

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

Adam L. Tyson^{1*}, Charly V. Rousseau^{1,2*}, Christian J. Niedworok^{1*}, Sepiedeh Keshavarzi¹, Chryssanthi Tsitoura¹, and Troy W. Margrie^{1,✉}

¹Sainsbury Wellcome Centre, University College London, 25 Howland Street, London, W1T 4JG, United Kingdom

²Present address: Institute Pasteur, 25 Rue du Dr Roux, 75015 Paris, France

*Equal contributions

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.

Correspondence: t.margrie@ucl.ac.uk

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. Although there are many relevant commercial and open-source bio-image analysis packages available (13–16), 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 (17–20), 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 (21), even if only 0.01% of cells in the brain are labelled, a manual approach becomes impractical for any kind of routine analysis. Furthermore, considering the difficulty and computational complexity of 3D analysis, a number of studies have used automated approaches to segment neurons across the brain in serial 2D sections (22–27).

These existing methods can suffer from several sources of bias. Manual analysis is subjective, and 2D analysis can under, or overestimate cell numbers depending on sampling in the third dimension. To our knowledge, one study has shown unbiased 3D cell identification in whole-brain images (28),

however in this case nuclear labels were used. 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. Detecting cells using cytosolic reporters in 3D throughout an entire brain therefore remains an unmet and highly desired goal within systems neuroscience. To overcome the limitations of traditional computer vision, machine learning — and particularly deep learning (29) — has revolutionised the analysis of biomedical and cellular imaging (30). 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 (22, 27) or to segment axons (31). 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 retro-

grade rabies virus labelling. Viral injections were performed into visual or retrosplenial cortex, causing many thousands of cells to be cytoplasmically labelled throughout the brain. Data was acquired using serial two-photon microscopy as previously described (18) (Fig. 1). Briefly, coronal sections are imaged at high-resolution ($2\mu\text{m} \times 2\mu\text{m} \times 5\mu\text{m}$ 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 $50\mu\text{m}$ of tissue and the process is repeated until the entire brain data set is collected. Light emitted from the specimen is filtered 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.




Fig. 1. Simplified overview of the serial two-photon data acquisition process.

A: $50\mu\text{m}$ of tissue (between $40\mu\text{m}$ to $90\mu\text{m}$ below the tissue surface) is imaged and then an in-built microtome physically removes a $50\mu\text{m}$ thick section from the surface. This process is repeated to generate a complete 3D dataset of the brain. B: The emitted light is split into two channels whereby the primary channel contains the fluorescence signal of interest from labelled cells (e.g. mCherry at $610\mu\text{m}$) and the second channel (e.g. at $450\mu\text{m}$) contains tissue autofluorescence signal that reveals gross anatomical structure.

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. (32)) 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 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. 2c). 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 (Fig. 2b-e).

Cell candidate classification using deep learning. A classification step, which uses a 3D adaptation of the ResNet (33) convolutional neural network (Figs. S1 & S2) 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. 2d). In total, $\sim 100,000$ cell candidates ($50,653$ cells and $56,902$ non-cells) were labelled from five brains. Small cuboids of $50 \times 50 \times 100\mu\text{m}$ 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. 2e). 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 $\sim 7,000$ annotated cell candidates.

Application. To illustrate the method, the cell detection software was applied to data which was not used to develop or train the classification network. Neurons presynaptic to inhibitory retrosplenial cortical cells were labeled using rabies virus tracing (expressing mCherry) in a GAD2-Cre mouse. On a desktop computer workstation (with 2×14 core Intel Xeon Gold 6132 CPUs and an NVIDIA TITAN RTX GPU) detection of $60,881$ cell candidates took 65 minutes and classification (resulting in $11,469$ cells) took 72 minutes (total of 2 hours and 17 minutes).

To assign detected cells to a brain region, the Allen Mouse Brain Reference Atlas (ARA (34)) annotations were registered to the secondary autofluorescence channel using brainreg (35), a Python port of the validated aMAP pipeline (36).




Fig. 2. Illustration of the cell detection process. A: Raw data (single coronal plane) B: Enlarged insert of cortical region from A showing structural features (artefacts) often erroneously detected. C: Detected cell candidates overlaid on raw data. Labelled cells as well as numerous artefacts are detected. D: Illustration of training data. A small number of cells (yellow) and artefacts (purple) selected to train the network. E: Classified cell candidates. The cell classification network correctly rejects the initial false positives.

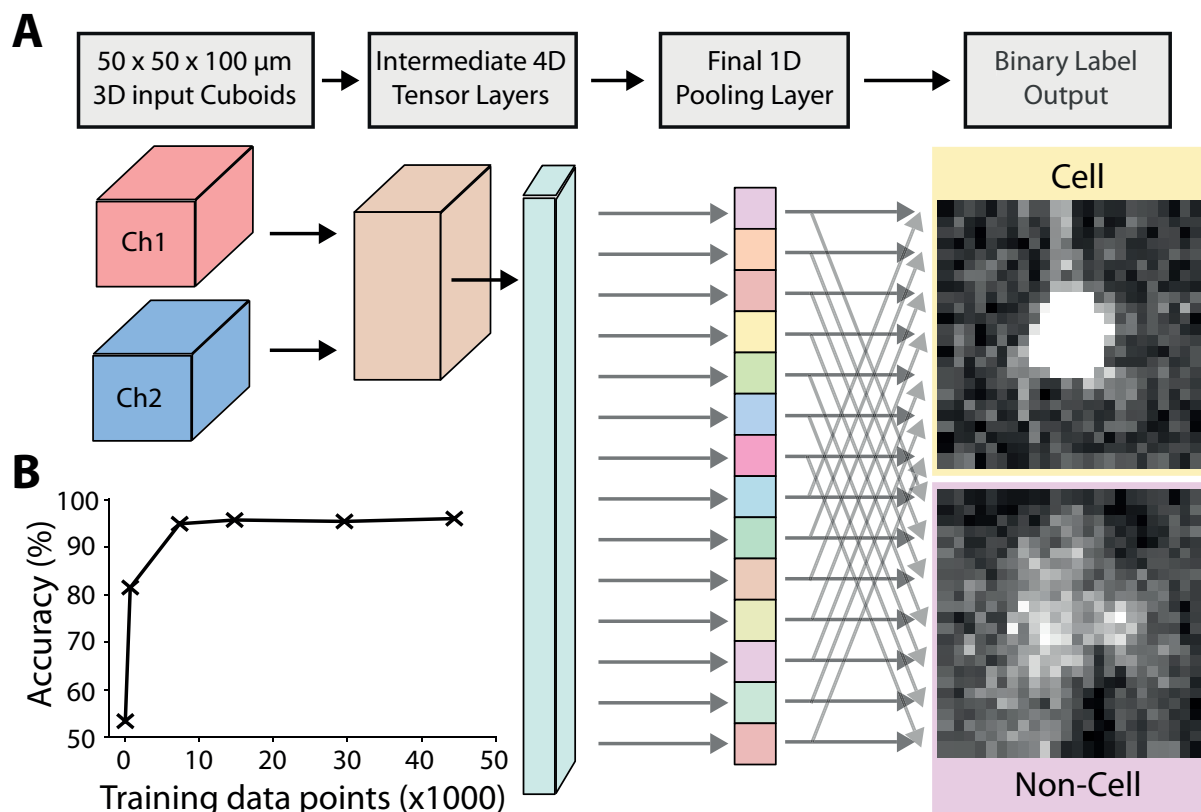


Fig. 3. Cell classification. A: The input data to the modified ResNet are 3D image cuboids ($50\ \mu\text{m} \times 50\ \mu\text{m} \times 100\ \mu\text{m}$) centered on each cell candidate. There are two cuboids, one from the primary signal channel, and one from the secondary autofluorescence channel. 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 Figs. S1 & S2 for more details of the 3D ResNet architecture. B: Classification accuracy as a function of training data quantity.

These annotations were overlaid on the raw data (Fig. 4a), and cell positions warped to the ARA coordinate space (Fig. 4b), to visualise the distribution of detected cells. The number of cells in each brain region were reported, allowing for quantitative analysis (Table 1).

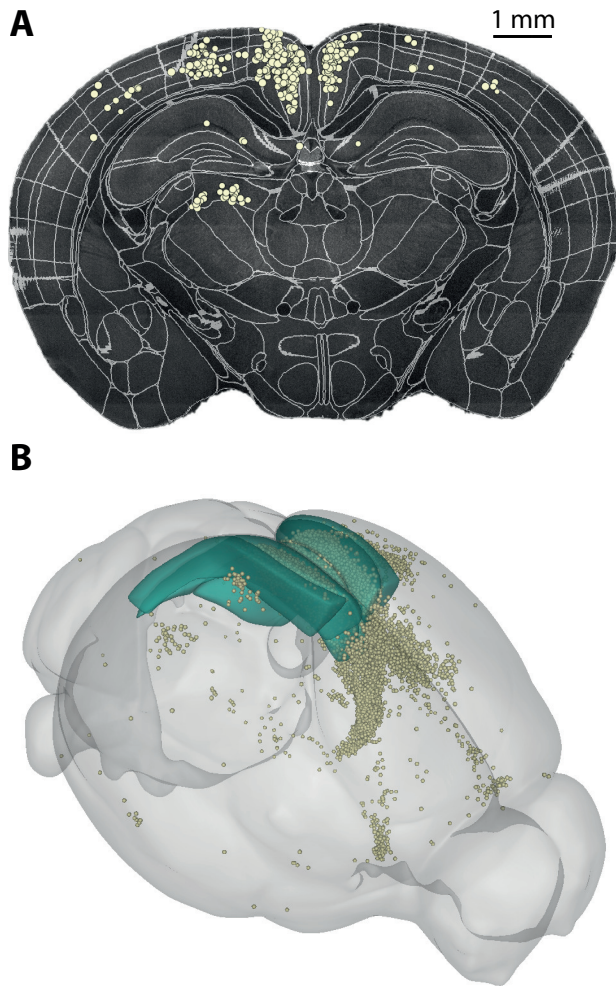


Fig. 4. Using atlas registration to visualise detected cells. A: Detected cells (yellow) overlaid on raw data, along with the brain region segmentation. B: Visualisation of detected cells warped to the ARA standard space in 3D (yellow), along with the injection site target (retrosplenial cortex, green).

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

Table 1. Number of cells in each brain region projecting to inhibitory GABA-ergic retrosplenial cortical neurons. Ten regions with the greatest number of cells shown (193 regions in total).

Brain structure	Left hemisphere count	Right hemisphere count	Total count
Retrosplenial area, ventral part, layer 5		1853	814 2667
Lateral dorsal nucleus of thalamus		1541	0 1541
Retrosplenial area, ventral part, layer 2/3		163	686 849
Retrosplenial area, dorsal part, layer 5		561	82 643
Retrosplenial area, dorsal part, layer 2/3		194	245 439
Ventral anterior-lateral complex of the thalamus		412	0 412
Anterior cingulate area, dorsal part, layer 5		340	40 380
Anteroventral nucleus of thalamus		374	0 374
Subiculum		304	0 304
Lateral posterior nucleus of the thalamus		275	0 275

be carried out in a timely fashion, and without relying on large-scale computing infrastructure. Processing time for the 240GB image in Fig. 4 on a desktop workstation was less than three hours, so eight datasets could be analysed in a single day, likely 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 (37), 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. Considering the classification performance is very good (95%, Fig. 3b) with ~7000 annotated cell candidates, sufficient training data can be annotated in a single day. The amount of training data required is also likely to be much lower if the new data is used to re-train the existing pre-trained network, rather than starting from scratch.

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 is also less biased than analysing a series of 2D planes, especially in regions with low cell densities, or differing cell sizes.

This software is fully open-source, and has been written with further development and collaboration in mind. In future we aim to adapt the network to be flexible as to the number of

input channels, and output labels. 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. 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.

DATA AVAILABILITY

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 and pre-built wheels at pypi.org/project/cellfinder. Documentation, including tutorials for training the network, and for analysing the data in Fig. 4 is available at docs.brainglobe.info/cellfinder.

ACKNOWLEDGEMENTS

This work was supported by grants from the Gatsby Charitable Foundation (GAT3361) and Wellcome Trust (090843/F/09/Z and 214333/Z/18/Z) to T.W.M.. This manuscript was typeset using a modified version of the [HenriquesLab bioRxiv template](https://github.com/HenriquesLab/bioRxiv-template).

Bibliography

1. Ian R. Wickersham, Stefan Finke, Karl Klaus Zenzelmann, and Edward M. Callaway. Retrograde neuronal tracing with a deletion-mutant rabies virus. *Nature Methods*, 4(1):47–49, 2007.
2. Leon G. Reijmers, Brian L. Perkins, Naoki Matsuo, and Mark Mayford. Localization of a stable neural correlate of associative memory. *Science*, 317(5842):1230–1233, 2007.
3. Sung Kim, Jae Hun Cho, Evan Murray, Naveed Bakh, Heejin Choi, Kimberly Ohn, Lizabeth Ruelas, Austin Hubbert, Meg McCue, Sara L. Vassallo, Philipp J. Keller, and Kwanghun Chung. Stochastic electrotransport selectively enhances the transport of highly electro-mobile molecules. *Proceedings of the National Academy of Sciences*, pages E6274–E6283, 2015.
4. Ali Ertürk, Klaus Becker, Nina Jähring, Christoph P. Mauch, Caroline D. Hojer, Jackson G. Egen, Farida Hellal, Frank Bradke, Morgan Sheng, and Hans-Ulrich Dodt. Three-dimensional imaging of solvent-cleared organs using 3DISCO. *Nature Protocols*, 7(11):1983–1995, 2012.
5. Kwanghun Chung, Jenelle Wallace, SungYon Kim, Sandhya Kalyanasundaram, Aaron S. Andalman, Thomas J. Davidson, Julie J. Mirzabekov, Kelly A. Zalocusky, Joannas Mattis, Aleksandra K. Denisin, Sally Pak, Hannah Bernstein, Charu Ramakrishnan, Logan Grosenick, Viviana Gradinaru, and Karl Deisseroth. Structural and molecular interrogation of intact biological systems. *Nature*, 497(7449):332–337, 2013.
6. Etsuo A. Susaki, Kazuki Tainaka, Dimitri Perrin, Fumiaki Kishino, Takehiro Tawara, Tomonobu M. Watanabe, Chihiro Yokoyama, Hirotsuka Onoe, Megumi Eguchi, Shun Yamaguchi, Takaya Abe, Hiroshi Kiyonari, Yoshihiro Shimizu, Atsushi Miyawaki, Hideo Yokota, and Hiroki R. Ueda. Whole-Brain Imaging with Single-Cell Resolution Using Chemical Cocktails and Computational Analysis. *Cell*, 157(3):726–739, 2014.
7. HU Dodt, Ulrich Leischner, and Anja Schierloh. Ultramicroscopy: three-dimensional visualization of neuronal networks in the whole mouse brain. *Nature Methods*, 4(4):331–336, 2007.
8. Timothy Ragan, Lolahon R. Kadiri, Kannan Umadevi Venkataraju, Karsten Bahlmann, Jason Sutin, Julian Taranda, Ignacio Arganda-Carreras, Yongsoo Kim, H. Sebastian Seung, and Pavel Osten. Serial two-photon tomography for automated ex vivo mouse brain imaging. *Nature Methods*, 9(3):255–258, 2012.
9. Pavel Osten and Troy W. Margrie. Mapping brain circuitry with a light microscope. *Nature Methods*, 10(6):515–523, 2013.
10. Qinyi Fu, Benjamin L. Martin, David Q. Matus, and Liang Gao. Imaging multicellular specimens with real-time optimized tiling light-sheet selective plane illumination microscopy. *Nature Communications*, 7:11088, 2016.
11. Kaoru Seiriki, Atsushi Kasai, Takeshi Hashimoto, Wiebke Schulze, Misaki Niu, Shun Yamaguchi, Takanobu Nakazawa, Ken ichi Inoue, Shiori Uezono, Masahiko Takada, Yuichiro Naka, Hisato Igarashi, Masato Arganda, James A. Waschek, Yukio Ago, Kenji F. Tanaka, Atsuko Hayata-Takano, Kazuki Nagayasu, Norihito Shintani, Ryota Hashimoto, Yasuto Kunii, Mizuki Hino, Junya Matsumoto, Hiroki Yabe, Takeharu Nagai, Katsumasa Fujita, Toshio Matsuda, Kazuhiro Takuma, Akemichi Baba, and Hitoshi Hashimoto. High-Speed and Scalable Whole-Brain Imaging in Rodents and Primates. *Neuron*, 94(6):1085–1100.e6, 2017.
12. Fabian F. Voigt, Daniel Kirschenbaum, Evgenia Platonova, Stéphane Pagès, Robert A. A. Campbell, Rahel Kastli, Martina Schaefflin, Ladan Egoif, Alexander van der Bourg, Philipp Bethge, Karen Haenraets, Noémie Frézel, Thomas Topilko, Paola Perin, Daniel Hillier, Sven Hildebrand, Anna Schueth, Alard Roebroek, Botond Roska, Esther T. Stoekli, Roberto Pizzala, Nicolas Renier, Hanns Ulrich Zellhofer, Theofanis Karayannis, Urs Ziegler, Laura Batti, Anthony Holtmaat, Christian Lüscher, Adriano Aguzzi, and Fritjof Helmchen. The mesoSPIM initiative: open-source light-sheet microscopes for imaging cleared tissue. *Nature Methods*, 16:1105–1106, 2019.
13. Johannes Schindelin, Ignacio Arganda-Carreras, Erwin Frise, Verena Kaynig, Mark Longair, Tobias Pletzsch, Stephan Preibisch, Curtis Rueden, Stephan Saalfeld, Benjamin Schmid, Jean-Yves Tinevez, Daniel James White, Volker Hartenstein, Kevin Eliceiri, Pavel Tomancak, and Albert Cardona. Fiji: an open-source platform for biological-image analysis. *Nature Methods*, 9(7):676–682, jul 2012. ISSN 1548-7105.
14. Fabrice De Chaumont, Stéphane Dallongeville, Nicolas Chenouard, Nicolas Hervé, Sorin Pop, Thomas Provoost, Yannary Meas-Yedid, Praveen Pankajakshan, Timothée Lecomte, Yoann Le Montagner, Thibault Lagache, Alexandre Dufour, and Jean Christophe Olivo-Marin. Icy: An open bioimage informatics platform for extended reproducible research. *Nature Methods*, 9(7):690–696, 2012.
15. Claire McQuin, Allen Goodman, Vasilij Chernyshev, Lee Kamentsky, Beth A. Cimini, Kyle W. Karhohs, Minh Doan, Liya Ding, Susanne N. Rafelski, Derek Thirstrup, Winfried Wieggers, Shantanu Singh, Tim Becker, Juan C. Caicedo, and Anne E. Carpenter. CellProfiler 3.0: Next-generation image processing for biology. *PLoS Biology*, 16(7):1–17, 2018.
16. Stuart Berg, Dominik Kutra, Thorben Kroeger, Christoph N. Straehle, Bernhard X. Kausler, Carsten Haubold, Martin Schiegg, Janez Ales, Thorsten Beier, Markus Rudy, Kemal Eren, Jaime I. Cervantes, Buote Xu, Fynn Beuttenmueller, Adrian Wolny, Chong Zhang, Ulrich Koethe, Fred A. Hamprecht, and Anna Kreshuk. Ilastik: Interactive Machine Learning for (Bio)Image Analysis. *Nature Methods*, 16:1226–1232, 2019.
17. Mitsuko Watabe-Uchida, Lisa Zhu, Sachie K. Ogawa, Archana Vamanrao, and Naoshige Uchida. Whole-Brain Mapping of Direct Inputs to Midbrain Dopamine Neurons. *Neuron*, 74(5):858–873, 2012.
18. Mateo Véléz-Fort, Charly V. Rousseau, Christian J. Niedworok, Ian R. Wickersham, Ede A. Rancz, Alexander P.Y. Brown, Molly Strom, and Troy W. Margrie. The stimulus selectivity and connectivity of layer six principal cells reveals cortical microcircuits underlying visual processing. *Neuron*, 83(6):1431–1443, 2014.
19. Sachie K. Ogawa, Jeremiah Y. Cohen, Dabin Hwang, Naoshige Uchida, and Mitsuko Watabe-Uchida. Organization of monosynaptic inputs to the serotonin and dopamine neuromodulatory systems. *Cell Reports*, 8(4):1105–1118, 2014.
20. Martin K. Schwarz, Annemarie Scherbarth, Rolf Sprengel, Johann Engelhardt, Patrick Theer, and Guenter Giese. Fluorescent-Protein Stabilization and High-Resolution Imaging of Cleared, Intact Mouse Brains. *PLoS one*, 10(5):e0124650, 2015.
21. Suzanaerculano-Houzel, Bruno Mota, and Roberto Lent. Cellular scaling rules for rodent brains. *Proceedings of the National Academy of Sciences of the United States of America*, 103(32):12138–12143, 2006.
22. Yongsoo Kim, Kannan Umadevi Venkataraju, Kith Pradhan, Carolin Mende, Julian Taranda, Srinivas C. Turaga, Ignacio Arganda-Carreras, Lydia Ng, Michael J. Hawrylycz, Kathleen S. Rockland, H. Sebastian Seung, and Pavel Osten. Mapping social behavior-induced brain activation at cellular resolution in the mouse. *Cell Reports*, 10(2):292–305, 2015.
23. William Menegas, Joseph F. Bergan, Sachie K. Ogawa, Yoh Isogai, Kannan Umadevi Venkataraju, Pavel Osten, Naoshige Uchida, and Mitsuko Watabe-Uchida. Dopamine neurons projecting to the posterior striatum form an anatomically distinct subclass. *eLife*, 4:1–30, 2015.
24. Casper Bo Gravesen Salinas, Tess Tsai Hsiu Lu, Sanaz Gabery, Kasper Marstal, Tomas Alanentalo, Aaron Jeffrey Mercer, Anda Cornea, Knut Conradsen, Jacob Hecksher-Sørensen, Anders Bjørholm Dahl, Lotte Bjerre Knudsen, and Anna Secher. Integrated Brain Atlas for Unbiased Mapping of Nervous System Effects Following Intraligature Treatment. *Scientific Reports*, 8(1):1–12, 2018.
25. Daniel Furth, Thomas Vaissiere, Ourania Tzortzi, Yang Xuan, Antje Martin, Iakovos Lazaridis, Giada Spigalò, Gilberto Fisone, Raju Tomer, Karl Deisseroth, Marie Carlen, Courtney A. Miller, Gavin Rumbaugh, and Konstantinos Meletis. An interactive framework for whole-brain maps at cellular resolution. *Nature Neuroscience*, 21:139–149, 2018.
26. Jin, Michelle and Nguyen, Joseph D and Weber, Sophia J and Mejias-Aponte, Carlos A and Madangopal, Rajtarun and Golden, Sam A. SMART: An open source extension of WholeBrain for iDISCO+ LSFM intact mouse brain registration and segmentation. *bioRxiv*, page 727529, 2019. doi: 10.1101/727529.
27. Asim Iqbal, Astandiyar Sheikh, and Theofanis Karayannis. DeNeRD: high-throughput detection of neurons for brain-wide analysis with deep learning. *Scientific Reports*, 9(1):1–13, 2019. ISSN 20452322.
28. Nicolas Renier, Eliza L. Adams, Christoph Kirst, Zhuhao Wu, Ricardo Azevedo, Johannes Kohl, Anita E. Autry, Lolahon Kadiri, Kannan Umadevi Venkataraju, Yu Zhou, Victoria X. Wang, Cheuk Y. Tang, Olav Olsen, Catherine Dulac, Pavel Osten, and Marc Tessier-Lavigne. Mapping of Brain Activity by Automated Volume Analysis of Immediate Early Genes. *Cell*, 165(7):1789–1802, 2016.
29. Yann Lecun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521:436 – 444, 2015.
30. Erick Moen, Dylan Bannon, Takamasa Kudo, William Graf, Markus Covert, and David Van Valen. Deep learning for cellular image analysis. *Nature Methods*, 2019.
31. Drew Friedman, Albert Pun, Eliza L. Adams, Jan H. Lui, Justus M. Kobschull, Sophie M. Grutzner, Caitlin Castagnola, Marc Tessier-Lavigne, and Liqun Luo. Mapping mesoscale axonal projections in the mouse brain using a 3d convolutional network. *Proceedings of the National Academy of Sciences*, 117(20):11068–11075, 2020. ISSN 0027-8424. doi: 10.1073/pnas.1918465117.
32. Nobuyuki Otsu. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979.
33. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2016-December, pages 770–778, 2016. ISBN 9781467388504. doi: 10.1109/CVPR.2016.90.
34. Quanxin Wang, Song Lin Ding, Yang Li, Josh Royall, David Feng, Phil Lesnar, Nile Graddis, Maitham Naeemi, Benjamin Facer, Anh Ho, Tim Dolbeare, Brandon Blanchard, Nick Dee, Wayne Wakeman, Karla E. Hirokawa, Aaron Szafer, Susan M. Sunkin, Seung Wook Oh, Amy Bernard, John W. Phillips, Michael Hawrylycz, Christof Koch, Hongkui Zeng, Julie A. Harris, and Lydia Ng. The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell*, 181(4):936–953.e20, 2020. ISSN 10974172.
35. Adam L. Tyson, Charly V. Rousseau, and Troy W. Margrie. brainreg: automated 3D brain registration with support for multiple species and atlases, August 2020.
36. Christian J. Niedworok, Alexander P.Y. Brown, M. Jorge Cardoso, Pavel Osten, Sebastian Urselinn, Marc Modat, and Troy W. Margrie. AMAP is a validated pipeline for registration and segmentation of high-resolution mouse brain data. *Nature Communications*, 7:1–9, 2016.
37. Helen Spiers, Harry Songhurst, Luke Nightingale, Joost de Folter, Roger Hutchings, Christopher J. Peddie, Anne Weston, Amy Strange, Steve Hindmarsh, Chris Lintott, Lucy M. Collinson, and Martin L. Jones. Citizen science, cells and cnns – deep learning for automatic segmentation of the nuclear envelope in electron microscopy data, trained with volunteer segmentations. *bioRxiv*, 2020. doi: 10.1101/2020.07.28.223024.
38. Stephan Preibisch, Stephan Saalfeld, and Pavel Tomancak. Globally optimal stitching of tiled 3D microscopic image acquisitions. *Bioinformatics*, 25(11):1463–1465, 2009. ISSN 13674803.
39. François Chollet et al. Keras. <https://keras.io>, 2015.
40. Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Good-

fellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

41. Nicholas Sofroniew, Talley Lambert, Kira Evans, Juan Nunez-Iglesias, Kevin Yamauchi, Ahmet Can Solak, Grzegorz Bokota, ziyangczi, Genevieve Buckley, Philip Winston, Tony Tung, Draga Doncila Pop, Hector, Jeremy Freeman, Matthias Bussonnier, Peter Boone, Loic Royer, Hagai Har-Gil, Shannon Axelrod, Ariel Rokem, Bryant, Justin Kiggins, Mars Huang, Pranathi Vemuri, Reece Dunham, Trevor Manz, jakirkham, Chris Wood, Alexandre de Siqueira, and Bhavya Chopra. napari/napari: 0.3.8, September 2020.
42. Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv*, 2014.
43. Federico Claudi, Luigi Petrucco, Adam L. Tyson, Tiago Branco, Troy W. Margrie, and Ruben Portugues. Brainglobe atlas api: a common interface for neuroanatomical atlases. *Journal of Open Source Software*, 5(54):2668, 2020. doi: 10.21105/joss.02668.
44. Marc Modat, Gerard R. Ridgway, Zeike A. Taylor, Manja Lehmann, Josephine Barnes, David J. Hawkes, Nick C. Fox, and Sébastien Ourselin. Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine*, 98(3): 278–284, 2010.
45. S. Ourselin, A. Roche, G. Subsol, X. Pennec, and N. Ayache. Reconstructing a 3D structure from serial histological sections. *Image and Vision Computing*, 19(1-2):25–31, 2001.
46. Federico Claudi, Adam L. Tyson, and Tiago Branco. Brainrender. a python based software for visualisation of neuroanatomical and morphological data. *bioRxiv*, 2020. doi: 10.1101/2020.02.23.961748.

Materials and methods

All experiments were carried out in accordance with the UK Home Office regulations (Animal Welfare Act 2006) and the local Animal Welfare and Ethical Review Body.

Sample preparation. All mice used were transgenic Cre-reporter (Ntsr1-Cre, GAD2-IRES-Cre & Rbp4-Cre) mice bred on a C57BL/6 background. The mice were anaesthetized and an AAV Cre-dependent helper virus encoding both the envelope protein receptor and the rabies virus glycoprotein 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. For imaging, the fixed brain was 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 (18). 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 subtraction of an average image from each physical section. Tiles were then stitched using a custom FIJI (13) plugin (modified from (38)) and downsampled to 2 μm x 2 μm x 5 μm voxel size.

Cell candidate detection. To detect cell candidates (broadly defined as anything of sufficient brightness and of approximately the correct size to be a cell), initially data from the primary signal channel was processed in 2D. 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 thresholded, and passed to 3D filters. An ellipsoidal filter was first used to remove noise. Every position of this spatial filter in which the majority of the filter overlaps with thresholded voxels was saved as a potential cell candidate. This is used to remove noise (from e.g. neurites). Individually detected cell candidates were then merged (if close together) or split (if the total cell volume is too large), and the coordinates were saved as an xml file. These steps were all carried out using the default software parameters (Table 2).

Table 2. Default cell detection parameters

Parameter	Value
X/Y cell size (\varnothing)	16 μm
Gaussian smoothing sigma	0.2 \varnothing
Intensity threshold	Mean + 10 x SD
X/Y ellipsoidal filter	6 μm
Z ellipsoidal filter	15 μm
Ellipsoidal filter overlap threshold fraction	0.6

Cell candidate classification using deep learning. Cell candidates were classified using a ResNet (33), implemented in Keras (39) for TensorFlow (40). 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 Figs. S1 & S2. To generate training data, output from the candidate detection step (cell candidate coordinates) were manually classified using a custom FIJI (13) plugin (for users of the released software, an integrated tool using napari (41) is supplied). 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:

- Flipping in any of the three axes
- Rotation around any of the three axes (between -45° to 45°)
- Circular translation along any of the three axes (up to 5% of the axis length)

The networks were trained with a batch size of 32 and the Adam (42) 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 (34), provided by the BrainGlobe Atlas API (43)) was registered to the autofluorescence channel. This was carried out using brainreg (35), a Python

port of the automatic mouse atlas propagation (aMAP) software (36), which itself relies on the medical image registration library, niftyreg (44). Firstly the sample brain was downsampled to the same voxel spacing as the atlas (10 μm isotropic) and was reoriented to 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* (45)), followed by freeform non-linear registration (*reg_f3d* (44)). 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.

Analysis. For numerical analysis the atlas annotations in the sample coordinate space were used to determine the number of cells in each brain region. 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.

Visualisation. All data visualisation was performed using napari (41) or brainrender (46).

Supplementary figures

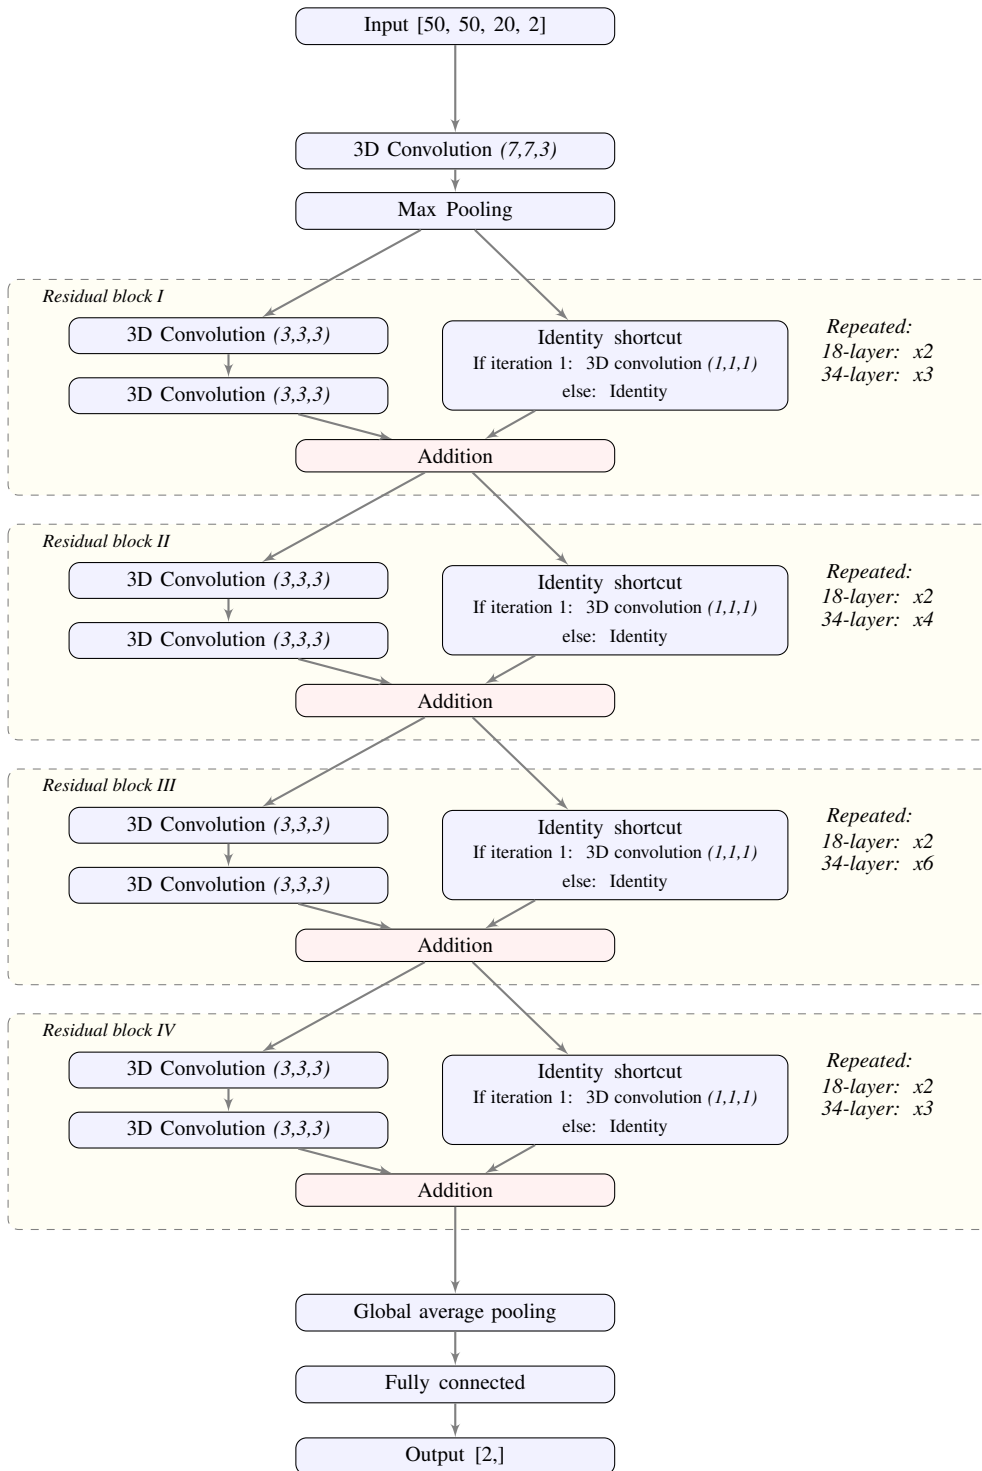


Fig. S1. Architecture of the 3D resnet

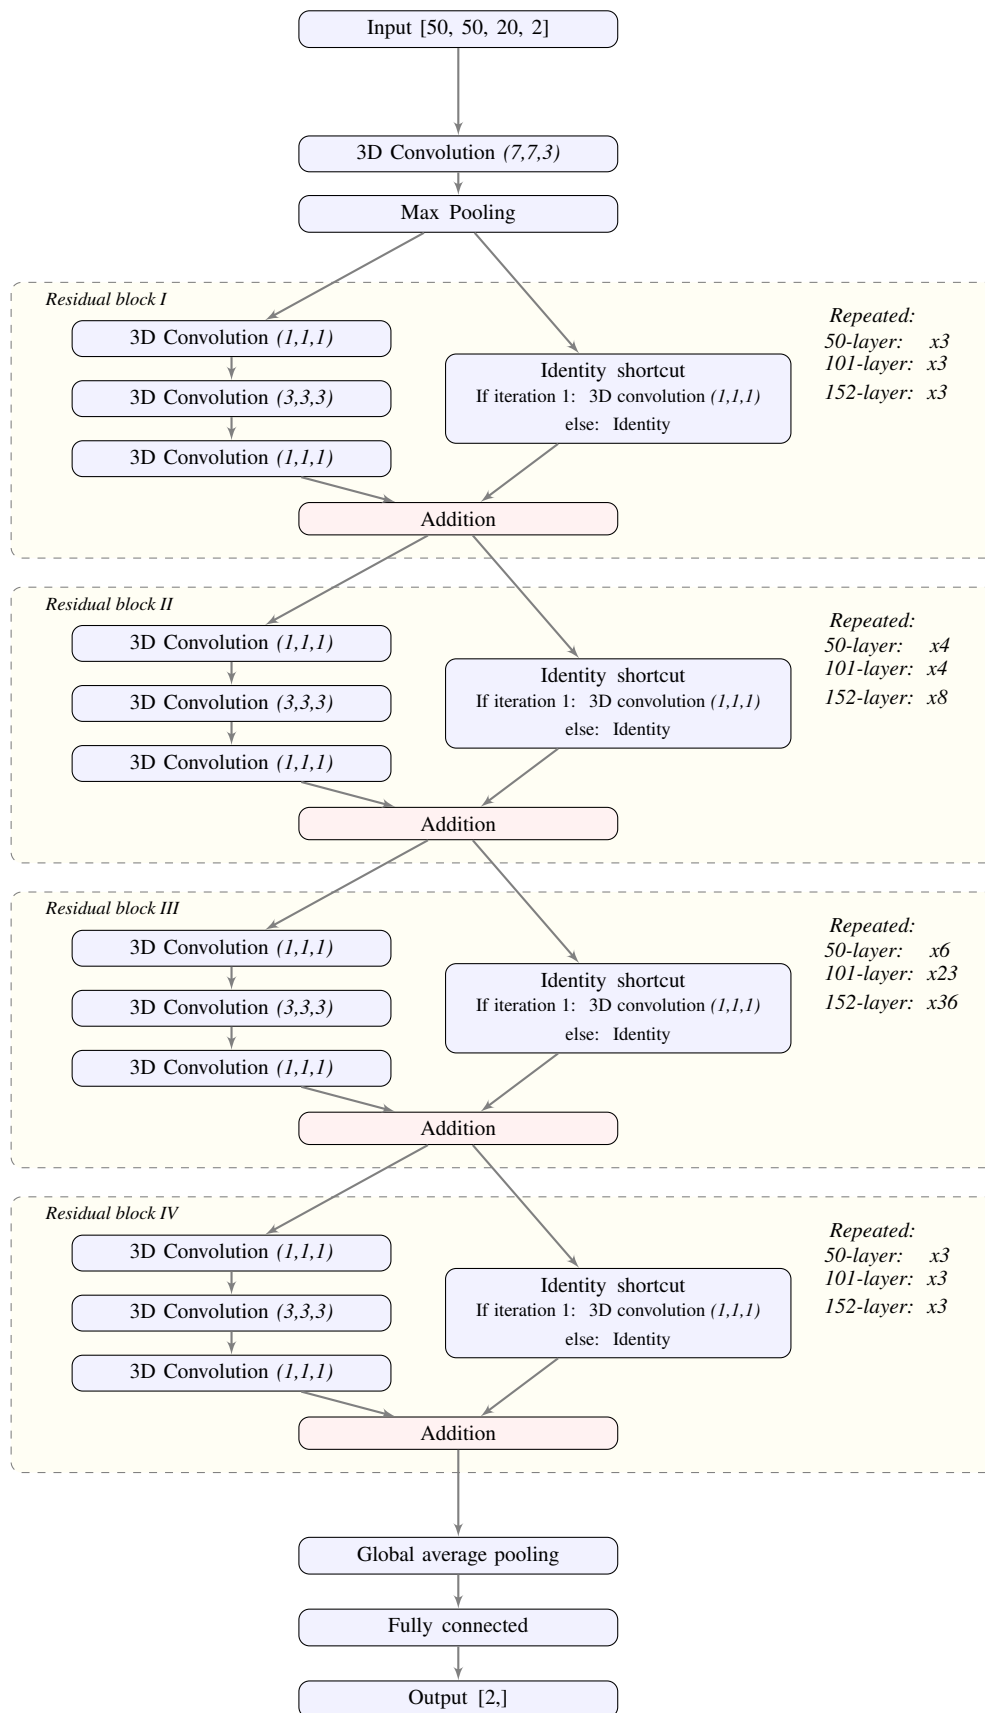


Fig. S2. Architecture of the bottleneck 3D resnet