



HAL
open science

Biom3d, a modular framework to host and develop 3D segmentation methods

Guillaume Mougeot, Sami Safarbati, Hervé Alégot, Pierre Pouchin, Nadine Field, Sébastien Almagro, Émilie Pery, Aline Probst, Christophe Tatout, David Evans, et al.

► **To cite this version:**

Guillaume Mougeot, Sami Safarbati, Hervé Alégot, Pierre Pouchin, Nadine Field, et al.. Biom3d, a modular framework to host and develop 3D segmentation methods. 2024. hal-04747016

HAL Id: hal-04747016

<https://hal.science/hal-04747016v1>

Preprint submitted on 21 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

1 **Biom3d, a modular framework to host and develop 3D**

2 **segmentation methods**

3 **Guillaume Mougeot (1), Sami Safarbaty (2), Hervé Alégot (2), Pierre Pouchin (2), Nadine**
4 **Field (4), Sébastien Almagro (5), Émilie Pery (3), Aline Probst (2), Christophe Tatout (2),**
5 **David E. Evans (4), Katja Graumann (4), Frédéric Chausse (3), Sophie Desset (2)**

6 (1) Aarhus University (Department of Ecoscience, Biodiversity and Conservation, C.F.
7 Møllers Allé 8, Building 1110, 8000 Aarhus C, Denmark)

8 (2) [iGReD - Génétique, Reproduction et Développement](#), [UCA - Université Clermont](#)
9 [Auvergne](#), [CNRS - Centre National de la Recherche Scientifique](#) UMR6293, [INSERM -](#)
10 [Institut National de la Santé et de la Recherche Médicale](#) U1103, (Facultés de Médecine et de
11 Pharmacie, TSA 50400, 28 Place Henri Dunant, 63001 Clermont-Ferrand - France)

12 (3) [IP - Institut Pascal](#), [UCA - Université Clermont Auvergne](#), [CNRS - Centre National de la](#)
13 [Recherche Scientifique](#) UMR6602 (Campus Universitaire des Cézeaux, 4 avenue Blaise
14 Pascal, TSA 60026 / CS 60026, 63178 Aubière Cedex - France)

15 (4) [Oxford Brookes University](#), (Department for Biological and Medical Sciences,
16 Headington Campus, Gypsy Lane, Oxford OX3 0BP, UK - Royaume-Uni)

17 (5) Université de Reims Champagne-Ardenne, CNRS, MEDYC, Reims, France

18

19 **U-Net is a convolutional neural network model developed in 2015 and has proven to be**
20 **one of the most inspiring deep-learning models for image segmentation. Numerous U-Net-**
21 **based applications have since emerged, constituting a heterogeneous set of tools that**
22 **illustrate the current reproducibility crisis in the deep-learning field. Here we propose a**
23 **solution in the form of Biom3d, a modular framework for deep learning facilitating the**
24 **integration and development of novel models, metrics, or training schemes for 3D image**
25 **segmentation. The new development philosophy of Biom3D provides an improved code**
26 **sustainability and reproducibility in line with the FAIR principles and is available as a**
27 **graphical user interface and an open-source deep-learning framework to target a large**
28 **community of users, from end users to deep learning developers.**

29

30 **Introduction**

31 The biomedical field produces many three-dimensional images (3D), whether to follow the fate
32 of a tumor in an organ or to measure shape properties of cells, nuclei or other organelles. Such
33 properties can be extracted using a method in image analysis called semantic segmentation.
34 Semantic segmentation classifies each pixel from an image as belonging to the object of interest
35 or to the background and thus produces a binary mask of the image. The increased massification
36 of data has led to the development of innovative solutions to automate segmentation. While
37 these solutions present some challenges, especially when it comes to 3D biomedical images,
38 they offer great potential. The variability in imaging modalities, noise-to-signal ratio, size,
39 voxel spacing or volume distribution of the objects to segment presents exciting opportunities
40 for further development. This is where deep learning methods come into play, being both
41 performant and flexible on a broad range of applications, from denoising¹ to image
42 segmentation², and of modalities, from microscopy to medical imaging, both for 2D and 3D
43 images.

44 U-Net, a convolutional neural network model developed in 2015³, has been demonstrated to be
45 one of the most effective deep-learning models for 3D semantic segmentation. Numerous U-
46 Net-based applications² have since emerged but are still not routinely used for biomedical image
47 analysis. This can be attributed, in part, to the difficulty encountered by non-programmers in
48 mastering such tools. Furthermore, there is a lack of evidence to demonstrate the superiority of
49 deep learning methods in comparison with conventional ones. Some methods, such as
50 CellPose^{4,5}, StarDist⁶ and DeepImageJ⁷, are ready to use, with the necessary code,
51 documentation, and information about the computing environment². Unfortunately, they have
52 lost some of their flexibility, and adaptation to other datasets. Their possibility of new training,
53 especially for 3D images, is not always straightforward or even possible. On the other hand,
54 nnU-Net⁸, with its ability to auto-configure itself on training sets, performs well for many image
55 acquisition modalities, but is not accessible to users without programming experience. At the
56 other end of the skills spectrum, deep learning developers have at their disposal a bank of
57 perfectly documented and up-to-date functions thanks to MONAI⁹, but without an assembly
58 plan, such as proposed for 2D images by OpenMMLab¹⁰, development of any innovation is a
59 time-consuming process, and there is currently no available method for evaluating them in a
60 timely manner. The objective of this article is to present the development of a deep learning
61 segmentation framework that combines the ease of use of DeepImageJ with the flexibility of

62 nnU-Net, while also serving as a development base for improving methods for developers in
63 the style of OpenMMLab.

64 In this work, in addition to performance, flexibility and reusability, we propose a new prism
65 through which to evaluate a deep learning method: **sustainability**. Our definition of method
66 sustainability encapsulates two main objectives: method accessibility and adaptability. Method
67 accessibility relates to the range of users concerned with the tool. Increasing this range would
68 not only result in an enlarged user community, but would also facilitate the emergence of
69 multidisciplinary interactions, a goal that is shared by numerous researchers¹¹. This is
70 particularly true for software inherited from fundamental research, which attempts to reach
71 applied research. Involving end-users will ensure that the software meets the needs of an
72 external community and has concrete applications, giving it both meaning and feedback for
73 improvements. Involving external developers will attract novel ideas while strengthening a
74 community of maintainers. Method adaptability relates to this last community of developers
75 and is the ability of a tool to quickly evolve, especially in the fast-moving world of deep
76 learning. It can be reasonably assumed that the more adaptable a tool is, and the broader its user
77 community, the longer it should last.

78 Following this, we defined a new code development philosophy. To reach a wider audience, our
79 aim is to create a series of tools that bridge the gap between end-users and developers in a gentle
80 way, meaning that end-users wishing to deepen their understanding of the method right down
81 to the lines of code need to be able to do this in small steps only. For each user profile, the
82 objective is thus to substantially reduce the time required to invest in using a deep learning
83 method, while simultaneously expanding the scope of accessible functions (Figure 1). The
84 intention is to prioritize usability to the same extent as functionality. Current state-of-the-art
85 methods usually target only one of these user profiles: nnU-Net⁸, Cellpose⁵ or StarDist⁶ target
86 Python Programmers, while ZeroCostDL4Mic¹² or DeepImageJ⁷ are mainly addressed to Non-
87 Programmers (Figure 1). Moreover, during the code implementation, the emphasis should be
88 made on code clarity, which implies that the code should remain comprehensible from its
89 overall architecture to its individual lines by, for example, being presented in a hierarchical,
90 tree-like structure. This approach starts with high-level, general functions and progressively
91 delves into more granular details and functionalities, facilitating a gradual and systematic
92 comprehension of the code.

93 Such a method can also be made adaptive to novel innovations by incorporating an additional
94 objective into the development process: code modularity. Code modularity means that

95 components should be as self-contained and independent as possible (Figure 1). In computer
96 science, this concept is also known as code cohesion and is opposed to code coupling¹³. This
97 approach allows programmers to seamlessly incorporate or remove specific code components
98 without causing undue disruptions to the overall system.

99 Following all these sustainability constraints, we developed a series of open-access tools called
100 Biom3d. By default, it is optimized to segment 3D images potentially having multiple channels
101 or simultaneously presenting objects from different classes. These volumetric images can
102 originate from CT-scan, MRI, confocal, X-ray or electron microscope. Biom3d integrates
103 offline and online interfaces (GUI, CLI, and API), automated configuration processes, editable
104 configuration files and a modular deep learning framework facilitating the choice of models,
105 metrics, data loaders or training routines.

106

107 **Results**

108 **Biom3D flexibility and performance.** Biom3d was first benchmarked with medical datasets
109 using part of the Medical Segmentation Decathlon (MSD)¹⁴ and the Multi-organ Abdominal
110 challenge (often called BTCV)¹⁵. These datasets have served to develop the flexibility of nnU-
111 Net and constitutes thus the primary goal of Biom3d performance (Figure 2). On both datasets,
112 Biom3d automatically adjusts its training configuration, resulting in Dice scores on the test sets
113 exceeding those of nnU-Net (Figure 3). To demonstrate its performance against conventional
114 methods, Biom3d was compared with NucleusJ¹⁶, a method specialized in 3D segmentation of
115 individual nuclei. A custom dataset of 93 images of individual plant nuclei of various shapes
116 and fluorescence intensity captured with a structured illumination microscope was created¹⁷. In
117 this case, in addition to nnU-Net and NucleusJ, Biom3d could also be compared to methods
118 designed for 3D nucleus segmentation in fluorescent images like DeepCell¹⁸, QCANet¹⁹,
119 Cellpose⁵ and StarDist⁶ (Figure 3). While being slightly better than nnU-Net, Biom3d
120 significantly outperforms the other methods, even the conventional method NucleusJ. Our
121 hypothesis is that the pretraining of Cellpose, DeepCell and QCANet and the star-convex
122 polyhedron constraint of StarDist introduce strong biases and prevent the model from
123 appropriately fitting to the strong variation of intensity within the nuclei. To further assess its
124 capabilities, Biom3d was also evaluated with other modalities (Figure 2). For electron
125 microscope (EM) images, Biom3d achieved a Dice score of 92% when evaluated with the two
126 3D images from EPFL public database²⁰ (Figure 3). Finally, Biom3d was applied to an EM
127 image of plant root nuclei, to synchrotron images of aortic lamellae and to segment nonconvex
128 objects such as membranes of epithelial cells. In all these cases, Biom3d successfully
129 segmented the objects of interest (Figure 2).

130 **A modular architecture.** During Biom3d development, a significant amount of effort was
131 dedicated to ensuring code modularity. It would be beneficial for programmers to have the
132 opportunity to experiment with a broad degree of freedom in the deep learning hyper-
133 parameters by altering the type of deep learning model, the learning rate, or the data loading
134 process. These challenges are tackled with the modular structure of Biom3d, which first core
135 components are Configuration Files. The Configuration Files of Biom3d depart from those of
136 OpenMMLab being both simplified and completed. They include the definition of two types of
137 hyper-parameters. First, stand-alone hyper-parameters can be integers, float numbers, or lists.
138 They include parameters defined in the Graphical User Interface (patch size, number of epochs
139 etc.) as well as many additional stand-alone hyper-parameters such as the initial learning rate

140 or whether to use half-precision float-point format or not. Second, module hyper-parameters
141 are key-value dictionaries. In this dictionary, the first key-value pair precises the name of the
142 module being used, and the second key-value pair defines the parameters of this module. For
143 example, “UNet3DVGGDeep” is the name of the 3D U-Net model in Biom3d, and its set of
144 parameters is the number of pooling layers of the model and the number of classes of objects
145 in the images (Figure 4).

146 The second modularity component of Biom3d is the Module Register. Modules are code
147 components classified in eight categories (Figure 4): Preprocessors, preparing the datasets
148 before training, Datasets, loading the datasets during training, Models, defining the deep
149 learning models, Metrics, defining the loss functions for model performance assessment,
150 Trainers, defining the training and validation loops, Predictors, defining the prediction loops,
151 Postprocessors, processing the model output before saving and Callbacks, decoupling all
152 recurrent events such as model saving, learning rate updates, logs printing and saving, or loss
153 updates from the training loop (Figure 4). We expect that such a paradigm will simulate a
154 significant number of deep learning methods, with a broader application spectrum than just
155 image segmentation. Except for Callbacks which are defined in the Builder (see below), all
156 available Modules in Biom3d are named and listed in the Module Register. While editing the
157 Configuration File, a user can thus select from the Register the modules that fit the most to the
158 expected task. As modules are Python classes or functions, if the user does not find the
159 appropriate modules, new modules can be easily integrated into Biom3d by importing them and
160 adding a new entry to the Register. It must be noticed that Biom3d accepts Pytorch modules
161 from any origin if they respect the expected inputs and outputs of the surrounding modules. For
162 example, if intending to add a new Metric to the default Biom3d configuration, the user must
163 only make sure that it includes a few arguments like a name, a value and an average value. Once
164 added to the Register, the new Module can directly be incorporated in Biom3d workflow by
165 simply editing the Configuration File.

166 If Modules are the organs of Biom3d, the Builder is its skeleton, providing a general structure
167 for the Module to properly work together. The Builder reads the Configuration File, retrieves
168 and instantiates the selected modules from the Register, and prepares them for training or
169 prediction. Once a Builder is instantiated with a Configuration File, the user can then easily
170 start the training or prediction process by calling a single function. The Builder stores its
171 configuration and the training status (training curves, etc.) which can then be reused to carry on
172 interrupted training or to fine-tune models with a new Configuration File.

173 **An easy-to-use tool.** Biom3d integrates one local Graphical User Interface (GUI) and one
174 online GUI in the form of a Google Colab notebook. An important innovation with Biom3d
175 interfaces is the “auto-configuration” button. Once the dataset folders path has been entered and
176 this button pressed, Biom3d proceeds to a cascade of hidden operations (Supplementary Figure
177 1). The dataset is first scanned to check its quality and to retrieve its characteristics, such as the
178 median image size or the median sampling. Biom3d is flexible on the input image format, as it
179 works with most medical formats (Nifti, DICOM etc.) and works with TIFF format, largely
180 used for microscopy data. Biom3d data scanning process is equipped with many safeguards to
181 automatically correct most user annotator mistakes, which seems to be decisive in avoiding the
182 "first-click abandonment" phenomenon. These data characteristics are then used to normalize
183 the images and to automatically configure the training process of the deep learning model.

184 The result of this automated training configuration, including the definition of the patch size or
185 the batch size, is stored in a YAML file, is partially displayed in the interface, and is left editable
186 in its entirety if needed. Finally, pressing on “start” initiates the training process. The local
187 interface can also be employed in conjunction with a remote server, where computing resources
188 are installed, and with OMERO²¹, a hosting platform for microscopy images. While the training
189 proceeds, all relevant information, such as the training curves or the trained model, is
190 periodically stored. Once trained, a model can be either fine-tuned on another dataset, or used
191 for prediction. The prediction process is straightforward for the end user (Supplementary Figure
192 1): the training configuration file is used to load the pretrained model, to normalize the new
193 dataset, and to execute the prediction workflow.

194 **A toolbox for bioimage analysts.** Biom3d has been packaged as an easy-to-use Python library
195 that can serve as an Application Programming Interface (API) and integrates a Command Line
196 Interface (CLI). Exploiting the CLI enlarges Biom3d potentialities. Users may for instance
197 execute Biom3d on a High-Performance Computing (HPC) resource in a batch process while
198 applying it on a large variety of datasets. They could also isolate individual processes such as
199 data preprocessing, auto-configuration, or model training. Moreover, to fuse capabilities of
200 multiple models trained on different datasets or with different configurations, Biom3d allows
201 multi-model predictions. This last capability is particularly important to assemble the work of
202 different teams across the world, working under different conditions. All these CLI options are
203 directly accessible after installation of Biom3d Package, referenced in the Python Package
204 Index (PyPI).

205 Python programmers may also be interested in exploiting the potential of the Configuration File
206 and the eight types of modules. Configuration Files of Biom3d allows fine-tuning of more
207 hyper-parameter settings more precisely than is feasible through the GUI alone, without the
208 need to delve into Biom3d code. Biom3d modularity allows for rapid experimentation with
209 different configurations to identify the optimal setup for specific tasks or datasets. For example,
210 the default deep learning model of Biom3d, which corresponds to the dynamic 3D U-Net
211 available in nnU-Net framework, can easily be changed for another one. As a result, we
212 demonstrated that a 3D EfficientU-Net model can outperform the 3D U-Net on the Pancreas
213 dataset (Figure 3). Additionally, Biom3d is compatible with MONAI⁹, public library integrating
214 the latest biomedical models for multidimensional images, which enable quick use cutting-edge
215 models.

216 **A framework for developers.** Biom3d has also been meticulously crafted at a granular code
217 level, allowing deep learning programmers, well-versed in working with libraries like PyTorch
218 or TensorFlow, to easily understand, edit, remove, or add code components. It is also worth
219 noting that Biom3d plays an important role in the process of comparing old and new code
220 segments when creating new code based on existing works. This development strategy was, for
221 instance, exploited to entirely redesign the original Dataset Module of nnU-Net. As Biom3d
222 easily accepts alien code incorporation, the entire and original “Dataset Module” of nnU-Net
223 was integrated as a Biom3d Module. Following this integration, a series of intermediate and
224 hybrid Dataset Modules were developed and tested. Comparisons between old and new Dataset
225 Modules could easily be done by editing a single line in the Configuration File. The weak points
226 of the prototype Dataset Module were rapidly spotted and improved. The final Dataset Module,
227 which relies on the TorchIO library²², is likely the reason why Biom3d outperforms nnU-Net
228 on multiple tasks (Figure 3). This development strategy can be followed for all Biom3d
229 Modules. Last, as each piece of Biom3d code can be isolated from the others, they can easily
230 be extracted and reused for another independent project. This is particularly true not only for
231 Biom3d Modules but even for individual functions such as the auto-configuration or the image
232 loading and saving.

233 **Discussion**

234 Biom3d is a new, high-performing, and flexible set of tools applicable to various 3D
235 segmentation problems, with numerous biological and medical applications. It surpasses in
236 accuracy specialized tools, such as NucleusJ, while demonstrating state-of-the-art results on a

237 wide variety of new problems, such as plant chromocenter segmentation or aorta lamella
238 segmentation. It responds to the needs of end users through easy to install and easy to use
239 interfaces, enabling auto-configuration of the deep learning model training, and compatible
240 with several image formats and with OMERO software²¹. To make Biom3d an even bigger part
241 of the landscape of tools now available to biologist users, it will be made compatible with
242 Napari²³, ZeroCostDL4Mic^{12,24} and DeepImageJ⁷. In line with FAIR principles for AI²⁵, future
243 models will then be shared online in Open Neural Network Exchange (ONNX)²⁶ format, on
244 website such as HuggingFace.co or Bioimage.io, and future developments streamlined on
245 MLflow (<https://mlflow.org>). To broaden Biom3d's spectrum of applications, further
246 developments should include instance segmentation, object tracking, image denoising as well
247 as processing images in proprietary formats or with large N-dimension such as light sheet
248 images and to offer the possibility of using images stored with the next generation file formats
249 (NGFF) such as OME-Zarr²⁸.

250 The implementation of all these developments will be facilitated by the modular code
251 architecture of Biom3d. Biom3d has been designed as a sandbox, within which developers are
252 strongly encouraged to integrate new code elements. It is coded in the Python programming
253 language and based primarily on the Pytorch²⁹ deep learning framework. This language and
254 framework were chosen for their current popularity in the field of deep learning, so code
255 implemented with them, such as MONAI, can be added to Biom3d with little to no effort. These
256 constraints may yet pose some integration challenges for developers using other programming
257 languages or deep learning libraries, such as TensorFlow³⁰, with which StarDist was created, or
258 JAX³¹, an emerging framework. While being highly modular, Biom3d code remains coupled to
259 some extent. This occurs because it is needed in some instances to improve efficiency, by using
260 shared internal functions across different components. It is minimized however to keep the
261 advantages of modularity as interconnection can make the codebase more sensitive to changes
262 as alterations to one section of the code might unintentionally impact other areas, complicating
263 maintenance, and updates. The development of Biom3d represents a significant advance in the
264 modularity of deep learning code, while maintaining performance and avoiding code
265 redundancy. While further increases in modularity are possible, they may necessitate a
266 significant additional investment of time and resources.

267 Finally, through the benchmarking performed on nucleus images in this article, we have
268 pinpointed one of the main remaining limitations of deep learning methods: they depend on
269 large, manually annotated datasets. While conventional image analysis tools can help provide

270 partial annotations, manual interventions are inevitable, which also impedes the spread of deep
271 learning methods. The modular code of Biom3d could be exploited to reduce this time-
272 consuming process, by including methods such as active learning³², generative methods³³,
273 weakly supervised learning³⁴, and self-supervised learning^{35,36}. Biom3d has, for instance,
274 successfully performed self-supervision for 3D images, which involved both image
275 classification and image segmentation.

276 Most importantly, we hope that the intuitive design of Biom3d, from its graphical interfaces to
277 its modular code architecture, will foster collaboration among diverse communities, including
278 biologists, radiologists, microscopists, image analysts and developers, and will render it both
279 reusable and sustainable.

280 **Online content**

281 Biom3d is available at <https://github.com/GuillaumeMougeot/biom3d> and as a notebook at
282 https://colab.research.google.com/github/GuillaumeMougeot/biom3d/blob/master/docs/biom3d_colab.ipynb.
283

284 **References**

- 285 1. Buchholz, T. O. *et al.* Content-aware image restoration for electron microscopy.
286 *Methods Cell Biol.* **152**, 277–289 (2019).
- 287 2. Mougeot, G. *et al.* Deep learning — promises for 3D nuclear imaging: a guide for
288 biologists. *J. Cell Sci.* **135**, jcs258986 (2022).
- 289 3. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical
290 image segmentation. *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif.*
291 *Intell. Lect. Notes Bioinformatics)* **9351**, 234–241 (2015).
- 292 4. Stringer, C., Wang, T., Michaelos, M. & Pachitariu, M. Cellpose: a generalist
293 algorithm for cellular segmentation. *Nat. Methods* **18**, 100–106 (2021).
- 294 5. Pachitariu, M. & Stringer, C. Cellpose 2.0: how to train your own model. *Nat. Methods*
295 **19**, 1634–1641 (2022).
- 296 6. Weigert, M., Schmidt, U., Haase, R., Sugawara, K. & Myers, G. Star-convex polyhedra
297 for 3D object detection and segmentation in microscopy. *Proc. - 2020 IEEE Winter*
298 *Conf. Appl. Comput. Vision, WACV 2020* 3655–3662 (2020).

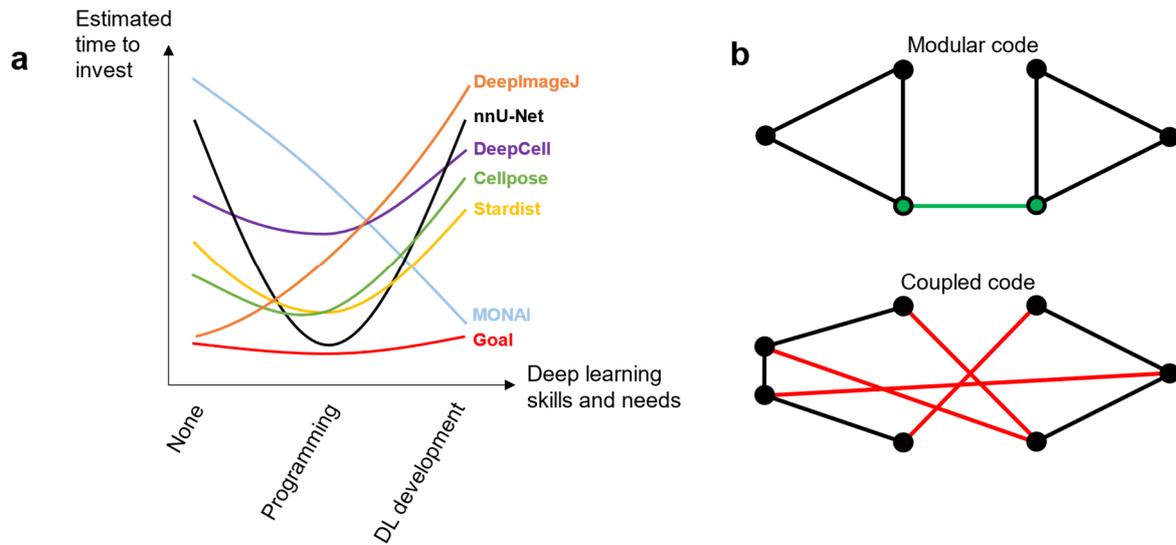
- 299 doi:10.1109/WACV45572.2020.9093435
- 300 7. Gómez-de-Mariscal, E. *et al.* DeepImageJ: A user-friendly environment to run deep
301 learning models in ImageJ. *Nat. Methods* **18**, 1192–1195 (2021).
- 302 8. Isensee, F., Jaeger, P. F., Kohl, S. A. A., Petersen, J. & Maier-Hein, K. H. nnU-Net: a
303 self-configuring method for deep learning-based biomedical image segmentation. *Nat.*
304 *Methods* **18**, 203–211 (2021).
- 305 9. MONAI Consortium. MONAI: Medical Open Network for AI. (2023).
306 doi:<https://doi.org/10.5281/zenodo.4323058>
- 307 10. MMSegmentation. Available at: <https://github.com/open-mmlab/mms Segmentation>.
- 308 11. Nogare, D. D., Hartley, M., Deschamps, J., Ellenberg, J. & Jug, F. Using AI in
309 bioimage analysis to elevate the rate of scientific discovery as a community. *Nat.*
310 *Methods* **20**, 973–975 (2023).
- 311 12. Chamier, L. von *et al.* ZeroCostDL4Mic: An open platform to use deep-learning in
312 microscopy. *bioRxiv* 2020.03.20.000133 (2020). doi:10.1101/2020.03.20.000133
- 313 13. Gamma, E., Helm, R., Johnson, R. & Vlissides, J. *Design Patterns: Elements of*
314 *Reusable Software. Addison-Wesley Professional Computing Series* (Addison-Wesley
315 Longman Publishing Co., Inc., 1996).
- 316 14. Antonelli, M. *et al.* The Medical Segmentation Decathlon. *Nat. Commun.* **13**, 4128
317 (2022).
- 318 15. Gibson, E. *et al.* Multi-organ Abdominal CT Reference Standard Segmentations.
319 *Zenodo, 22-Feb-2018* **26**, 1–7 (2018).
- 320 16. Dubos, T. *et al.* Automated 3D bio-imaging analysis of nuclear organization by
321 NucleusJ 2.0. *Nucleus* **11**, 315–329 (2020).
- 322 17. iGReD, U. C. A. Dataset of individual plant nuclei. Available at:
323 <https://omero.bio.fsu.edu/webclient/?show=project-5001>.
- 324 18. Greenwald, N. F. *et al.* Whole-cell segmentation of tissue images with human-level
325 performance using large-scale data annotation and deep learning. *Nat. Biotechnol.*
326 2021.03.01.431313 (2021). doi:10.1038/s41587-021-01094-0
- 327 19. Tokuoka, Y. *et al.* 3D convolutional neural networks-based segmentation to acquire

- 328 quantitative criteria of the nucleus during mouse embryogenesis. *npj Syst. Biol. Appl.* **6**,
329 1–12 (2020).
- 330 20. Lucchi, A., Li, Y. & Fua, P. Learning for structured prediction using approximate
331 subgradient descent with working sets. *Proc. IEEE Comput. Soc. Conf. Comput. Vis.*
332 *Pattern Recognit.* 1987–1994 (2013). doi:10.1109/CVPR.2013.259
- 333 21. Allan, C. *et al.* OMERO: Flexible, model-driven data management for experimental
334 biology. *Nat. Methods* **9**, 245–253 (2012).
- 335 22. Pérez-García, F., Sparks, R. & Ourselin, S. TorchIO: A Python library for efficient
336 loading, preprocessing, augmentation and patch-based sampling of medical images in
337 deep learning. *Comput. Methods Programs Biomed.* **208**, 106236 (2021).
- 338 23. Sofroniew, N. *et al.* napari. (2022). doi:10.5281/zenodo.5848842
- 339 24. Hidalgo-Cenalmor, I. *et al.* DL4MicEverywhere: Deep learning for microscopy made
340 flexible, shareable, and reproducible. *bioRxiv* (2023). doi:10.1101/2023.11.19.567606
- 341 25. Huerta, E. A. *et al.* FAIR for AI: An interdisciplinary and international community
342 building perspective. *Sci. Data* **10**, 487 (2023).
- 343 26. Bai, J., Lu, F. & Zhang, K. ONNX: Open Neural Network Exchange. *GitHub*
344 *repository* (2019).
- 345 27. MLflow: An Open Source Platform for Machine Learning. (2024).
- 346 28. Miles, A. *et al.* zarr-developers/zarr-python: v2.17.1. (2024).
347 doi:10.5281/zenodo.10790679
- 348 29. Paszke, A. *et al.* PyTorch: An imperative style, high-performance deep learning library.
349 *arXiv* **32**, (2019).
- 350 30. Abadi, M. *et al.* TensorFlow: Large-Scale Machine Learning on Heterogeneous
351 Distributed Systems. (2016).
- 352 31. Bradbury, J., Frostig, R., Hawkins, P., Johnson, M., Leary, C., Maclaurin, D., Necula,
353 G., Paszke, A., Vanderplas, J., WandermanMilne, S., Zhang, Q. JAX: composable
354 transformations of Python + NumPy programs. 1–8 (2018).
- 355 32. Budd, S., Robinson, E. C. & Kainz, B. A survey on active learning and human-in-the-
356 loop deep learning for medical image analysis. *Medical Image Analysis* **71**, 102062

- 357 (2021).
- 358 33. Fu, C. *et al.* Three dimensional fluorescence microscopy image synthesis and
359 segmentation. *IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Work.* **2018–**
360 **June**, 2302–2310 (2018).
- 361 34. Zhao, Z. *et al.* Deep Learning Based Instance Segmentation in 3D Biomedical Images
362 Using Weak Annotation BT - Medical Image Computing and Computer Assisted
363 Intervention – MICCAI 2018. in (eds. Frangi, A. F., Schnabel, J. A., Davatzikos, C.,
364 Alberola-López, C. & Fichtinger, G.) 352–360 (Springer International Publishing,
365 2018).
- 366 35. Sahasrabudhe, M. *et al.* Self-supervised Nuclei Segmentation in Histopathological
367 Images Using Attention BT - Medical Image Computing and Computer Assisted
368 Intervention – MICCAI 2020. in (eds. Martel, A. L. et al.) 393–402 (Springer
369 International Publishing, 2020).
- 370 36. Kirillov, A. *et al.* Segment Anything. (2023).

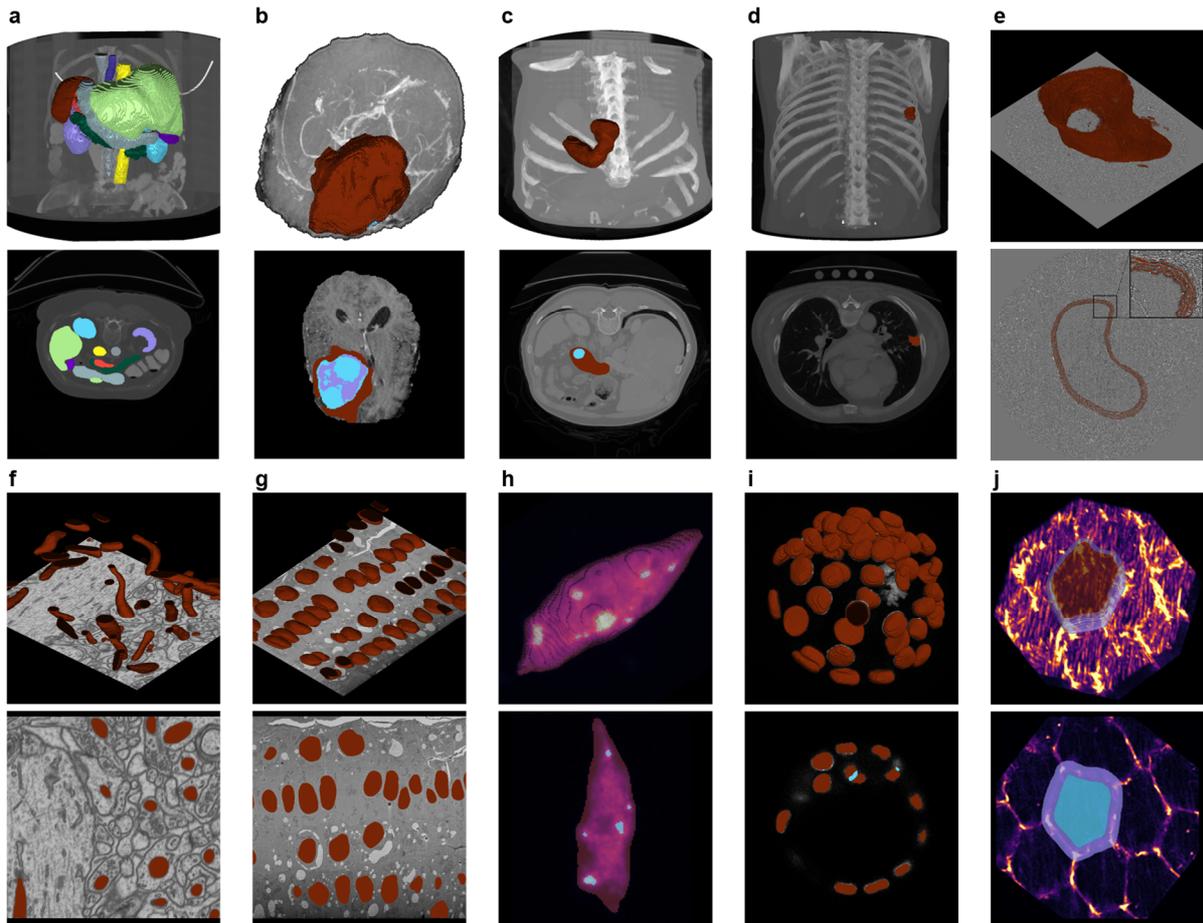
371

372



373

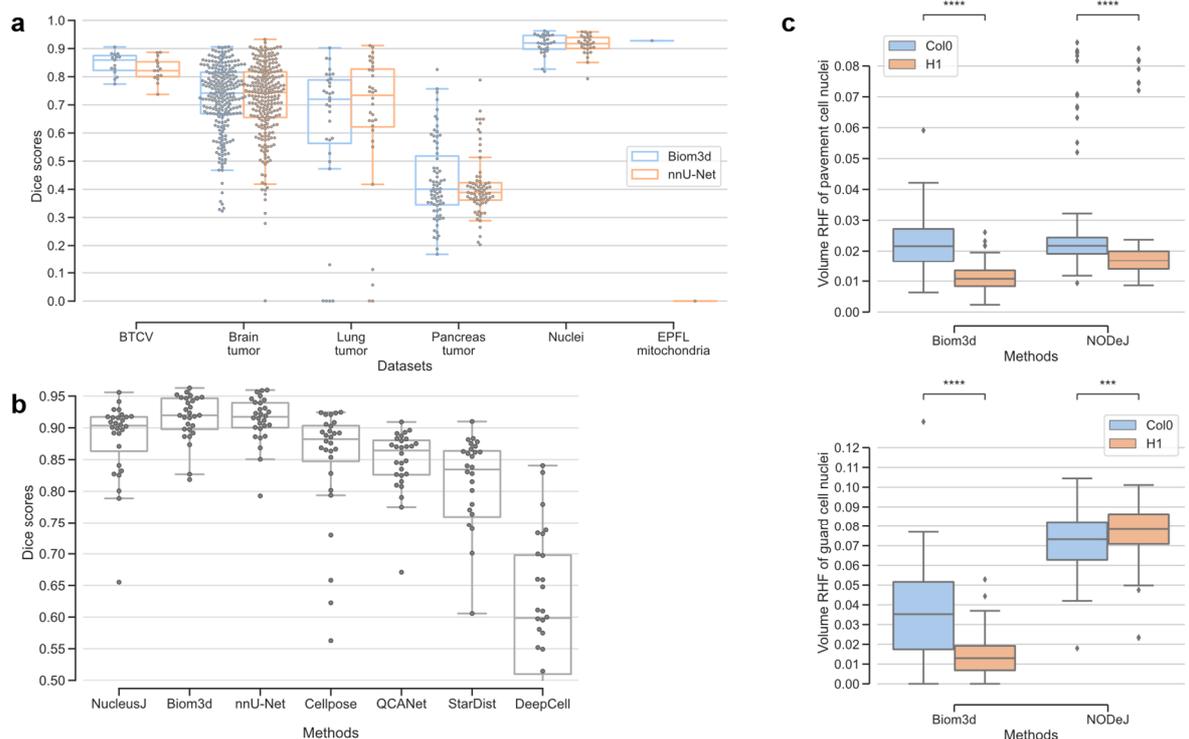
374 Figure 1 – **Method accessibility for different user profiles.** **a**, Fulfilling the needs of a
 375 continuum of user profiles. In the abscissa are conjointly represented the skills and needs of
 376 deep learning users while in the ordinate is represented the estimated time required for them to
 377 use the method and fulfil their needs. On the left, non-programmers may expect easy-to-use
 378 software with graphical user interfaces and better segmentation results than with non-deep
 379 learning methods, without the need for delving into deep learning theory. In the middle, Python
 380 programmers may look for a command line interface or a Python package as well as a
 381 performant, flexible and easily reconfigurable deep learning tool. On the right, deep learning
 382 developers may be interested in a well-documented and modular deep learning framework,
 383 where pieces of code can easily be changed, removed, or added. Most of the state-of-the-art
 384 methods are only targeting one user profile, noticeably omitting deep learning developers. **b**,
 385 Difference between a modular code and a coupled code. Each piece of code represents a node
 386 in these graphs. Each edge is an interaction between two distinct pieces of code. A modular
 387 code (*top*) means that the few code components are clearly isolated and have few clearly defined
 388 links. A coupled code (*bottom*) has code components with a low internal cohesion and are
 389 strongly intertwined.



390

391 Figure 2 – Samples of Biom3d predictions over medical datasets (*top row*) and biological
 392 datasets (*bottom row*). All views originate from testing images, unseen during the training of
 393 Biom3d default model. **a-d** and **f** originate from public datasets while the rest are custom
 394 datasets. For each example, a 3D view is depicted on top of a 2D z-slice, both representing an
 395 overlay of the segmentation result and the raw image. The rendering was done by Napari. **a**,
 396 Multi-organ abdominal challenge (BTCV) representing thirteen segmented organs in CT-scan.
 397 **b**, Three brain tumor sub-regions in MRI. **c**, A pancreas (red) and a tumor (blue) in CT-scan. **d**,
 398 A lung tumor in CT-scan. **e**, An aortic lamella in a synchrotron microscope. **f**, EPFL dataset of
 399 mitochondria in an electron microscope. **g**, Plants cell nuclei in an electron microscope. **h**, A
 400 plant cell nucleus and its chromocenters in a structured illumination microscope. **i**, Nuclei of a
 401 mouse embryo in a confocal microscope. To help the model separate nuclei clusters, frontier
 402 regions between nuclei have been segmented (blue). **j**, An epithelial cell of a *Drosophila*
 403 embryo. Both the inner cell and cell membranes have been segmented.

404

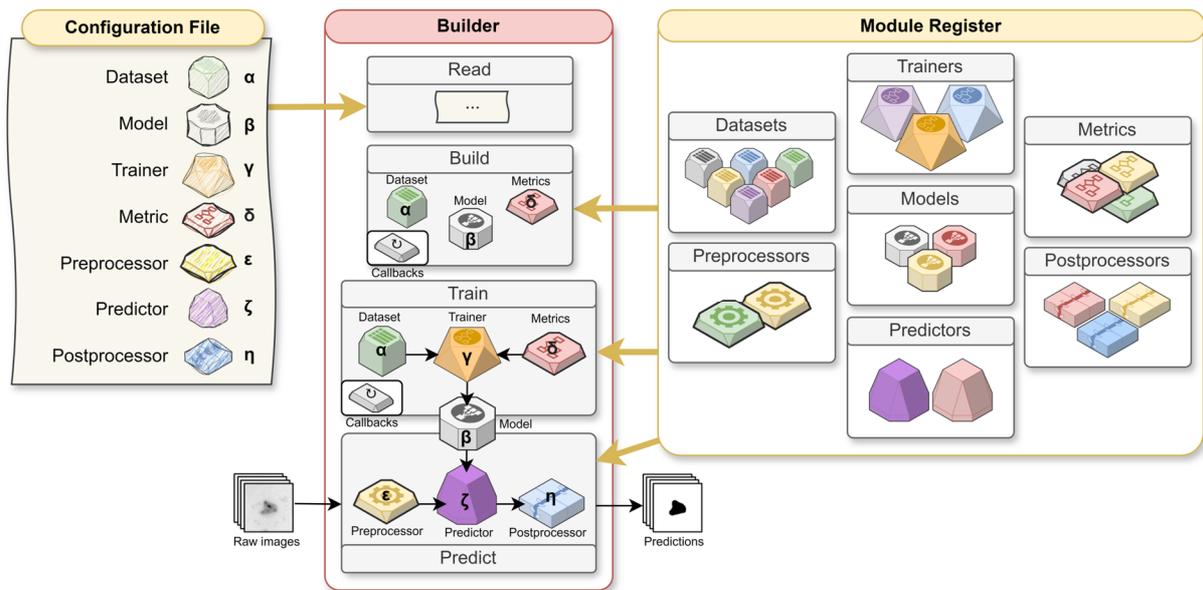


405

406 **Figure 3 – Biom3d performances.** **a**, Biom3d reaches similar Dice scores as nnU-Net on
 407 various 3D segmentation datasets. The Dice score, when applied to 3D segmentations, measures
 408 similarity by calculating twice the volume of intersection divided by the total volume of both
 409 the ground truth mask and the prediction. The higher the Dice score is, the better. The images
 410 in these datasets were captured with CT-scan (BTCV, Lung tumor, Pancreas tumor), MRI (Brain
 411 tumor), fluorescent microscope (Nuclei), and electron microscope (EPFL mitochondria). Public
 412 datasets from medical segmentation challenges (BTCV, Brain tumor, Lung tumor, and Pancreas
 413 tumor) were selected for their diversity in the number of object classes to segment. These
 414 datasets were split in two, with the first half of the images and masks used for training and the
 415 second half for testing. For the custom dataset Nuclei, 65 images and masks were used for
 416 training, while the remaining 28 were used for testing. The EPFL dataset has only one publicly
 417 available training image and one testing image. Training and testing images were identically
 418 chosen for Biom3d and nnU-Net. **b**, Benchmark of 7 segmentation methods on a custom dataset
 419 of 3D plant nuclei. NucleusJ is a non-deep learning method specialized in 3D nucleus
 420 segmentation. Only Biom3d, nnU-Net and StarDist provided the requirement to train a new 3D
 421 segmentation model. Without a way to retrain them, Cellpose, QCANet and DeepCell could
 422 only be tested with their pretrained models. Only Biom3d and nnU-Net exceeded the average
 423 Dice score obtained with NucleusJ. **c**, Example of Biom3d application to obtain biological
 424 insights. The volume of relative heterochromatin fraction (RHF) refers to the ratio between the

425 volume of voxels located in chromocenter regions and of voxels located elsewhere in the
426 nucleus. The box plots represent samples of *A. thaliana* nuclei sorted between guard cells (*top*)
427 and pavement cells (*bottom*). Biom3d is compared to NODeJ, a non-deep learning add-on of
428 NucleusJ, on the task of segmenting chromocenters in Col0 and in H1 mutant nuclei. While
429 obtaining similar results on pavement cells, Biom3d predicted a shift in the volume distribution
430 inverse to that predicted by NODeJ. The validity of this prediction was manually verified by
431 experts. This result illustrates the importance of precise segmentations when assessing a
432 biological hypothesis.

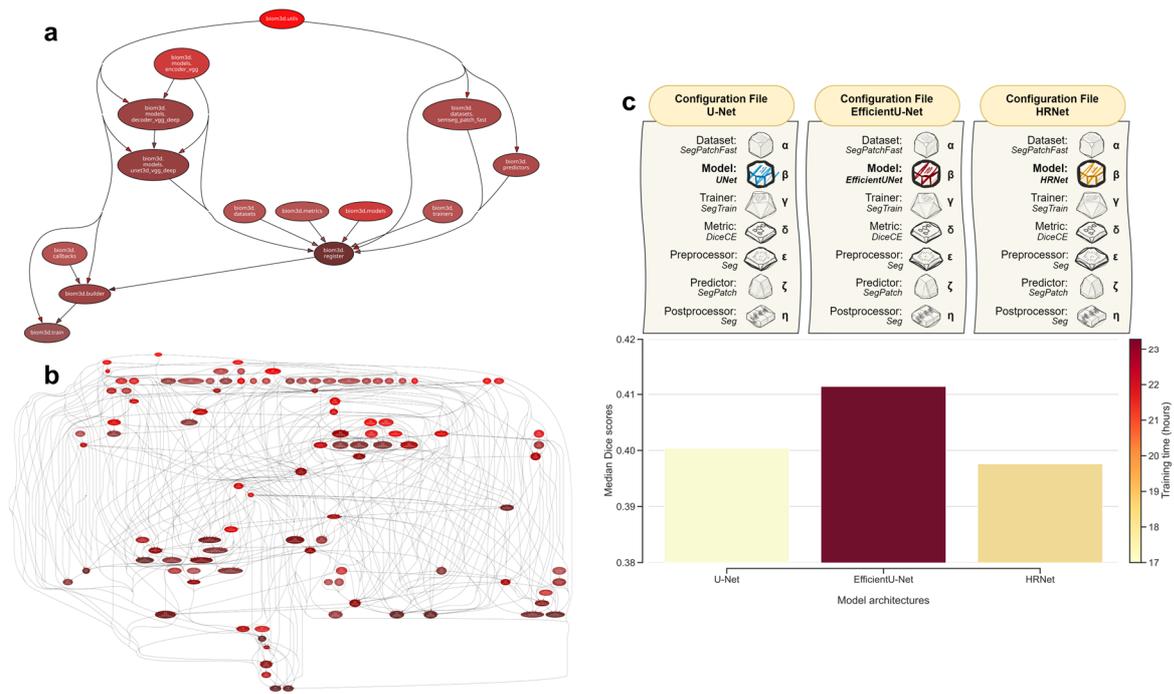
433



434

435 Figure 4 – **Configuration File, Builder and Module Register.** (Left) The Configuration File
 436 lists the names of existing Modules appearing in the Module Register and defines their
 437 parameters (*Greek letters*). (Middle) The Builder reads parameters and Module names in the
 438 Configuration File. It then retrieves the corresponding modules from the Module Register and
 439 builds them with their parameters. The Builder can then be used to train a Model with a Dataset,
 440 a Trainer, and a Metric. Once trained, the Model can be used by the Builder to predict the
 441 segmentation masks of some raw images using a Preprocessor, a Predictor, and a Postprocessor.
 442 (Right) The Module Register lists all existing modules of Biom3d, except for Callbacks which
 443 are built in the Builder. There are currently seven different types of modules in the Register,
 444 and each type has different variants (*colour shades*).

445



446

447 **Figure 5 – Code modularity of Biom3d.** **a**, Representation of Biom3d code architecture. Each
 448 node in this directed acyclic graph is one Python script in Biom3d. Each directed edge is one
 449 call of one script by another. This graph was automatically generated by *pydeps*, a Python
 450 dependency visualizer, and by calling the default training workflow of Biom3d (*bottom left*
 451 *node*). The training process instantiates a Builder which then consults the Register to instantiate
 452 requested modules from the Configuration file. **b**, Representation of nnU-Net code architecture.
 453 For comparison with Biom3d, this graph was also generated by *pydeps* and depicts the nnU-
 454 Net training workflow. This graph is more coupled than Biom3d graph, indicating that nnU-Net
 455 code will probably be harder to understand or modify. **c**, Exploiting modularity to benchmark
 456 model architectures. Once new model architectures have been added to Biom3d, a benchmark
 457 experiment can be conducted by only changing the model's name in the configuration files,
 458 while leaving the rest of the training hyper-parameters untouched. The default U-Net model in
 459 Biom3d is here compared to an EfficientU-Net and a HRNet, both being integrated into
 460 Biom3d. These model architectures are trained and tested on the Pancreas tumor dataset. The
 461 results of this experiment indicate that, despite being slower, the EfficientU-Net model is more
 462 appropriate for this dataset than the others.

463

465 **Methods**

466 **Requirements.** Biom3d is a Python3 package essentially relying on numpy and Pytorch
467 packages. Image reading and saving is ensured via scikit-image, SimpleITK and tiff file. Deep
468 learning computation can be either executed on a CPU or a GPU. GPU execution requires the
469 installation of CUDA and CuDNN libraries compatible with the Pytorch library as well as a
470 GPU having at least 10 Gb of VRAM for training and 4 Gb for prediction.

471 **Graphical User Interfaces.** The online interface is based on Google Colab and thus does not
472 require accessing a local GPU. Its use requires the connection to a Google account. The training
473 and prediction datasets must be uploaded to a Google Drive accessible via the interface. The
474 code of the online interface is hidden for ergonomic reasons, but is accessible, and it has been
475 made as simple as possible, exclusively using Biom3d framework. The online interface allows
476 to preprocess a training dataset, to auto-configure a training process, to execute the training and
477 to make prediction with a trained model. To take advantage of Biom3d modularity, the
478 Configuration file is directly editable in Google Colab. While requiring access to a computing
479 GPU, the local interface allows a higher degree of flexibility (Supplementary Figure 1 and
480 Supplementary Figure 2). The user can start training or predicting on a local computer or on a
481 private remote Linux server. In both cases, the interface is dynamic depending on user choices.
482 If the user has a trained model and intends to fine-tune it, the interface displays a field for the
483 path of the trained model and a field for the new dataset. If the user has a preprocessed dataset,
484 only the configuration path will be asked. During prediction, if the user has access to an
485 OMERO server (Supplementary Figure 3), the interface will ask for the user credentials and the
486 OMERO dataset number. The OMERO dataset is then downloaded on the computing server,
487 used for prediction, and the resulting images are eventually uploaded to a new OMERO dataset.
488 For the remote interface, the user is additionally asked to provide a name to the training dataset
489 which will automatically be sent to the computing server. After prediction, resulting images can
490 either be downloaded locally or sent to an OMERO server.

491 **Preprocessing.** Preprocessing consists in adapting a dataset to the training workflow and,
492 inversely, adapting the training workflow to the dataset. The default preprocessing methodology
493 has been designed and tested for 3D image segmentation. In the following explanation, masks
494 are manually segmented images. The default preprocessing in Biom3d is a four-step process:
495 (1) data reading and scanning to extract the data characteristics, (2) auto-configuration to create

496 a Configuration file adapted to the data characteristics, (3) data normalization to uniformize
497 image and mask dimensions and intensities, and (4) data splitting to separate training and
498 validation sets. Preprocessing is illustrated in Supplementary Figure 4.

499 *Data scanning.* For any processing requiring image reading (data-scanning, data-loading,
500 preprocessing, etc.), Biom3d adaptively reads 3D images stored in NumPy format (.npy), in
501 TIFF format (.tif or .tiff) or in any other medical formats read by SimpleITK library (see list of
502 available format on <https://simpleitk.readthedocs.io/en/master/IO.html>), such as Nifti format.
503 After preprocessing or prediction, Biom3d saves 3D images or masks along with necessary
504 meta-data. While compressed formats are appropriate for long term storage (such as Nifti or
505 TIFF format), fast reading formats are preferred for deep learning applications (such as NumPy
506 format). After preprocessing, training images and masks are thus stored in NumPy format. After
507 prediction, Biom3d stores the resulting masks using the input image format and meta-data. As
508 TIFF tagging system is unrestrictive regarding meta-data formatting, Biom3d image reader
509 might not properly read meta-data originating from some proprietary formats, except if TIFF
510 images have successfully been processed with Bio-format. The current image reader of Biom3d
511 does not support proprietary formats, such as CZI. Once read, Biom3d extracts the dataset
512 fingerprint from the images. This fingerprint includes the median image size, the mean and
513 standard deviation of the image voxel intensities within the mask, the 5% and 95% percentiles
514 of these intensities, and, if available, the median sampling, which represents the voxel
515 dimensions in meters.

516 *Auto-configuration.* The auto-configuration process of Biom3d mainly consists in finding the
517 best patch size and batch size to provide to the deep learning model, as well as the number and
518 dimensions of the successive pooling layers occurring in the U-Net model. Biom3d heuristics
519 produce similar results as nnU-Net ones, yet they have been simplified. Biom3d auto-
520 configuration process only uses the median size of the input images. The ideal patch size is set
521 to respect the median size proportions, while the product of its dimensions is smaller than a
522 maximum patch size, by default defined as $(128, 128, 128) = (d_1^*, d_2^*, d_3^*)$, a value chosen for
523 being adapted to GPUs with less than 10 Gb of VRAM. More formally, the goal is to find the
524 exponents α_i in $\prod_{i=1}^n \frac{D_i}{(1+\epsilon)^{\alpha_i}} = \prod_{i=1}^n d_i^*$, where $n = 3$ is the number of dimensions, D_i , $d_i =$
525 $\frac{D_i}{(1+\epsilon)^{\alpha_i}}$ and d_i^* are the i -th dimension of the median size, the patch size and the maximum patch
526 size respectively, and ϵ is a small number typically equal to 10^{-3} . If all exponents α_i were

527 assumed to be identical, their ideal value is $\alpha^* = \frac{1}{n \log(1+\epsilon)} \sum_{i=1}^n \log\left(\frac{D_i}{d_i^*}\right)$. This value would
528 ensure the preservation of the median size proportions. Yet, for large and strongly anisotropic
529 images, this may cause some dimensions of the ideal patch size to be set smaller than 1. To
530 prevent this, an upper bound u_i on α_i 's values is set to be $d_i = \frac{D_i}{(1+\epsilon)^{\alpha_i}} \geq 1 \Leftrightarrow \alpha_i \leq \frac{\log(D_i)}{\log(1+\epsilon)} =$
531 u_i . If $S = \{i \in \llbracket 1, n \rrbracket, \alpha_i > u_i\} \neq \emptyset$, then $\forall i \in S, \alpha_i = u_i$, and $\forall i \notin S, \alpha_i = \frac{1}{n-|S|} (n\alpha^* -$
532 $\sum_{i \in S} u_i)$, so to preserve $\sum \alpha_i = n\alpha^*$. This process is reiterated until $S = \emptyset$. The final values of
533 α_i can then be used to obtain $d_i = \frac{D_i}{(1+\epsilon)^{\alpha_i}}$, the patch size. For each dimension, the number of
534 the 2-pooling in the U-Net model is then defined by $p_i = \max\left(0, \min\left(\left\lfloor \log_2\left(\frac{d_i}{\max_i(d_i^*)}\right) + \right.$
535 $\left. p^* \right\rfloor, p^*\right)$, where p^* is the user-defined maximum number of 2-pooling, by default set to 5. To
536 prevent features maps having non-integer dimensions after dimension reduction during 2-
537 pooling, the patch size dimensions d_i are readjusted one last time to the closest multiple of 2^{p_i} .
538 The batch is set to 2 and eventually increased if the GPU VRAM allows it. The auto-
539 configuration results are stored in a new Configuration file in Python format, based on the
540 default Configuration file included in Biom3d.

541 *Data normalization.* Data normalization encompasses image reshaping and resizing and voxel
542 intensity normalizing (Supplementary Figure 5). Images are automatically reshaped to conform
543 to the standard (channel, height, width, depth) dimensions. Biom3d accepts a wide range of
544 dimension variants, even within the same dataset. If not specified by the user, Biom3d
545 automatically detects the location of the channel dimension in 4-dimensional images. For each
546 mask, a series of check-ups are performed to automatically spot and, eventually, correct
547 annotation mistakes. For instance, within the same dataset, it can be found manual annotations
548 with an inconsistent number or order of dimensions. If annotating with levels of grey, users
549 could also mistakenly decide to use inconsistent levels of grey or more levels of grey than
550 existing classes of objects. Depending on the mistake, Biom3d either corrects or warns the user.
551 For instance, if two classes of object are expected and more than two were found in the mask,
552 Biom3d automatically applies a threshold by considering as background the most recurrent
553 class. In case where three or more classes are expected and even more classes are found in the
554 mask, Biom3d displays an error. Biom3d is compatible with masks having 4 or 3 dimensions.
555 Image intensities are then Z-normalized using either the mean and standard deviation of the
556 image voxel intensities, or, if available, using the median mean and standard deviation of the

557 dataset voxel intensities retrieved during data scanning. If the median sampling could be
558 retrieved, images and masks are resized to all match the media sampling. Images are resized
559 with trilinear interpolation while masks with nearest neighborhood. Finally, for each object
560 class in mask, a random sampling is performed to extract some foreground locations. These
561 values will be used during data loading to rapidly locate foreground regions during image patch
562 cropping. Once preprocessed, the images, masks and foreground location are stored in
563 automatically created output folders. The output folder paths are added to the Configuration
564 file.

565 *Data splitting.* Following the cross-validation strategy, the dataset is by default split into 5
566 subsets called folds (Supplementary Figure 5). More formally, each image filename is
567 associated with a random integer between 0 and 4. During training, if fold 0 is selected, images
568 associated with 0 will be considered as validation images while remaining images will be
569 considered as training images. If less than 10 images are present in the dataset, Biom3d reduces
570 the number of folds to $\max\left(\left\lfloor \frac{N}{2} \right\rfloor, 2\right)$, where N is the number of images in the dataset. If only
571 one image is found in the dataset, Biom3d split the image and mask in two along the largest
572 image dimension – 80% of the image will be used for training and 20% for validation.
573 Filenames with associated fold indices are stored in a CSV file that will be loaded during
574 training. The CSV file path is added to the Configuration file.

575 **Data loading.** Data loading consists in loading a batch of preprocessed data into computer
576 memory and eventually performing data augmentation, for it to be prepared for the training
577 workflow (Supplementary Figure 6). Several data loading modules for 3D image segmentation
578 are available in Biom3d: one based on nnU-Net batchgenerators package, one based on TorchIO
579 SubjectsDataset, and one based on Pytorch Dataset. The latter one is the default and will be
580 detailed here (Supplementary Figure 6). The data loading module initialization first loads the
581 image and mask filenames of training and validation sets using the CSV file created during data
582 splitting. To fasten data loading, images and masks can be, on demand, loaded into computer
583 memory. The initialization then creates the data augmentation transformations used during
584 training with TorchIO package: random cropping, random affine transform, random anisotropy,
585 random flipping, random intensity variation, random blurring, random noise, random patch
586 swapping, and random contrast variation. The axes of the rotation in the random affine
587 transform and of the random anisotropy depends on the patch size anisotropy. By default,
588 anisotropy transforms, and rotation transforms are applied to every axis. If any patch size
589 dimensions are bigger than three times the smallest patch size dimension, then only these

590 dimensions are considered as valid axes to apply random anisotropy. In this scenario, only the
591 smallest dimension is considered a valid axis for rotation transforms. Once this initialization is
592 completed, the data loader is ready to use. During training data loading, a batch of foreground
593 locations, images and masks are loaded into computer memory. The image and mask are then
594 randomly cropped using the patch size. To do so, several constraints are considered. First, if
595 rotation transforms are applied, the patch size is temporally enlarged beforehand to avoid empty
596 corners appearing in images and masks once rotated. The enlarged patch size is set to be cubic
597 and its dimensions equal to the largest diagonal of the patch. Second, foreground locations are
598 used to crop one image and mask pair out of three in the batch. Among the list of possible
599 foreground locations, one is selected to be the location of the center of the patch. Third, if the
600 cropped image and mask are smaller than the patch size, then they are padded with zeros equally
601 in all three dimensions. Once cropped, the images and masks are transformed with the rest of
602 the data augmentations before being output. During validation data loading, only the random
603 cropping is performed without any other form of data augmentation.

604 **Dynamic models.** Biom3d includes several deep learning model definitions. The default deep
605 learning model is a standard 3D U-Net model for semantic segmentation¹. It has been
606 implemented in a modular fashion, meaning that the encoder, the decoder and the
607 convolutional blocks can work independently. More specifically, the default VGG encoder²
608 can straightforwardly serve as an independent classification model or can be replaced by any
609 other encoder in the 3D U-Net model, such as the 3D EfficientNet³ included in Biom3d, the
610 MONAI models, or any Pytorch model compatible with 3D images. The default 3D U-Net
611 model architecture is dynamic and depends on the number p_i of successive pooling layers found
612 for each dimension during the auto-configuration. The maximum number of 2-pooling along a
613 given dimension k is determined by $\min(p_k, \max_i p_i)$. For each case where $p_k < \max_i p_i$, there
614 will be $\max_i p_i - p_k$ 1-pooling layers, evenly distributed between the head and tail of the
615 encoder. For instance, if the auto-configuration sets the ideal number of successive pooling
616 layers to be (3,5,5), the pooling layers of the 3D U-Net will have the following list of kernel
617 dimensions: ((1,2,2), (2,2,2), (2,2,2), (2,2,2), (1,2,2)). The definition of the rest of the model
618 layers follows the original U-Net architecture.

619 **Training.** Losses, metrics, callbacks as well as training and validation routines are all
620 independent modules in Biom3d, the default ones being designed for 3D semantic
621 segmentation. The default training loss is the sum of the class Dice score and the cross-entropy

622 between the predicted and annotated masks. Default validation metrics include the Dice score
623 and the intersection over union between the predicted and annotated masks. Even if the default
624 behavior of these metrics requires the 3D annotated masks to be formatted such that each pixel
625 value of the mask represent a single object class (0, 1, 2, etc.), Biom3d metrics are also
626 compatible with 4D annotated masks formatted with an additional channel dimension
627 representing each individual object class. In such setup, Biom3d metrics accept input pixel to
628 be associated with multiple object classes. The default optimizer is the stochastic gradient
629 descent with a Nesterov momentum of 0.99 and a weight decay of $3e-5$. The default training
630 routine uses the data loader to get a batch of data, passes it to the model then to the loss,
631 computes the gradients and clips their norms, before updating the model parameters using the
632 optimizer. The training routine periodically calls the callbacks (Supplementary Figure 7) to
633 update the learning rate with cosine annealing, to print and store information about the training
634 and validation, and to store the model parameters of the best performing model. Training and
635 validation stored information includes the training and validation losses per epoch and
636 prediction snapshot on the validation set. After loading the Configuration file, the Builder
637 oversees the instantiation of the losses, metrics, callbacks and optimizer before executing the
638 training and validation routines. This execution can be done using mixed precision and in
639 parallel on multiple GPUs. For reproducibility reasons, once the training is finished, the output
640 folder includes the Configuration file and the data splitting file, in addition to all the other
641 information stored by the training callbacks. If interrupted, training can thus be restarted by
642 solely using this output folder. Biom3d also allows to perform model finetuning or retraining
643 by instantiating a new Builder using both an output folder and a new Configuration file.

644 **Predicting.** Predicting with Biom3d encompasses three steps: (1) pre-processing, by default
645 including image reading and normalization, (2) predicting, by default involving tiling the
646 images and passing the tiles to the model, and (3) post-processing, by default implying
647 discretizing model outputs, removing noise and saving the predicted masks. Each of these steps
648 is an isolated Biom3d module. The modularity of Biom3d allows to reuse the exact same
649 function for pre-processing as the one used to normalize training data. Before being passed to
650 the deep learning model, input images are tiled with TorchIO grid sampler. For one input image,
651 this function creates a series of patches equally distributed and overlapping by an overlap equal
652 to half of the patch dimensions. To respect these two tiling constraints, the input image is
653 eventually padded with zeros. Batch of patches are constituted and given to the model. To
654 increase the prediction accuracy, each batch is augmented by flipping along the seven possible

655 combinations of (x, y, z)-axis. Augmented predictions are then flipped back before being
656 averaged. Predicted tiles are then aggregated using Hann filtering to reduce edge artefacts in
657 patch overlapping regions. As data pre-processing involved data resampling, the aggregated
658 prediction is finally resized back to the original image dimensions. As Biom3d framework is
659 compatible with ensemble learning, it is also possible to aggregate predictions coming from
660 different models. If such case, model outputs are simply averaged before being post-processed.
661 Post-processing then starts by discretizing the prediction. For 4D masks, discretizing means to
662 apply a threshold of 0.5 to the output of the sigmoid function applied to the model output. For
663 3D masks, discretizing means to retrieve the argmax of the output of the SoftMax function
664 applied to the model output. On user request, two distinct noise removal strategies can then be
665 applied to remove too small, segmented regions. Connected components can thus be computed
666 to retrieve either the biggest segmented object or objects which volumes are higher than an Otsu
667 threshold determined using the volume distribution of all connected components. Finally, the
668 post-processed prediction is automatically saved along with input image metadata.

669 **Evaluation.** Controlling the quality of a trained model can be done with Biom3d by either using
670 the local graphical user interface or the application programming interface. For 3D
671 segmentation, one folder containing 3D ground truth masks of images different from the
672 training set and another folder with the corresponding predictions can be passed to Biom3d to
673 retrieve the average Dice score on this set. To fasten experiments and benchmarking, Biom3d
674 also includes scripts that allow the preprocessing, training, prediction, and evaluation of a new
675 deep learning model on a new dataset to be executed with a single command

676 **Data availability**

677 All public datasets displayed on Figure 2 or used to benchmark Biom3d in Figure 3 can be
678 accessed via their respective website: Beyond The Cranial Vault (BTCV)-Abdomen,
679 <https://www.synapse.org/Synapse:syn3193805>; the Medical Segmentation Decathlon (Brain
680 tumor, Lung tumor and Pancreas tumor), <http://medicaldecathlon.com>; EPFL electron
681 microscopy dataset, <https://www.epfl.ch/labs/cvlab/data/data-em>. The rest of the data created
682 in the frame of this work is hosted by Mésocentre UCA on a public OMERO server
683 <https://omero.mesocentre.uca.fr/webclient/userdata/?experimenter=352>. Trained models and
684 predictions created by Biom3d on the public datasets can be found in the following OMERO
685 project: <https://omero.mesocentre.uca.fr/webclient/?show=project-2005>. Trained models and
686 custom training datasets used to create the illustration in Figure 2 for the X-ray microscopy

687 image of aorta, for the electron microscopy image of plant root nuclei, for the confocal
688 microscopy images of the mouse embryo and the Drosophila embryo can be found in the
689 following OMERO project: <https://omero.mesocentre.uca.fr/webclient/?show=project-2007>.
690 For all custom datasets, the specific images and predictions represented on Figure 2 have been
691 set aside in the following OMERO project:
692 <https://omero.mesocentre.uca.fr/webclient/?show=project-2006>. The plant nucleus dataset
693 used to benchmark the nucleus segmentation methods in Figure 3 can be found in the
694 following OMERO project: <https://omero.mesocentre.uca.fr/webclient/?show=project-2002>.
695 For the previous dataset, the training data and the trained model of Biom3d can be found in
696 the following OMERO project: [https://omero.mesocentre.uca.fr/webclient/?show=project-](https://omero.mesocentre.uca.fr/webclient/?show=project-2001)
697 [2001](https://omero.mesocentre.uca.fr/webclient/?show=project-2001). The plant chromocenter dataset used to benchmark Biom3d and NODeJ in Figure 3 can
698 be found in the following OMERO project:
699 <https://omero.mesocentre.uca.fr/webclient/?show=project-2004>. For the previous dataset, the
700 training data and the trained model of Biom3d can be found in the following OMERO project:
701 <https://omero.mesocentre.uca.fr/webclient/?show=project-2003>.

702 **Code availability**

703 Biom3d is a public python package referenced in the Python Package Index
704 (<https://pypi.org/project/biom3d/>). The latest version of the code and documentation of Biom3d
705 can be found on GitHub (<https://github.com/GuillaumeMougeot/biom3d>).

706 **References**

- 707 1. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical
708 image segmentation. *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif.*
709 *Intell. Lect. Notes Bioinformatics)* **9351**, 234–241 (2015).
- 710 2. Simonyan, K. & Zisserman, A. Very Deep Convolutional Networks for Large-Scale
711 Image Recognition. *CoRR* **abs/1409.1**, (2014).
- 712 3. Tan, M. & Le, Q. V. EfficientNet: Rethinking model scaling for convolutional neural
713 networks. *36th Int. Conf. Mach. Learn. ICML 2019* **2019–June**, 10691–10700 (2019).

714 **Acknowledgements**

715 We would like to thank Adama Nana for his support with benchmarking nucleus segmentation
716 methods, the bioinformatics core facility of the Au-Bi (Mesocentre UCA) and the iGReD CLIC

717 microscopy facilities. This work was partially achieved using HPC resources from GENCI–
718 IDRIS (Grant 2022-AD011013709) on the supercomputer Jean Zay's A100 partition. We thank
719 Dr. Fredy Barneche (CNRS, IPBS Sorbonne) and Sara Farrona (University of Galway, Ireland)
720 for providing us with seeds; Xiaowen Liang (Université de Reims Champagne Ardenne) for
721 providing us with annotated images of rat aorta; Cynthia Dennis (iGReD) for providing us with
722 an annotated dataset of epithelial cells of drosophila ovaries and Nicolas Allègre (iGReD) for
723 stained mouse embryos. Students in the Master 1 Bioinformatics program (UCA) have
724 contributed to the annotation of nuclei images and the drafting of documentation for biologists
725 (graduating classes 2023-2024 and 2025).

726 This work was supported by Agence Nationale de la Recherche of the French government
727 through the programme ‘Investissements d’Avenir’ (16-IDEX-0001 CAP 20-25), ‘Dynam’Het’
728 ANR-11 JSV2 009 01 and ‘SINUDYN’ ANR-12 ISV6 000; Oxford Brookes University,
729 Université Clermont Auvergne, Centre National de la Recherche Scientifique, Institut National
730 de la Recherche et de la Santé, the Université Reims Champagne Ardenne, the European
731 Regional Development Fund (FEDER), EMERGENCE (16-IDEX-0001 CAP 20-25) and the
732 trainee grants funding GDR IMABIO (France). A.P., C.T., D.E.E., K.G. and S.D. are part of the
733 International Plant Nucleus Consortium (IPNC; <https://radar.brookes.ac.uk>) and the European
734 Cooperation in Science and Technology COST-Action CA16212 (INDEPTH;
735 <https://indepth.brookes.ac.uk/>). The authors acknowledge Synchrotron SOLEIL to have
736 provided beamtime for Figure 2 under project no. 20211303. ANATOMIX is an Equipment of
737 Excellence (EQUIPEX) funded by the Investments for the Future program of the French
738 National Research Agency (ANR), project NanoimagesX, grant no. ANR-11-EQPX-0031.

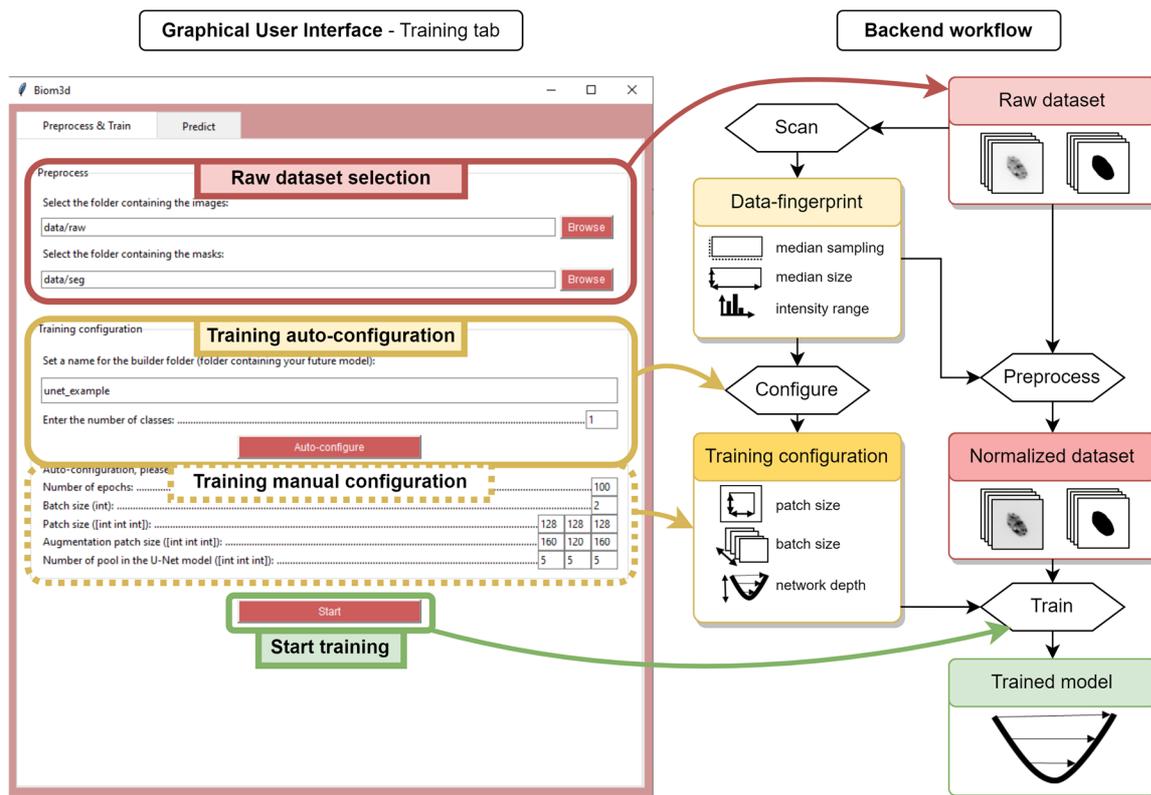
739 **Author contributions**

740 G.M. conceptualized the coding philosophy and undertook most of the programming work and
741 experiments. D.E.E., F.C., C.T., E.P. and S.D. provided ideas for the project and obtained
742 funding. S.S. helped debug and improve the user interface. G.M. and S.D. wrote the manuscript
743 with input from C.T., F.C., D.E.E. and E.P.

744 G.M. and S.D. performed data analysis and prepared the figures. A.P., H.A., P.P., S.D., S.A. and
745 N.F. provided annotated images and were beta testers of Biom3d.

746 **Competing interests**

747 The authors declare no competing or financial interests.



750

751 Supplementary Figure 1 – **Training workflow of Biom3d in the local graphical interface.**752 **(Left) Training tab of the local graphical interface.** The user specifies the path to the folders

753 containing training images and annotations, then defines a name for the Configuration file and

754 the future trained model, and the auto-configuration can start. Automatically defined parameters

755 can be adjusted manually, if needed, before starting the training. **(Right) Backend workflow.**

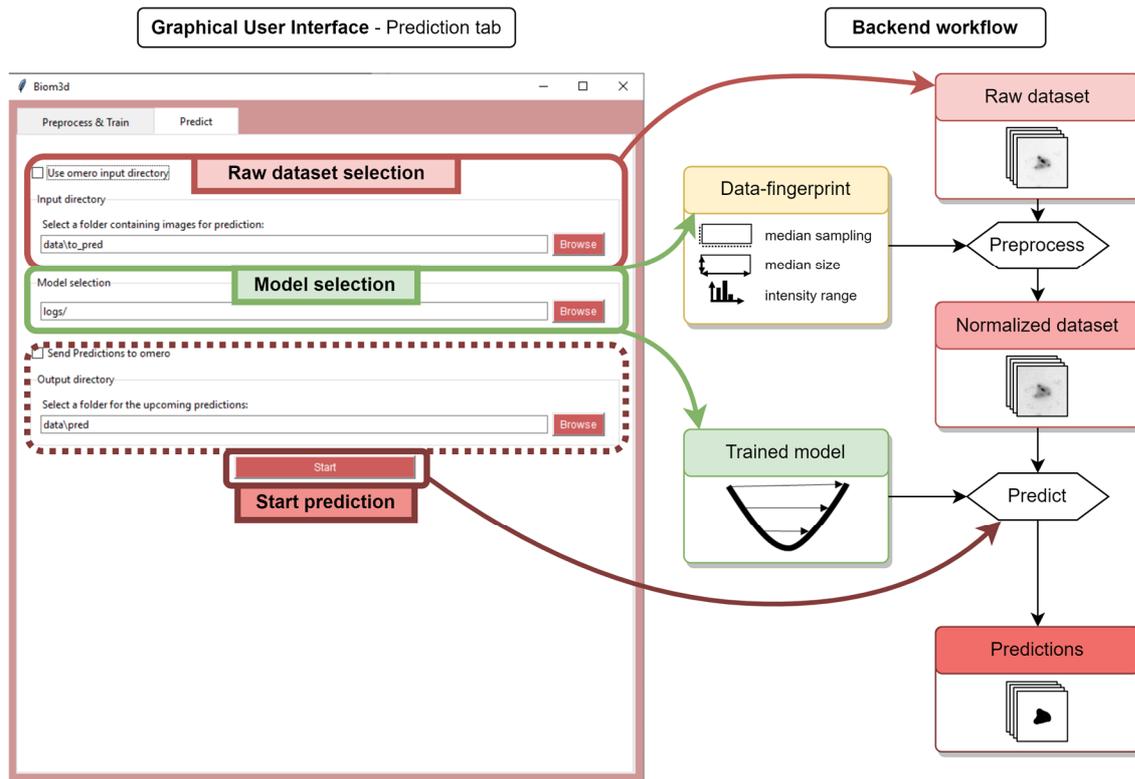
756 Once the “Auto-configuration” button is pressed, the data-preprocessing starts. The dataset key

757 elements (median shape etc.) are extracted and used to normalize all the images and to define

758 training configuration (patch size etc.). If the “Start” button is pressed, a deep learning model

759 will be trained and saved along with the pre-processing methodology (data-fingerprint).

760

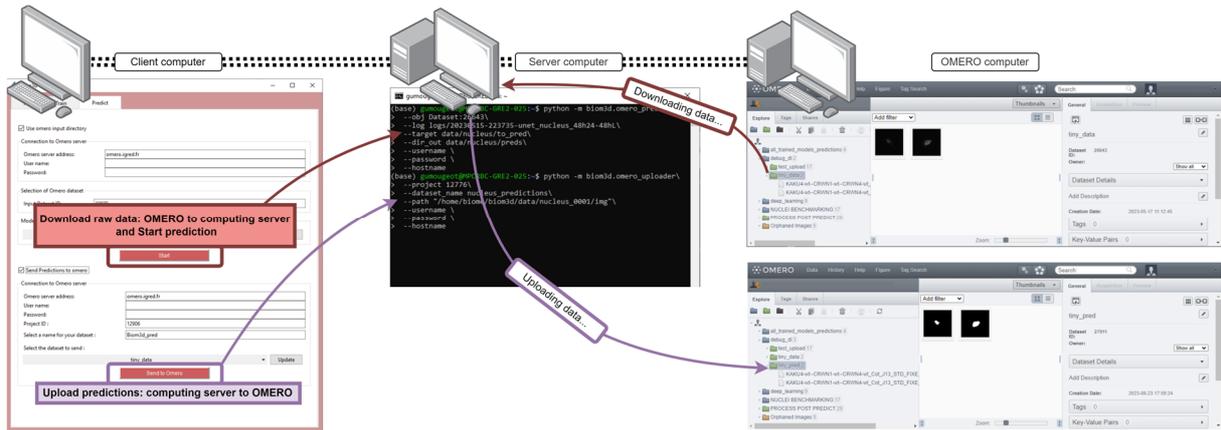


761

762 Supplementary Figure 2 – **Prediction workflow.** (Left) Prediction tab of the graphical
 763 interface. The user chooses a folder containing raw images, the path to a trained model and a
 764 folder for the future predictions. The predictions start when the “Start” button is pressed.

765 (Right) Backend workflow. The raw images are normalized using the data-fingerprint of the
 766 training dataset. The trained model is then loaded and used to compute predictions.

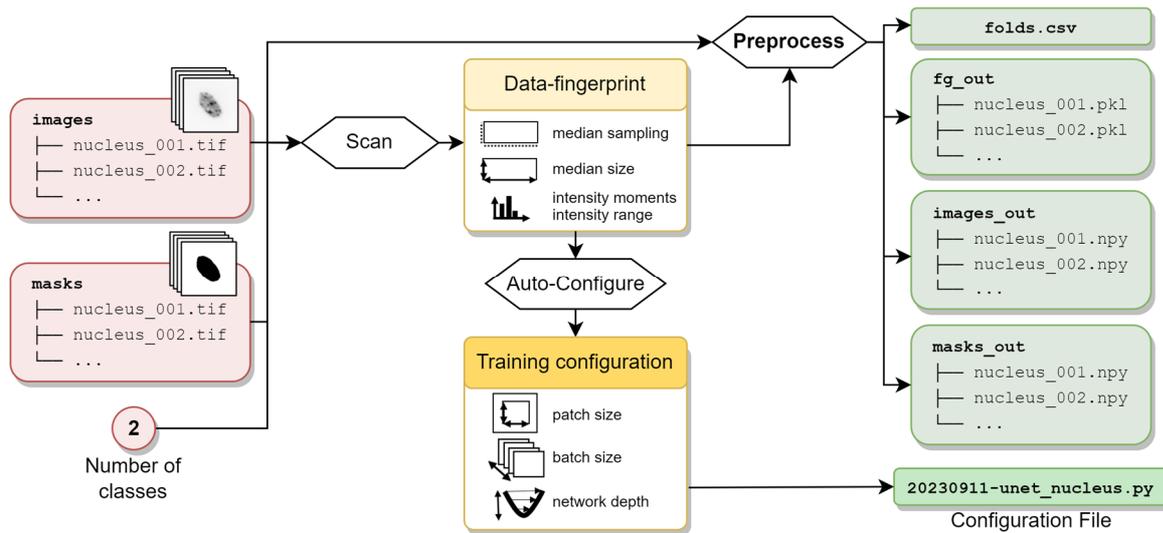
767



768

769 **Supplementary Figure 3 – Remote access and OMERO access to Biom3d.** Remote access
 770 and OMERO access can work together or independently. A combined use case is represented.
 771 The local interface (*left*) allows the access to a remote server remote (*middle*) running Biom3d
 772 training or prediction workflows. If used in combination with an OMERO server (*right*), the
 773 remote server will download raw images from an OMERO dataset and upload them back in a
 774 new OMERO dataset.

775



776

777 Supplementary Figure 4 – **Auto-configuration and Preprocessing of training dataset for**

778 **3D segmentation.** Image and mask folders (*red, left*) are scanned to extract their data-

779 fingerprint. The data-fingerprint is used to preprocess the images and masks and to

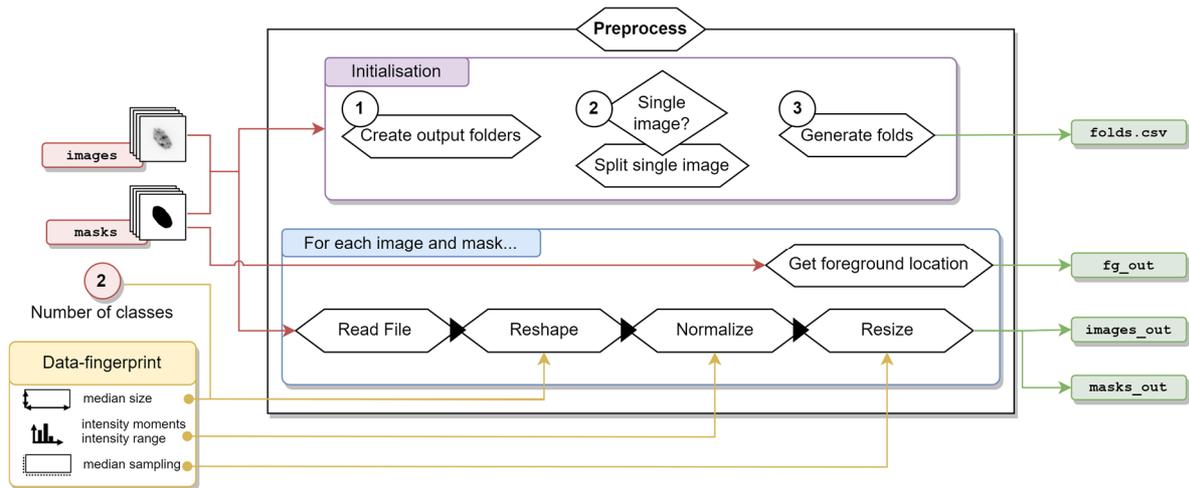
780 automatically configure the future training (*yellow, centre*). The outputs of these two steps

781 (*green, right*) are: a CSV file (*folders.csv*) describing which files will be used for training or

782 validation, a Configuration file, the pre-processed images (*images_out*) and masks

783 (*masks_out*), and the location of the foreground voxels (*fg_out*).

784



785

786 **Supplementary Figure 5 – Training data normalization and splitting for 3D segmentation.**

787 The preprocessing starts (*purple box, initialization*) by (1) creating the three output folders

788 (*fg_out, images_out, masks_out*), (2) optionally splitting single image/mask and (3) splitting

789 the dataset into training and validation folds (*folds.csv*). Afterwards, each image and mask

790 (*blue box*) are read (*Read File*), independently of their format, reshaped (*Reshape*) so to have

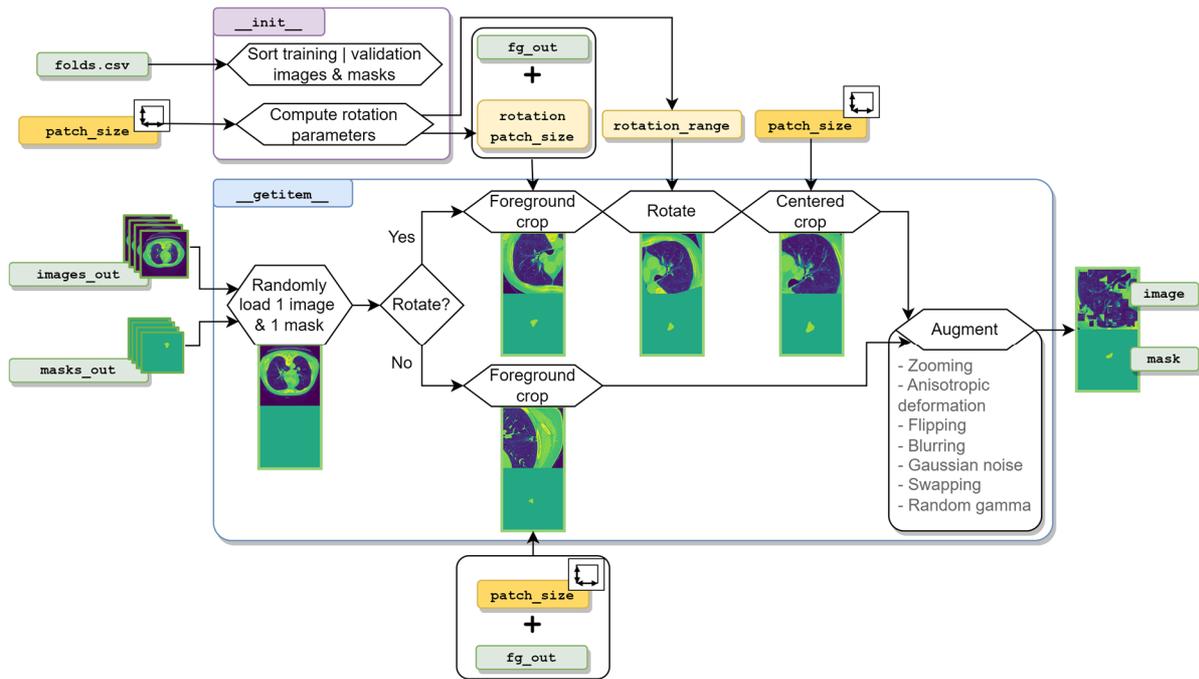
791 exactly 4 dimensions in (*channel, depth, heigh, width*) format, z-normalized for images and

792 uniformized for masks (*Normalize*), and resized (*Resize*) so all images and masks have the

793 same sampling. Finally, the locations of foreground voxels are extracted and stored in the

794 appropriate output folder (*fg_out*).

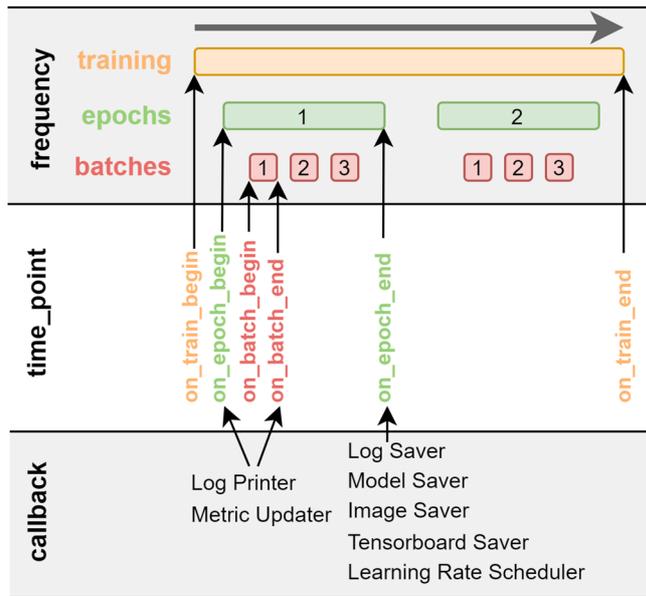
795



796

797 **Supplementary Figure 6 – Biom3d default Dataset Module for 3D segmentation.** In the
 798 `__init__` class function (*purple*), the CSV file is used to sort training images from validation
 799 images and the patch size is used to determine the parameters of the rotation transformation
 800 (rotation angle and rotation patch size, *yellow*). In the `__getitem__` class function (*blue*), one
 801 image and one mask are loaded into computer memory from their local folder. Those are then
 802 cropped in regions where foreground objects are located. If rotation augmentation is active,
 803 then the foreground crop is performed with a larger patch size before being cropped a second
 804 time to discard unwanted empty regions in the image corners. Another series of
 805 augmentations is finally applied to obtain a ready to use pair of image and mask patches.

806



807

808 Supplementary Figure 7 – **Callback Module principle**. The whole training (top row, orange)
 809 is divided into epochs (green) themselves divided into batches (red). Callbacks are Python
 810 classes that can have one or more class-functions, each representing one of 6 different time
 811 points (middle row). Biom3d currently has 7 types of Callback Modules (bottom row).

812