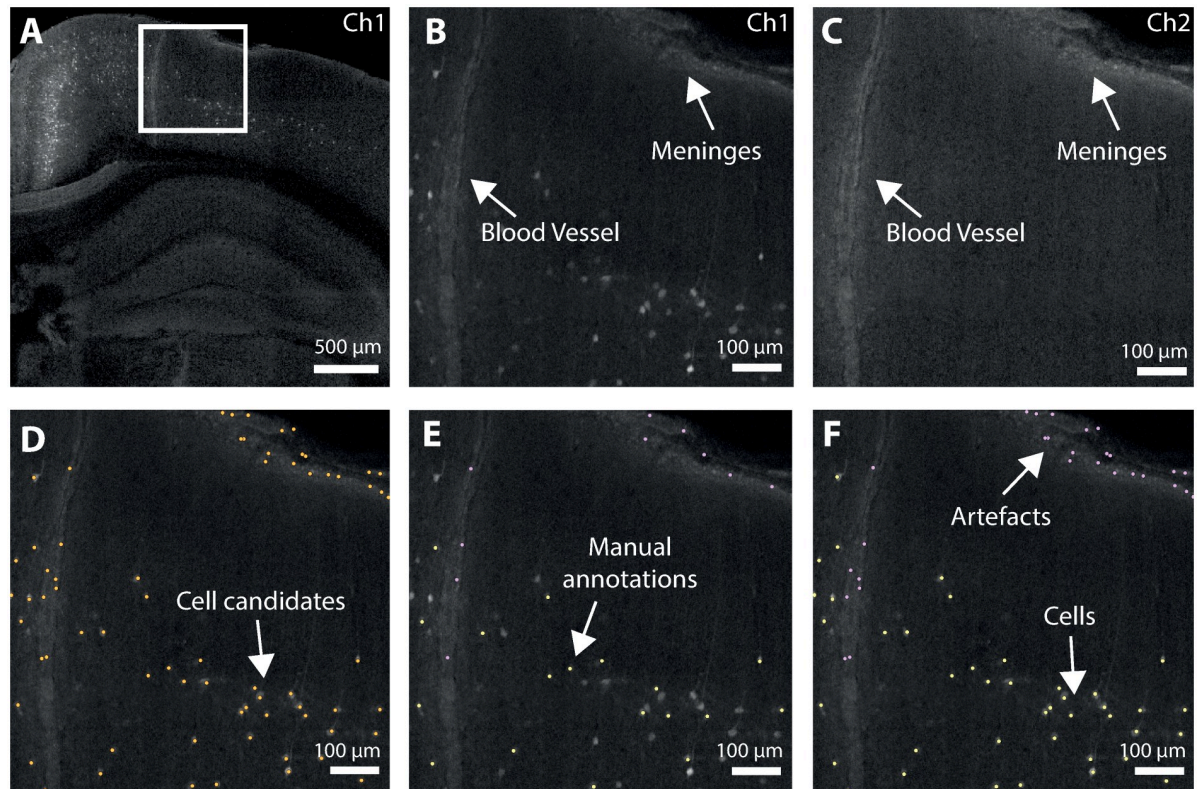
Fig 2. Illustration of the cell detection process. A: Single coronal plane of raw data (primary signal channel, Ch1). B: Enlarged insert of cortical region from A showing examples of structural features (artefacts) often erroneously detected. C: Cortical region shown in B, in the secondary autofluorescence channel (Ch2). Cells can only be seen in Ch1, but artefacts are visible in both channels. D: Detected cell candidates overlaid on raw data. Labelled cells as well as numerous artefacts are detected. E: Illustration of training data. A subset of detected cell candidates are classified as cells (yellow) or artefacts (purple). Cuboids of data centered on these selected cell candidates are then used to train the network. F: Classified cell candidates. The trained cell classification network is applied to all the cell candidates from (E) and correctly rejects the initial false positives.

https://doi.org/10.1371/journal.pcbi.1009074.g002

and detected via at least two channels, a primary signal channel containing the fluorescence signal from labelled target neurons and a secondary 'autofluorescence' channel that does not contain target signals but provides anatomical outlines. An example single-plane image is shown in Fig 2A.

Cell candidate detection

When developing any object detection algorithm, a balance must be struck between false positives and false negatives. In traditional (two-dimensional) histology, simple thresholding (e.g. [35]) can often work well for cell detection. This does not necessarily apply to whole brain images. In samples with bright, non-cellular structures (artefacts, Fig 2B) or lower signal to noise ratio, simple thresholding can detect many non-cellular elements. Image preprocessing and subsequent curation of detected objects can overcome some of these issues, but no single method works reliably across the brain in multiple samples. Either some cells are missed (false negatives), or many artefacts are also detected (false positives). To overcome this, a traditional image analysis approach was used to detect cell candidates, i.e. objects of approximately the correct brightness and size to be a cell. This list of candidates is then later refined by the deep learning step. Crucially, this refinement allows the traditional analysis to produce many false positives while minimising the number of false negatives. Images are median filtered and a

Table 1. Default cell detection parameters.

Parameter	Value
In-plane cell somata diameter (⌀)	16 μm
Gaussian smoothing sigma	0.2⌀
Intensity threshold	Mean + 10 x SD
In-plane (lateral) ellipsoidal filter width	6 μm
Axial (z) ellipsoidal filter width	15 μm
Ellipsoidal filter overlap threshold fraction	0.6

https://doi.org/10.1371/journal.pcbi.1009074.t001

Laplacian of Gaussian filter is used to enhance cell-like objects. The resulting image is thresholded, and objects of approximately the correct size are saved as candidate cell positions (Fig 2D). An overview of the cell candidate detection steps are shown in S1 and S2 Figs. The thresholding is tuned to pick up every detectable cell but this also results in the detection of many false positives that often appear as debris on the surface of the brain and in some cases unidentified objects within blood vessels. This initial detection step is based on a number of tunable parameters (Table 1) which were chosen based on trial and error to be biased to over detection, and so be relatively robust to changes in the input data. They may of course need to be tuned for very different data, in particular the in-plane cell somata diameter which may need to be tuned when different types of cells are labelled.

Cell candidate classification using deep learning

A classification step, which uses a 3D adaptation of the ResNet [36] convolutional neural network (S3 and S4 Figs) is then used to separate true from false positives. To classify cell candidates, a subset of cell candidate positions were manually annotated (e.g. Fig 2E). In total, ∼100,000 cell candidates (50,653 cells and 56,902 non-cells) were labelled from five brains. Small cuboids of 50 x 50 x 100 μm around each candidate were extracted from the primary signal channel along with the corresponding cuboid in the secondary autofluorescence channel (Fig 3A). This allows the network to "learn" the difference between neuron-based signals (only present in the primary signal channel), and other non-neuronal sources of fluorescence (potentially present in both channels).

The trained classification network is then applied to classify the cell candidates from the initial detection step (Fig 2F). The artefacts (such as those at the surface of the brain and in vessels) have been correctly rejected, while correctly classifying the labelled cells. To quantify the performance of the classification network, and to assess how much training data is required for good performance, the manually annotated training data was split up into a new training dataset from four brains, and a test dataset from the fifth brain. A new network was trained on subsets of the training data, and performance tested on the fifth brain (15,872 cells and 18,168 non-cells). Fig 3B shows that relatively little training data was required for good performance on unseen test data, with 95% of cell candidates classified correctly with ∼7,000 annotated cell candidates. Although ∼7,000 data points are required to train the network from scratch, we provide the network trained on the full dataset with the software. Users can then re-train this network with a much smaller amount of experiment-specific training data.

Application

To illustrate the method, the cell detection software was applied to data which was not used to develop or train the classification network. Data from a previous experiment [37] was acquired on a different microscope and was used to simulate real-world usage in which the SNR
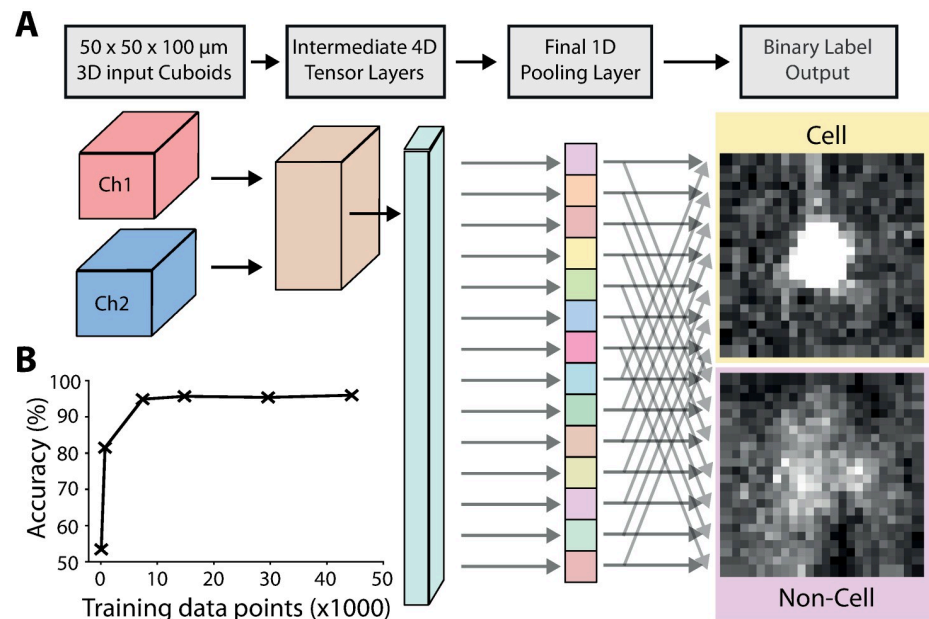
Fig 3. Cell classification. A: The input data to the modified ResNet are 3D image cuboids (50 μm x 50 μm x 100 μm) centered on each cell candidate. There are two cuboids, one from the primary signal channel, and one from the secondary autofluorescence channel. The data is then fed through the network, resulting in a binary label of cell or non-cell. During training the network "learns" to distinguish true cells, from other bright non-cellular objects. See S3 and S4 Figs for more details of the 3D ResNet architecture. B: Training the initial cell classification network: classification accuracy as a function of training data quantity.

https://doi.org/10.1371/journal.pcbi.1009074.g003

characteristics of the data may vary from that used to pre-train the supplied network. Neurons presynaptic to layer 2/3 primary visual cortical cells were labeled using rabies virus tracing (expressing mCherry) in two Penk-Cre mice.

The algorithm was run on these two brains, using the default candidate detection parameters (Table 1). A small number (619) of the detected cell candidates were manually annotated on a single brain (brain 1) including confirmations of correct classifications, and corrections of incorrect classifications. The pre-trained network was then retrained using these data points and 10% of the data was held back for validation during training. The network was trained until the validation loss function began to plateau, taking 73 minutes. The full cell detection algorithm was then repeated for both brains using the re-trained network on a laptop computer. Total time for cell detection was 83 minutes for brain 1 and 91 minutes for brain 2 (for full timings see Table 2).

To assign detected cells to a brain region, the Allen Mouse Brain Reference Atlas (ARA [38]) annotations were registered to the secondary autofluorescence channel using brainreg

Table 2. Algorithm timings on a laptop computer with Intel i9–9900K CPU, 32GB RAM and an NVIDIA RTX2080 GPU. Data stored on an external solid-state drive.

	Brain 1	Brain 2
Number of cell candidates	63995	80122
Number of candidates classified as cells	4266	4021
Time for cell candidate detection	50 minutes	48 minutes
Time for classification	33 minutes	43 minutes
Total time	83 minutes	91 minutes

https://doi.org/10.1371/journal.pcbi.1009074.t002

[39], a Python port of the validated aMAP pipeline [40]. These annotations were overlaid on the raw data (Fig 4A), and the number of cells in each brain region were reported, allowing for quantitative analysis (Table 3).

The new data was noisier than that used to pre-train the network (possibly due to the use of resonant vs galvanometer scanning), which lead to false-positives throughout the brain. A small amount of new training data, taking approximately five minutes to generate was sufficient to significantly improve performance of the algorithm (Fig 4B). The re-trained network removed many of the false-positives, while still correctly classifying the labelled cells.

Validation

To quantify the accuracy of the algorithm brain-wide, we generated ground truth data for both the brain used to generate data for re-training (brain 1) and the "unseen" brain (brain 2). Two experts manually annotated all labelled cell somata throughout the brains. These cells were assigned to regions in the ARA in the same way as the automated cell counts and an average of the two experts was taken. The comparison between the automated counts and the manual counts is shown in Fig 4C. Using the pre-trained network, the algorithm detects false positives, including in many areas with no labelled cells. Re-training the network significantly reduces the number of false positives, producing a best-fit line close to an exact match to the manual cell counts. For both brains, a linear fit to the algorithm and manual cell counts is above 1 (brain 1—1.113, brain 2—1.053), suggesting a small number of false positives (Table 4).

Although the results from the two brains look similar when the detected cells are warped to the ARA coordinate space (Fig 4D), there is still significant biological variability. For this reason, most experiments quantify relative, rather than absolute, cell counts [18, 20, 30, 41]. It is therefore important than the correlation between the automated cell counts and ground truth is as high as possible. The correlation for both the brain used for training (brain 1), and the "unseen" brain (brain 2) is very high (Pearson correlation coefficient, $\rho = 0.999$), and higher than the correlation between the automated cell counts for both brains ($\rho = 0.982$).

Effect of varying axial sampling

All the data presented was acquired with high axial sampling (5 μm), but this is not always possible or desirable due to imaging time, data storage requirements, or biological considerations such as photobleaching. To assess the performance of the algorithm with varying optical slice thicknesses, we downsampled the data from brain 2, to generate synthetic datasets with axial sampling of 5, 10, 20 and 40 μm, and the algorithm was applied as before.

The cell counts were compared to the mean of expert counts (Table 5). Performance in terms of absolute cell numbers (best-fit line slope, and Pearson correlation coefficient) was comparable for 5, 10 and 20 μm, although the 5 μm dataset was most highly correlated to the ground truth. The performance on the 40 μm dataset was much worse. The best-fit line slope of 0.675 means that many cells have been missed, although the high correlation ($\rho = 0.942$) suggests that this effect is relatively uniform throughout the brain.

The results are not surprising, as even with 20 μm axial sampling, most cells can still be visualised in multiple image planes, and so 3D cell detection is still possible. At 40 μm axial sampling, this is not possible, and so many cells are missed. Although the correlation coefficients at 10 and 20 μm axial spacing are relatively high ($\rho = 0.978, 0.981$), they are lower than the correlation between different brains ($\rho = 0.982$). This may reduce the likelihood of detecting small biological effects using data with low axial sampling. There are also many other factors that affect how many image planes a cell will appear in, such as the cell somata size, and the axial point spread function of the microscope.

Fig 4. Application of the algorithm to unseen data. Presynaptic neurons labelled by rabies viral injection into layer 2/3 primary visual cortex in Penk-Cre mice. A: Detected cells overlaid on raw data along with the brain region segmentation for Brain 1 (blue) and 2 (orange). The size of the coloured disk represents the proximity of the cell centroid to the image plane displayed. Brain regions with detected cells shown: RSP—Retrosplenial area, VISp—Primary visual area, VISl—lateral visual area, VISli—Laterointermediate area, TEa—Temporal association areas. B: Comparison of cell detection before and after re-training the pre-trained network in different regions of the image. Cells with different morphologies are correctly detected in both dense and sparse regions, and artefacts are rejected. VISli—Laterointermediate area, LGd—Dorsal part of the lateral geniculate complex, DR—Dorsal nucleus raphe, SC/M—Superior colliculus & meninges. C: Comparison of cell counts per ARA brain region between the algorithm and the mean of the two expert counts. Lighter points represent results before re-training and darker points after re-training. Best fit shown after re-training. Five regions with the highest number of detected cells coloured, VISp2/3—Primary visual area, layer 2/3, VISp5—Primary visual area, layer 5, LGd-co—Dorsal part of the lateral geniculate complex, core, LP—Lateral posterior nucleus of the thalamus, VISp4—Primary visual area, layer 4. D: Visualisation of detected cells from both brains warped to the ARA coordinate space in 3D, along with the rabies virus injection site target (primary visual cortex, wireframe).

Table 3. Total number of cells in each region, per brain, projecting to layer 2/3 neurons in primary visual cortex.
Ten regions with the greatest number of cells across both brains shown.

Brain structure name	Brain 1	Brain 2
Primary visual area, layer 2/3	965	1222
Primary visual area, layer 5	650	558
Dorsal part of the lateral geniculate complex, core	371	340
Lateral posterior nucleus of the thalamus	240	215
Primary visual area, layer 4	207	205
Retrosplenial area, ventral part, layer 5	162	135
Dorsal part of the lateral geniculate complex, shell	122	124
Lateral dorsal nucleus of thalamus	122	107
Retrosplenial area, dorsal part, layer 5	110	103
Posteromedial visual area, layer 5	46	74

https://doi.org/10.1371/journal.pcbi.1009074.t003

Table 4. Comparison between algorithm and expert cell counting.

	Brain 1		Brain 2	
	Pearson correlation coefficient	Linear best-fit slope	Pearson correlation coefficient	Linear best-fit slope
Algorithm vs expert 1	0.999	1.054	0.998	0.964
Algorithm vs expert 2	0.999	1.178	0.997	1.151
Expert 1 vs expert 2	0.999	1.116	0.994	1.188
Algorithm vs expert mean	0.999	1.113	0.999	1.053

https://doi.org/10.1371/journal.pcbi.1009074.t004

Table 5. Algorithm performance compared to mean of expert manual counts with varying axial sampling.

Axial sampling [μm]	Pearson correlation coefficient	Linear best-fit slope
5	0.999	1.053
10	0.978	0.906
20	0.981	1.011
40	0.942	0.675

https://doi.org/10.1371/journal.pcbi.1009074.t005

Discussion

Mapping the distribution of labelled neurons across the brain is critical for a complete understanding of the information pathways that underlie brain function. Many existing methods for cell detection in whole-brain images rely on classical image processing, which can be affected by noise, and may not detect complex cell morphologies. DNNs can be used for highly sensitive image processing, but often require laborious generation of training data and are prohibitively slow for the analysis of large, 3D images. The presented method here overcomes these limitations by combining traditional image processing methods for speed, with a DNN to improve accuracy.

Recent developments in microscopy technology (e.g. [12]) now allow for quicker, more routine acquisition of whole-brain datasets. It is important that the image analysis can be carried out in a timely fashion, and without relying on large-scale computing infrastructure. Processing time for the ∼ 180GB images in Fig 4 on a laptop was around 90 minutes, so sixteen datasets could be analysed in a single day, much quicker than the sample preparation and

imaging steps. Once parameters are optimised, and the classification network is trained, the software can run entirely without user intervention.

In traditional DNN approaches for image analysis, generation of training data is often a major bottleneck. While large-scale "citizen science" approaches can be used to generate large amounts of training data [42], this is not practical for the majority of applications, e.g. when anatomical expertise is required. Our method overcomes this by requiring only a binary label (cell or non-cell) for each cell candidate in the training dataset, rather than a painstaking 3D outline of each cell. The software is released with a pre-trained network, and so the network can be re-trained for specific datasets very quickly. The total time to generate the new network used to classify images in Fig 4 was less than two hours, including generating training data and retraining the network.

We show that the results of the proposed method compare well to expert manual counts, particularly the correlation between counts in different brain areas. However our results also show the importance of re-training the pre-trained network for new datasets (even if they are superficially similar). It is important to also bear in mind that although we quantified performance of the algorithm across the brain, we did not label every type of cell, and some areas of the brain with densely packed neurons (e.g. hippocampus) were only sparsely labelled. When applying this method to very different data, users should ensure that they re-train the model with representative training data (including different brain areas, cell types and image artefacts if applicable), and check the results in detail. The data in Table 5 shows that although the method is relatively robust to the axial sampling, true 3D imaging (i.e. labelled cells appear in at least two axial planes) is required for accurate cell detection using our method.

The ability to quickly detect, visualise and analyse cytoplasmically labelled cells across the mouse brain brings a number of advantages over existing methods. Analysing an entire brain rather than 2D sections has the potential to detect many more cells, increasing the statistical power and the likelihood of finding novel results, particularly when studying rare cell types. Whole-brain analysis may also be less biased than analysing a series of 2D planes, especially in regions with low cell densities, or differing cell sizes.

In future we aim to adapt the network to be flexible as to the number of input channels, and output labels. The classification network currently relies on using both the primary signal and the secondary autofluroescence channel. In the future it would be valuable to train a network that could achieve a similar level of performance using a single input channel. Analysing a single channel would allow half as much data to be collected (although autofluorescence channels are optimal for atlas registration). Training a network to produce multiple labels (rather than just cell or non-cell) would allow for cell-type classification based on morphology, or based on gene or protein expression levels if additional signal channels were supplied. Although the ResNet architecture was chosen based on performance [36] and flexibility in new contexts (e.g. [43]), there are many newer network architectures that could be implemented to improve performance (e.g. [44, 45]). Different types of network (such as a 3D U-Net [46]) could also be used to segment the cell boundaries for additional analyses such as measuring cell size and shape, or quantifying fluorescence expression. Lastly, although this approach was designed for fast analysis of large whole-brain datasets, the proposed two-step approach could be used for any kind of large-scale 3D object detection.

This software is fully open-source, and has been written with further development and collaboration in mind. Although the algorithm was designed for neuron detection, the software could be applied to any 3D image data, and we provide a graphical user interface (as a napari [47] plugin) to allow easy adoption by as many users as possible.

Materials and methods

Ethics statement

All experiments were carried out in accordance with the UK Home Office regulations (Animal Welfare Act 2006) and approved by the establishments Animal Welfare and Ethical Review Board.

Sample preparation

All mice used were transgenic Cre-reporter (Ntsr1-Cre, GAD2-IRES-Cre, Rbp4-Cre & Penk-Cre) mice bred on a C57BL/6 background. The mice were anesthetized and an AAV Cre-dependent helper virus encoding both the envelope protein receptor and the rabies virus gly-coprotein was stereotactically injected into visual cortex or retrosplenial cortex. Four days later, a glycoprotein deficient form of the rabies virus expressing mCherry was delivered into the same site. After ten further days, the animal was deeply anaesthetized and transcardially perfused with cold phosphate buffer (0.1 M) followed by 4% paraformaldehyde (PFA) in PB (0.1 M) and the brain left overnight in 4% PFA at 4˚C.

Imaging

All data was acquired using serial section two-photon tomography [8]. To generate the data to pre-train the deep-learning model, fixed brains were embedded in 4% agar and placed under a two-photon microscope containing an integrated vibrating microtome and a motorized x-y-z stage as previously described [19]. Coronal images were acquired via two optical pathways (red and blue) as a set of 6 by 9 tiles with a voxel size of 1 μm x 1 μm obtained every 5 μm using an Olympus 10x objective (NA = 0.6) mounted on a piezoelectric element (Physik Instrumente, Germany). Following acquisition, image tiles were corrected for uneven illumination by sub-traction of an average image from each physical section. Tiles were then stitched using a cus-tom FIJI [14] plugin (modified from [48]) and downsampled to 2 μm x 2 μm x 5 μm voxel size.

To generate the data to test the algorithm, data was acquired using a different, custom-built resonant-scanning system controlled by ScanImage (v5.6, Vidrio Technologies, USA) using BakingTray [49], a custom software wrapper for setting up the imaging parameters. Images were assembled using StitchIt [50]. Both brains were imaged in a single acquisition using a Nikon 16x objective (NA = 0.8), with a voxel size of 2.31 μm x 2.31 μm x 5 μm.

Cell candidate detection

To detect cell candidates (broadly defined as anything of sufficient brightness and of approxi-mately the correct size to be a cell), initially data from the primary signal channel was pro-cessed in 2D (S1 Fig). Images were median filtered, and then a Laplacian of Gaussian filter was performed to enhance small, bright structures (e.g. cells). This filtered image was binarised using a threshold calculated for each image plane (mean of image plane + 10 x image plane standard deviation). The thresholded image was then passed to an ellipsoidal filter to remove noise. Every position of this spatial filter in which the majority (given by an input parameter, ellipsoidal filter overlap threshold fraction) of the filter overlaps with thresholded voxels was saved as a potential cell candidate. This is used to remove noise (from e.g. neurites).

All cell candidates that form continuous spatial structures were merged together, and classed as a single cell candidate. If the resulting cluster was too large to be a single cell (based on the input cell somata size parameter), then this cluster was split into individual cell candi-dates using an iterative ellipsoidal filter. Briefly, the ellipsoidal filter was applied to all voxels within the cell candidate cluster, and any resulting cell candidate coordinate positions were

recorded. The thresholded image was eroded, and the filter was reapplied on the eroded set of candidate voxels. This process was repeated for ten iterations, or until there were no cell candidates remaining. This process ensures that densely labelled cells are split into individual cell candidates.

Once the final list of candidate cells is determined, the centroid of each cell candidate (based on the 3D mean coordinate of thresholded voxels) was calculated, and the coordinates were saved. All of these steps were all carried out using the default software parameters (Table 1), with the exception of the 40 μm axial spacing dataset for which the axial extent of the ellipsoid filter was increased to 30 μm.

Cell candidate classification using deep learning

Cell candidates were classified using a ResNet [36], implemented in Keras [51] for TensorFlow [52]. 3D adaptations of all networks from the original paper are implemented in the software (i.e. 18, 34, 50, 101 and 152-layer) and can be chosen by the user, but the 50-layer network was used throughout this study. The general architecture of these networks is shown in S3 and S4 Figs.

To generate data to train the classification network, output from the candidate detection step (cell candidate coordinates) were manually classified using a custom FIJI [14] plugin, or a napari [47] plugin that is supplied with the software. Expert annotators were presented with the raw data, with cell candidates represented by hollow spheres, centered on the cell candidate. By scrolling through the 2D planes of the 3D dataset, the experts could view the cell candidate in 3D before marking it as a cell or artefact. Candidates were determined to be cells or artefacts based on size, shape (including the presence of neurites) and the fluorescence level compared to the secondary autofluorescence channel (which was visualised at the same time). Candidates were labelled by three experts, labelling different brains, and a subset of labelled candidates were cross-checked between experts.

Image cuboids of 50 μm x 50 μm x 100 μm (resampled to 50 x 50 x 20 voxels) were extracted from both the primary signal, and secondary autofluorescence channels, centered on the coordinates of the manually classified cell candidate positions. To increase the size of the training set, data were randomly augmented. Each of the following transformations were applied with a 10% likelihood: (i) flipping in any of the three axes, (ii) rotation around any of the three axes (between −45˚ to 45˚) and (iii) circular translation along any of the three axes (up to 5% of the axis length). The networks were trained using an NVIDIA TITAN RTX GPU with a batch size of 32 and the Adam [53] method was used to minimise the loss (categorical cross entropy), with a learning rate of 0.0001. Cell candidates were classified using the trained network, and saved as an XML file with a cell or artefact label.

Image registration and segmentation

To allow detected cells to be assigned an anatomical label, and for them to be analysed in a common coordinate framework, a reference atlas (Allen Mouse Brain Atlas, ARA, [38], provided by the BrainGlobe Atlas API [54]) was registered to the autofluorescence channel. This was carried out using brainreg [39], a Python port of the automatic mouse atlas propagation (aMAP) software [40], which itself relies on the medical image registration library, NiftyReg [55]. Firstly the sample brain was downsampled to the same voxel spacing as the atlas (10 μm isotropic) and was reoriented to match the atlas template brain. These two images were then filtered to remove-high frequency noise (greyscale opening and flat-field correction). The images were firstly registered using an affine transform (*reg_aladin* [56]), followed by freeform non-linear registration (*reg_f3d* [55]). The resulting transformation was applied to the atlas

brain region annotations (and a custom hemispheres atlas) to bring it into the same coordinate space as the sample brain.

Validation

To compare results of the algorithm to ground truth, two experts manually annotated each cell in two whole-brain images using the same napari plugin used to generate the training data. The experts were shown the full-resolution images, with both channels displayed in different colors. The experts could scroll through the 3D images plane by plane, zoom in and out, and adjust contrast settings to best visualise cells in different brain areas. Experts annotated cells based on the same criteria as for generating the training data (shape, size and fluorescence intensity). Annotating a cell position displayed a hollow sphere on top of the image, around the cell coordinate. This was visible in multiple image planes to ensure that individual cells were not labelled more than once.

The cell classification network was retrained for the new datasets, and the algorithm was run with the default cell candidate detection parameters. The images were also registered to the ARA, and cell coordinates were assigned to brain regions for both the automated and manual cell counts. The different cell counting approaches were firstly compared by calculating the Pearson correlation coefficient using Pandas [57, 58]. To assess the bias of the different approaches, they were compared by calculating the slope of the best fit line by fitting a linear model using scikit-learn [59].

Effect of varying axial sampling

To generate synthetic datasets with varying axial sampling, a single brain was downsampled in 3D by selecting every Nth image plane. For example, to generate a dataset with 20 μm sampling from the original dataset sampled at 5 μm, every fourth plane was used.

Visualisation

For visualisation of data in standard space, detected cells must be transformed to the atlas coordinate space. Firstly, the affine transform from the initial registration was inverted (using *reg_transform*). The sample brain was then registered non-linearly to the atlas (again using *reg_f3d*) and a deformation field (mapping points in the sample brain to the atlas) was generated (using *reg_transform*). This deformation field was applied to the coordinates of the detected cells for each sample, transforming them into atlas coordinate space.

Plots were generated using Matplotlib [60], and image visualisation was performed using napari [47] and brainrender [61].

Software implementation

The methods outlined in this manuscript are available within the cellfinder software, part of the BrainGlobe suite of computational neuroanatomy tools. The software is open-source, written in Python 3 and runs on standard desktop computing hardware (although a CUDA compatible GPU allows for a considerable reduction in processing time). Source code is available at github.com/brainglobe/cellfinder-core and pre-built wheels at pypi.org/project/cellfinder-core. To aid adoption of the method, a plugin for napari [47] is available at github.com/brainglobe/cellfinder-napari, and an integrated pipeline for the analysis of whole-brain microscopy data (including e.g. atlas registration) is available at github.com/brainglobe/cellfinder. Documentation, tutorials, and the data underlying Fig 4 are available at docs.brainglobe.info.

Supporting information

S1 Fig. Overview of the initial cell candidate detection steps, from raw data to a cuboid of data fed into the classification network. Upper row: from left to right, the raw image is median filtered to remove noise. A Laplacian of Gaussian is then performed to enhance small, bright structures such as the cell soma. Lower row: from right to left, the image is thresholded and a 3D ellipsoidal filter is used to remove small, non-cellular objects (not shown in this image plane). The centroid of the resulting object is then used to center the cuboid of data that it passed to the deep learning classification network. Images shown are 100 μm x 100 μm, and the cuboid is 50 μm x 50 μm (and 100 μm in the third dimension).
(TIF)

S2 Fig. Assignment of signal to individual cell candidates in the third dimension. Saggital sections (i.e. orthogonal to acquisition) shown. Upper row: from left to right, median and Laplacian of Gaussian filtering, as in S1 Fig. Lower row: from right to left, thresholding, filtering and centroid calculation, as in S1 Fig. Unlike 2D analysis, individual objects (cells) are correctly distinguished, and not merged, or erroneously split. Images shown are 205 μm x 143 μm, and the cuboid is 50 μm x 100 μm (and 50 μm in the third dimension).
(TIF)

S3 Fig. Architecture of the 3D ResNet. 3D adaptation of the 2D networks from [38] which are available for use in the software.
(TIF)

S4 Fig. Architecture of the bottleneck 3D ResNet. 3D adaptation of the bottleneck 2D networks from [38] which are available for use in the software. The 50-layer bottleneck network is used throughout this study, and is used for the pre-trained model supplied with the software.
(TIF)

Author Contributions

Conceptualization: Adam L. Tyson, Charly V. Rousseau, Christian J. Niedworok, Troy W. Margrie.

Data curation: Adam L. Tyson, Charly V. Rousseau, Christian J. Niedworok.

Formal analysis: Adam L. Tyson, Lee Cossell, Troy W. Margrie.

Investigation: Sepiedeh Keshavarzi, Chryssanthi Tsitoura, Molly Strom.

Methodology: Adam L. Tyson, Charly V. Rousseau, Christian J. Niedworok, Sepiedeh Keshavarzi, Chryssanthi Tsitoura.

Project administration: Troy W. Margrie.

Software: Adam L. Tyson, Charly V. Rousseau, Christian J. Niedworok.

Validation: Adam L. Tyson, Lee Cossell.

Visualization: Adam L. Tyson.

Writing – original draft: Adam L. Tyson, Troy W. Margrie.

Writing – review & editing: Adam L. Tyson, Charly V. Rousseau, Christian J. Niedworok, Sepiedeh Keshavarzi, Chryssanthi Tsitoura, Lee Cossell, Molly Strom, Troy W. Margrie.

References

1. Wickersham IR, Finke S, Conzelmann KK, Callaway EM. Retrograde neuronal tracing with a deletion-mutant rabies virus. Nature Methods. 2007; 4(1):47–49. https://doi.org/10.1038/NMETH999 PMID: 17179932

2. Reijmers LG, Perkins BL, Matsuo N, Mayford M. Localization of a stable neural correlate of associative memory. Science. 2007; 317(5842):1230–1233. https://doi.org/10.1126/science.1143839 PMID: 17761885

3. Kim S, Cho JH, Murray E, Bakh N, Choi H, Ohn K, et al. Stochastic electrotransport selectively enhances the transport of highly electromobile molecules. Proceedings of the National Academy of Sciences. 2015; p. E6274–E6283. https://doi.org/10.1073/pnas.1510133112 PMID: 26578787

4. Ertürk A, Becker K, Jährling N, Mauch CP, Hojer CD, Egen JG, et al. Three-dimensional imaging of solvent-cleared organs using 3DISCO. Nature Protocols. 2012; 7(11):1983–1995. https://doi.org/10.1038/nprot.2012.119 PMID: 23060243

5. Chung K, Wallace J, Kim S, Kalyanasundaram S, Andalman AS, Davidson TJ, et al. Structural and molecular interrogation of intact biological systems. Nature. 2013; 497(7449):332–337. https://doi.org/10.1038/nature12107 PMID: 23575631

6. Susaki EA, Tainaka K, Perrin D, Kishino F, Tawara T, Watanabe TM, et al. Whole-Brain Imaging with Single-Cell Resolution Using Chemical Cocktails and Computational Analysis. Cell. 2014; 157(3):726–739. https://doi.org/10.1016/j.cell.2014.03.042 PMID: 24746791

7. Dodt H, Leischner U, Schierloh A. Ultramicroscopy: three-dimensional visualization of neuronal networks in the whole mouse brain. Nature Methods. 2007; 4(4):331–336. https://doi.org/10.1038/nmeth1036 PMID: 17384643

8. Ragan T, Kadiri LR, Venkataraju KU, Bahlmann K, Sutin J, Taranda J, et al. Serial two-photon tomography for automated ex vivo mouse brain imaging. Nature Methods. 2012; 9(3):255–258. https://doi.org/10.1038/nmeth.1854 PMID: 22245809

9. Osten P, Margrie TW. Mapping brain circuitry with a light microscope. Nature Methods. 2013; 10 (6):515–523. https://doi.org/10.1038/nmeth.2477 PMID: 23722211

10. Fu Q, Martin BL, Matus DQ, Gao L. Imaging multicellular specimens with real-time optimized tiling light-sheet selective plane illumination microscopy. Nature Communications. 2016; 7:11088. https://doi.org/10.1038/ncomms11088 PMID: 27004937

11. Seiriki K, Kasai A, Hashimoto T, Schulze W, Niu M, Yamaguchi S, et al. High-Speed and Scalable Whole-Brain Imaging in Rodents and Primates. Neuron. 2017; 94(6):1085–1100.e6. https://doi.org/10.1016/j.neuron.2017.05.017 PMID: 28641108

12. Voigt FF, Kirschenbaum D, Platonova E, Pagès S, Campbell RAA, Kastli R, et al. The mesoSPIM initiative: open-source light-sheet microscopes for imaging cleared tissue. Nature Methods. 2019; 16:1105–1106. https://doi.org/10.1038/s41592-019-0554-0 PMID: 31527839

13. Tyson AL, Margrie TW. Mesoscale microscopy for micromammals: image analysis tools for understanding the rodent brain. arXiv. 2021;(2102.11812).

14. Schindelin J, Arganda-Carreras I, Frise E, Kaynig V, Longair M, Pietzsch T, et al. Fiji: an open-source platform for biological-image analysis. Nature Methods. 2012; 9(7):676–682. https://doi.org/10.1038/nmeth.2019 PMID: 22743772

15. De Chaumont F, Dallongeville S, Chenouard N, Hervé N, Pop S, Provoost T, et al. Icy: An open bio-image informatics platform for extended reproducible research. Nature Methods. 2012; 9(7):690–696. https://doi.org/10.1038/nmeth.2075 PMID: 22743774

16. McQuin C, Goodman A, Chernyshev V, Kamentsky L, Cimini BA, Karhohs KW, et al. CellProfiler 3.0: Next-generation image processing for biology. PLoS Biology. 2018; 16(7):1–17. https://doi.org/10.1371/journal.pbio.2005970 PMID: 29969450

17. Berg S, Kutra D, Kroeger T, Straehle CN, Kausler BX, Haubold C, et al. Ilastik: Interactive Machine Learning for (Bio)Image Analysis. Nature Methods. 2019; 16:1226–1232. https://doi.org/10.1038/s41592-019-0582-9 PMID: 31570887

18. Watabe-Uchida M, Zhu L, Ogawa SK, Vamanrao A, Uchida N. Whole-Brain Mapping of Direct Inputs to Midbrain Dopamine Neurons. Neuron. 2012; 74(5):858–873. https://doi.org/10.1016/j.neuron.2012.03.017 PMID: 22681690

19. Vélez-Fort M, Rousseau CV, Niedworok CJ, Wickersham IR, Rancz EA, Brown APY, et al. The stimulus selectivity and connectivity of layer six principal cells reveals cortical microcircuits underlying visual processing. Neuron. 2014; 83(6):1431–1443. https://doi.org/10.1016/j.neuron.2014.08.001 PMID: 25175879

20. Ogawa SK, Cohen JY, Hwang D, Uchida N, Watabe-Uchida M. Organization of monosynaptic inputs to the serotonin and dopamine neuromodulatory systems. Cell Reports. 2014; 8(4):1105–1118. https://doi.org/10.1016/j.celrep.2014.06.042 PMID: 25108805

21. Schwarz MK, Scherbarth A, Sprengel R, Engelhardt J, Theer P, Giese G. Fluorescent-Protein Stabilization and High-Resolution Imaging of Cleared, Intact Mouse Brains. PloS one. 2015; 10(5):e0124650. https://doi.org/10.1371/journal.pone.0124650 PMID: 25993380

22. Herculano-Houzel S, Mota B, Lent R. Cellular scaling rules for rodent brains. Proceedings of the National Academy of Sciences of the United States of America. 2006; 103(32):12138–12143. https://doi.org/10.1073/pnas.0604911103 PMID: 16880386

23. Kim Y, Venkataraju KU, Pradhan K, Mende C, Taranda J, Turaga SC, et al. Mapping social behavior-induced brain activation at cellular resolution in the mouse. Cell Reports. 2015; 10(2):292–305. https://doi.org/10.1016/j.celrep.2014.12.014 PMID: 25558063

24. Furth D, Vaissiere T, Tzortzi O, Xuan Y, Martin A, Lazaridis I, et al. An interactive framework for whole-brain maps at cellular resolution. Nature Neuroscience. 2018; 21:139–149. https://doi.org/10.1038/s41593-017-0027-7 PMID: 29203898

25. Salinas CBG, Lu TTH, Gabery S, Marstal K, Alanentalo T, Mercer AJ, et al. Integrated Brain Atlas for Unbiased Mapping of Nervous System Effects Following Liraglutide Treatment. Scientific Reports. 2018; 8(1):1–12. https://doi.org/10.1038/s41598-018-28496-6 PMID: 29985439

26. Iqbal A, Sheikh A, Karayannis T. DeNeRD: high-throughput detection of neurons for brain-wide analysis with deep learning. Scientific Reports. 2019; 9(1):1–13. https://doi.org/10.1038/s41598-019-50137-9 PMID: 31554830

27. Song JH, Choi W, Song YH, Kim JH, Jeong D, Lee SH, et al. Precise Mapping of Single Neurons by Calibrated 3D Reconstruction of Brain Slices Reveals Topographic Projection in Mouse Visual Cortex. Cell Reports. 2020; 31(8):107682. https://doi.org/10.1016/j.celrep.2020.107682 PMID: 32460016

28. Renier N, Adams EL, Kirst C, Wu Z, Azevedo R, Kohl J, et al. Mapping of Brain Activity by Automated Volume Analysis of Immediate Early Genes. Cell. 2016; 165(7):1789–1802. https://doi.org/10.1016/j.cell.2016.05.007 PMID: 27238021

29. Young DM, Duhn C, Gilson M, Nojima M, Yuruk D, Kumar A, et al. Whole-Brain Image Analysis and Anatomical Atlas 3D Generation Using MagellanMapper. Current protocols in neuroscience. 2020; 94 (1):e104. https://doi.org/10.1002/cpns.104 PMID: 32981139

30. Menegas W, Bergan JF, Ogawa SK, Isogai Y, Umadevi Venkataraju K, Osten P, et al. Dopamine neurons projecting to the posterior striatum form an anatomically distinct subclass. eLife. 2015; 4:1–30. https://doi.org/10.7554/eLife.10032 PMID: 26322384

31. Goubran M, Leuze C, Hsueh B, Aswendt M, Ye L, Tian Q, et al. Multimodal image registration and connectivity analysis for integration of connectomic data from microscopy to MRI. Nature Communications. 2019; 10(1):1–17. https://doi.org/10.1038/s41467-019-13374-0

32. Lecun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015; 521:436–444. https://doi.org/10.1038/nature14539 PMID: 26017442

33. Moen E, Bannon D, Kudo T, Graf W, Covert M, Van Valen D. Deep learning for cellular image analysis. Nature Methods. 2019. https://doi.org/10.1038/s41592-019-0403-1 PMID: 31133758

34. Friedmann D, Pun A, Adams EL, Lui JH, Kebschull JM, Grutzner SM, et al. Mapping mesoscale axonal projections in the mouse brain using a 3D convolutional network. Proceedings of the National Academy of Sciences. 2020; 117(20):11068–11075. https://doi.org/10.1073/pnas.1918465117 PMID: 32358193

35. Otsu N. A threshold selection method from gray-level histograms. IEEE Transactions on Systems, Man, and Cybernetics. 1979; 9(1):62–66. https://doi.org/10.1109/TSMC.1979.4310076

36. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. vol. 2016-Decem; 2016. p. 770–778.

37. Brown APY, Cossell L, Strom M, Tyson AL, Vélez-Fort M, Margrie TW. Analysis of segmentation ontology reveals the similarities and differences in connectivity onto L2/3 neurons in mouse V1. Scientific Reports. 2021; 11 (4983). https://doi.org/10.1038/s41598-021-82353-7 PMID: 33654118

38. Wang Q, Ding SL, Li Y, Royall J, Feng D, Lesnar P, et al. The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. Cell. 2020; 181(4):936–953.e20. https://doi.org/10.1016/j.cell.2020.04.007 PMID: 32386544

39. Tyson AL, Rousseau CV, Margrie TW. brainreg: automated 3D brain registration with support for multiple species and atlases; 2020. Available from: https://doi.org/10.5281/zenodo.3991718.

40. Niedworok CJ, Brown APY, Jorge Cardoso M, Osten P, Ourselin S, Modat M, et al. AMAP is a validated pipeline for registration and segmentation of high-resolution mouse brain data. Nature Communications. 2016; 7:1–9. https://doi.org/10.1038/ncomms11879 PMID: 27384127

41. Vélez-Fort M, Bracy EF, Keshavarzi S, Rousseau CV, Cossell L, Lenzi SC, et al. A Circuit for Integration of Head- and Visual-Motion Signals in Layer 6 of Mouse Primary Visual Cortex. Neuron. 2018; 98 (1):179–191. https://doi.org/10.1016/j.neuron.2018.02.023 PMID: 29551490

42. Spiers H, Songhurst H, Nightingale L, de Folter J, Hutchings R, Peddie CJ, et al. Citizen science, cells and CNNs—deep learning for automatic segmentation of the nuclear envelope in electron microscopy data, trained with volunteer segmentations. bioRxiv. 2020.

43. Mathis A, Mamidanna P, Cury KM, Abe T, Murthy VN, Mathis MW, et al. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. Nature Neuroscience. 2018; 21(9):1281–1289. https://doi.org/10.1038/s41593-018-0209-y PMID: 30127430

44. Huang G, Liu Z, Weinberger KQ. Densely Connected Convolutional Networks. CoRR. 2016;abs/1608.06993.

45. Howard AG, Zhu M, Chen B, Kalenichenko D, Wang W, Weyand T, et al. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. CoRR. 2017;abs/1704.04861.

46. Çiçek Ö, Abdulkadir A, Lienkamp SS, Brox T, Ronneberger O. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In: Ourselin S, Joskowicz L, Sabuncu MR, Unal G, Wells W, editors. Medical Image Computing and Computer-Assisted Intervention—MICCAI 2016. Cham: Springer International Publishing; 2016. p. 424–432.

47. Sofroniew N, Lambert T, Evans K, Nunez-Iglesias J, Winston P, Bokota G, et al. napari/napari: 0.4.8; 2021. Available from: https://doi.org/10.5281/zenodo.4741360.

48. Preibisch S, Saalfeld S, Tomancak P. Globally optimal stitching of tiled 3D microscopic image acquisitions. Bioinformatics. 2009; 25(11):1463–1465. https://doi.org/10.1093/bioinformatics/btp184 PMID: 19346324

49. Campbell RAA. BakingTray: Serial-section automated anatomy extension for ScanImage; 2020. Available from: https://github.com/SainsburyWellcomeCentre/BakingTray.

50. Campbell RAA, Blot A, lguerard. StitchIt: Stitching of large tiled datasets; 2020. Available from: https://github.com/SainsburyWellcomeCentre/StitchIt.

51. Chollet F, et al. Keras; 2015. https://keras.io.

52. Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems; 2015. Available from: https://www.tensorflow.org/.

53. Kingma DP, Ba J. Adam: A Method for Stochastic Optimization. arXiv. 2014.

54. Claudi F, Petrucco L, Tyson AL, Branco T, Margrie TW, Portugues R. BrainGlobe Atlas API: a common interface for neuroanatomical atlases. Journal of Open Source Software. 2020; 5(54):2668. https://doi.org/10.21105/joss.02668

55. Modat M, Ridgway GR, Taylor ZA, Lehmann M, Barnes J, Hawkes DJ, et al. Fast free-form deformation using graphics processing units. Computer Methods and Programs in Biomedicine. 2010; 98(3):278–284. https://doi.org/10.1016/j.cmpb.2009.09.002 PMID: 19818524

56. Ourselin S, Roche A, Subsol G, Pennec X, Ayache N. Reconstructing a 3D structure from serial histological sections. Image and Vision Computing. 2001; 19(1-2):25–31. https://doi.org/10.1016/S0262-8856(00)00052-4

57. Wes McKinney. Data Structures for Statistical Computing in Python. In: Stéfan van der Walt, Jarrod Millman, editors. Proceedings of the 9th Python in Science Conference; 2010. p. 56–61.

58. Reback J, McKinney W, jbrockmendel, den Bossche JV, Augspurger T, Cloud P, et al. pandas-dev/pandas: Pandas 1.2.2; 2021. Available from: https://doi.org/10.5281/zenodo.4524629.

59. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research. 2011; 12:2825–2830.

60. Hunter JD. Matplotlib: A 2D graphics environment. Computing in Science & Engineering. 2007; 9 (3):90–95. https://doi.org/10.1109/MCSE.2007.55

61. Claudi F, Tyson AL, Petrucco L, Margrie TW, Portugues R, Branco T. Visualizing anatomically registered data with brainrender. eLife. 2021. https://doi.org/10.7554/eLife.65751 PMID: 33739286