



Original software publication

fastMONAI: A low-code deep learning library for medical image analysis

Satheshkumar Kaliyugarasan^{*}, Alexander S. Lundervold

Department of Computer Science, Electrical Engineering and Mathematical Sciences, Western Norway University of Applied Sciences, Bergen, Norway
MMIV, Department of Radiology, Haukeland University Hospital, Bergen, Norway



ARTICLE INFO

Keywords:

Deep learning
Medical imaging
Radiology

ABSTRACT

We introduce fastMONAI, an open-source Python-based deep learning library for 3D medical imaging. Drawing upon the strengths of fastai, MONAI, and TorchIO, fastMONAI simplifies the use of advanced techniques for tasks like classification, regression, and segmentation. The library's design addresses domain-specific demands while promoting best practices, facilitating efficient model development. It offers newcomers an easier entry into the field while keeping the option to make advanced, lower-level customizations if needed. This paper describes the library's design, impact, limitations, and plans for future work.

Code metadata

Current code version	v0.4.0
Permanent link to code/repository used for this code version	https://github.com/SoftwareImpacts/SIMPAC-2023-405
Permanent link to Reproducible Capsule	
Legal Code License	Apache-2.0 license
Code versioning system used	git
Software code languages, tools, and services used	Python. See the complete list of dependencies in the project repo's <code>setup.ini</code> .
Compilation requirements, operating environments & dependencies	See the <code>setup.ini</code> file in the project repository
If available Link to developer documentation/manual	https://fastmonai.no
Support email for questions	skaliyugarasan@hotmail.com

1. Introduction

Medical imaging is a complex domain. To simplify the use of deep learning in this field, we introduce fastMONAI, a Python-based open-source deep learning library. The library, built on top of fastai [1], MONAI Core [2], TorchIO [3], and Imagedata [4], is designed to simplify the use of deep learning techniques for solving medical classification, regression, and segmentation tasks. fastMONAI makes it easy to load data, preprocess it, train models, and interpret the results. At the same time, multiple standard best practices are implicitly incorporated without the need to delve into complex configurations.

This paper discusses the need for such a research tool and its intended use. We describe deliberate design choices that enhance user experience and demonstrate its diverse applications. This is also explored in an accompanying video walk-through.¹ A section on documentation, usability, and maintainability further highlights the library's design and practicality.

^{*} Corresponding author at: Department of Computer Science, Electrical Engineering and Mathematical Sciences, Western Norway University of Applied Sciences, Bergen, Norway.

E-mail addresses: skaliyugarasan@hotmail.com (S. Kaliyugarasan), allu@hvl.no (A.S. Lundervold).

¹ https://fastmonai.no/tutorial_beginner_video

<https://doi.org/10.1016/j.simpa.2023.100583>

Received 8 September 2023; Accepted 17 September 2023

2. Software description

Deep learning develops at breakneck speed, with new models, techniques, and tricks constantly appearing. As a result, it is easy to get stuck on something less-than-optimal when using deep learning to solve a particular set of problems while also being in danger of getting lost in minor technical details when constructing models for concrete tasks. The popular PyTorch-based fastai deep learning library [1] provides both a high-level API that automatically incorporates many established best practices and a low-level API in which more expert users can modify details related to model architectures, training strategies, data augmentation, and more.

While fastai provides a valuable general foundation, the peculiarities of medical imaging, including specific medical imaging formats, data storage and transfer, data labeling procedures, domain-specific data augmentation, evaluation methods, and more, necessitate specialized tools. MONAI Core [2] and TorchIO [3] are two PyTorch-based

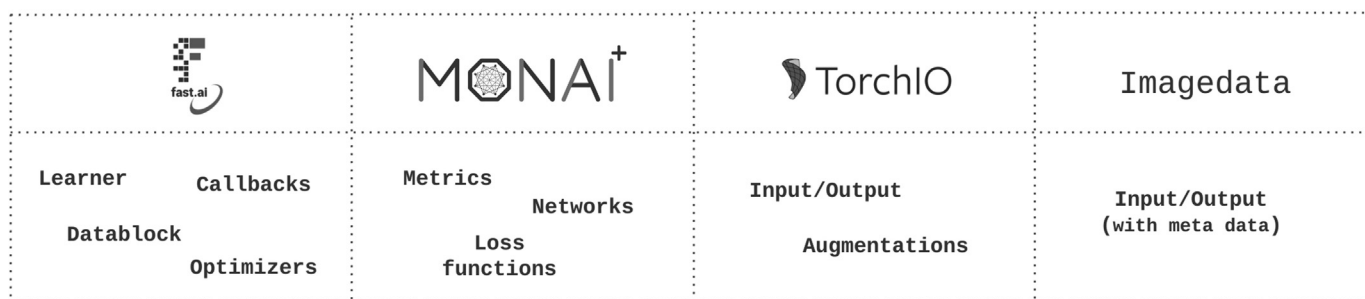


Fig. 1. MONAI Core, the primary library of MONAI (<https://monai.io/>), is built on top of PyTorch and provides domain-specific functionalities for medical imaging, including network architectures, metrics, and loss functions. TorchIO is a Python-based open-source library for efficiently loading, preprocessing, and augmenting 3D medical images.

libraries that target deep learning in healthcare imaging, incorporating multiple best practices.

fastMONAI merges fastai, MONAI Core, TorchIO, and the medical image library Imagedata [4], unifying their various capabilities. Together with its custom-made modules, this unification makes it possible to quickly construct, train, and use powerful models with a range of different architectures for a variety of medical imaging tasks while using established best practices for training, reading data, performing data augmentation, and for other domain-specific capabilities of these underlying libraries. This integration simplifies the construction of reliable baseline models and retains flexibility for more granular customizations when needed (see Fig. 1).

2.1. Documentation, usability, and maintainability

fastMONAI contains several practical tools to ensure the software's user-friendliness. Our focus on usability is reflected in the extensive documentation available at <https://fastmonai.no>, which includes a set of tutorials and source code from multiple research projects for various medical imaging tasks (e.g., classification, regression, and segmentation), together showcasing the features and use cases of the library.

We have written the fastMONAI library using nbdev, a tool for exploratory programming that allows you to write, test, and document a Python library in Jupyter Notebooks [5].

Tests are written directly in notebooks, and continuous integration with GitHub Actions runs the tests on each push, making software development easier with multiple collaborators.

To promote community contributions, we have added a short contribution guide to the project, available in its GitHub repository.

3. Impact

Using deep learning techniques in medical image analysis can be made quite straightforward when you have access to the right data. This is exemplified by our one-click tutorials for MedMNIST v2 [6] and Medical Segmentation Decathlon challenge [7] datasets. Such benchmark datasets are crucial for method development across diverse domains [8]. However, it is crucial to note that these benchmark datasets often undergo extensive preprocessing [9]. As a result, they might not always represent the intricacies of real-world data [10,11], which can be influenced by factors such as demographic nuance, image acquisition, disease spectrum, and more [12]. Furthermore, there is a pressing need for established standards in transparency and validation to build trust in deep learning-based systems [11].

In real-world clinical care, only a handful of studies employing deep learning systems have reached the deployment phase [13].

A significant contributing factor to these challenges is the limited collaboration between the medical and data science fields. To unlock the full potential of data-driven medical imaging, fostering interdisciplinary collaboration is crucial.




Fig. 2. A diagram illustrating the three primary components of fastMONAI: (1) Data pre-processing: Understanding the dataset, which includes aspects like voxel size, shape, orientations, and more. (2) Training and evaluation: Selecting the appropriate model and training strategies, and assessing model performance. (3) Deployment: Integrating the trained model into the existing infrastructure. See <https://github.com/MMIV-ML/fastMONAI/tree/master/research> for examples of models implemented in the Sectra research picture archiving and communication system (PACS) at the Western Norway Regional Health Authorities (Helse Vest RHF).

By being developed at the Mohn Medical Imaging and Visualization Centre (MMIV), part of the Department of Radiology at Haukeland University Hospital, fastMONAI benefits from direct insights and feedback from medical imaging professionals.

A key objective of MMIV and similar centers around the world is to explore and develop new methods for medical image analysis, aiming to improve decision-making and patient care. fastMONAI supports such efforts by easing the entry for new practitioners into medical AI and making it possible to quickly construct good baseline models while still being flexible enough to enable further optimizations (see Fig. 2).

Our use of the library within our institution suggests improved efficiency in model development, fostering a conducive environment for both newcomers and experienced practitioners in medical AI by lowering the threshold of entry. fastMONAI is well-suited for teaching deep learning for medical imaging, and we have successfully used the library in several hands-on workshops for a variety of audiences.

Picture archiving and communication system (PACS) are DICOM-driven medical systems with hardware and software built to facilitate storage, retrieval, and sharing of digital images in clinical settings [14]. In radiology departments, PACS are crucial for streamlining workflows. We have deployed two projects in research PACS at our hospital and shared the associated model weights on HuggingFace.² This encourages further collaboration and exploration in the community.

3.1. Ongoing research projects using fastMONAI

The fastMONAI framework was developed and evaluated through applications based on both public and clinical study data and involving multiple patient groups, and organ systems, including identifying the brain from surrounding tissue and structures [17], lung cancer [18], gynecological cancer [16], and low back pain [15]. In addition, we have replicated and further explored a recent endometrial cancer study [19], implementing the models in the research PACS infrastructure at our

² For more details, see the project's research folder: <https://github.com/MMIV-ML/fastMONAI/tree/master/research>.




Fig. 3. Illustrations of two deep learning applications deployed in our research PACS system at our local: (a) Spine segmentation and (b) Endometrial cancer segmentation. As shown, our segmentation algorithms can be run on-demand or automatically, without requiring the user to launch any external interface. For more details about the data and task, see our related publications [15,16], respectively.

hospital. In these projects, the main developers of fastMONAI played a major role. Other researchers have also adopted fastMONAI in their work:

- Segmentation of the substantia nigra for patients with Parkinson's disease by researchers at the University of Bergen (work in progress).
- A cardiothoracic radiologist from LNCT University has created a teaching unit on segmenting the left atrium using MRI scans of the whole heart. This will be highlighted at the forthcoming RSNA conference. Additional details can be found at https://fastmonai.no/tutorial_binary_segmentation.

3.2. Scholarly publications enabled by this software

- S. Kaliyugarasan, A. Lundervold, A.S. Lundervold, Brain age versus chronological age: A large scale MRI and deep learning investigation, Proceedings of European Congress of Radiology—ECR 2020 (2020) [17]
- S. Kaliyugarasan, M. Kocinski, A. Lundervold, A.S. Lundervold, 2D and 3D U-Nets for skull stripping in a large and heterogeneous set of head MRI using fastai, Proceedings of the NIK2020 (2020) [20]
- S. Kaliyugarasan, A. Lundervold, A.S. Lundervold, Pulmonary Nodule Classification in Lung Cancer from 3D Thoracic CT Scans Using fastai and MONAI, IJIMAI (2021) [18]
- E. Hodneland, S. Kaliyugarasan, et al. Fully Automatic Whole-Volume Tumor Segmentation in Cervical Cancer, Cancers 14 (10) (2022) [16]
- S. Kaliyugarasan, A.S. Lundervold, et al. Multi-Center CNN-based spine segmentation from T2w MRI using small amounts of data, 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI) (2023) [15]

4. Limitations and future improvements

- *Patch-Based Training and Inference:* As of now, fastMONAI does not offer patch-based training and inference out of the box. As this

method can enhance performance in specific scenarios, plans to integrate it in upcoming releases are underway.

- *Semi-Automatic Annotation Loop:* We have initiated research to develop a semi-automatic annotation loop for tumors. We can already integrate models developed using fastMONAI into our research PACS, ref. Fig. 3, which is currently used as part of a semi-automatic annotation loop. This system provides expert annotators with automatically segmented tumor instances for approval or further refinement. Future endeavors involve establishing, utilizing, and documenting an active learning pipeline [21,22], aiming to accelerate both the annotation process and subsequent model training.
- *PACS Integration:* Comprehensive documentation on integrating fastMONAI models with PACS remains absent. Our local research PACS system and workflow is documented at <https://github.com/mmiv-center/Research-Information-System/tree/master/components/Workflow-Image-AI>, but the integration with fastMONAI remains to be expanded upon.
- *Augmentation Documentation:* Presently, fastMONAI's documentation does not cover how to implement augmentation techniques beyond those provided by MONAI and TorchIO. Techniques such as MixUp [23] and CutMix [24] are notable omissions.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the Trond Mohn Research Foundation [grant number BFS2018TMT07], the Western Norway Regional Health Authority (Helse Vest RHF) [grant number F-12532], and The Research Council of Norway [grant number 309755] through our WIML project (<https://mmiv.no/wiml/>).

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.simpa.2023.100583>.

References

- [1] J. Howard, S. Gugger, Fastai: a layered API for deep learning, *Information* 11 (2) (2020) 108, <http://dx.doi.org/10.3390/info11020108>.
- [2] MONAI Consortium, MONAI: Medical open network for AI, 2022, <http://dx.doi.org/10.5281/zenodo.6639453>.
- [3] F. Pérez-García, R. Sparks, S. Ourselin, TorchIO: A Python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning, *Comput. Methods Programs Biomed.* 208 (2021) 106236, <http://dx.doi.org/10.1016/j.cmpb.2021.106236>.
- [4] E. Andersen, Imagedata: A Python library to handle medical image data in NumPy array subclass series, *J. Open Source Softw.* 7 (73) (2022) 4133, <http://dx.doi.org/10.21105/joss.04133>.
- [5] J. Howard, H. Husain, nbdev, <https://github.com/fastai/nbdev>.
- [6] J. Yang, R. Shi, D. Wei, Z. Liu, L. Zhao, B. Ke, H. Pfister, B. Ni, MedMNIST v2-a large-scale lightweight benchmark for 2D and 3D biomedical image classification, *Sci. Data* 10 (1) (2023) 41, <http://dx.doi.org/10.1038/s41597-022-01721-8>.
- [7] M. Antonelli, A. Reinke, S. Bakas, K. Farahani, A. Kopp-Schneider, B.A. Landman, G. Litjens, B. Menze, O. Ronneberger, R.M. Summers, et al., The medical segmentation decathlon, *Nature Commun.* 13 (1) (2022) 4128, <http://dx.doi.org/10.1038/s41467-022-30695-9>.
- [8] J. Thiyyagalingam, M. Shankar, G. Fox, T. Hey, Scientific machine learning benchmarks, *Nat. Rev. Phys.* 4 (6) (2022) 413–420, <http://dx.doi.org/10.1038/s42254-022-00441-7>.
- [9] L. Aroyo, M. Lease, P. Paritosh, M. Schaekermann, Data excellence for AI: why should you care? *Interactions* 29 (2) (2022) 66–69, <http://dx.doi.org/10.1145/3517337>.
- [10] N. Sambasivan, S. Kapania, H. Highfill, D. Akrong, P. Paritosh, L.M. Aroyo, "Everyone wants to do the model work, not the data work": Data cascades in high-stakes AI, in: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–15, <http://dx.doi.org/10.1145/3411764.3445518>.
- [11] P. Rajpurkar, E. Chen, O. Banerjee, E.J. Topol, AI in health and medicine, *Nat. Med.* 28 (1) (2022) 31–38, <http://dx.doi.org/10.1038/s41591-021-01614-0>.
- [12] K. Widner, S. Virmani, J. Krause, J. Nayar, R. Tiwari, E.R. Pedersen, D. Jeji, N. Hammel, Y. Matias, G.S. Corrado, et al., Lessons learned from translating AI from development to deployment in healthcare, *Nat. Med.* (2023) 1–3, <http://dx.doi.org/10.1038/s41591-023-02293-9>.
- [13] J. Wiens, S. Saria, M. Sendak, M. Ghassemi, V.X. Liu, F. Doshi-Velez, K. Jung, K. Heller, D. Kale, M. Saeed, et al., Do no harm: a roadmap for responsible machine learning for health care, *Nat. Med.* 25 (9) (2019) 1337–1340, <http://dx.doi.org/10.1038/s41591-019-0548-6>.
- [14] O.S. Pianykh, Digital imaging and communications in medicine (DICOM): A practical introduction and survival guide, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 3–5, http://dx.doi.org/10.1007/978-3-642-10850-1_1.
- [15] S. Kaliyugarasan, A. Lundervold, et al., Multi-center CNN-based spine segmentation from T2w MRI using small amounts of data, in: *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, ISBI, IEEE, 2023, pp. 1–5, <http://dx.doi.org/10.1109/ISBI53787.2023.10230454>.
- [16] E. Hodneland, S. Kaliyugarasan, et al., Fully automatic whole-volume tumor segmentation in cervical cancer, *Cancers* 14 (10) (2022) 2372, <http://dx.doi.org/10.3390/cancers14102372>.
- [17] S. Kaliyugarasan, A. Lundervold, A. Lundervold, Brain age versus chronological age: A large scale MRI and deep learning investigation, in: *Proceedings of European Congress of Radiology-ECR 2020, 2020*, <http://dx.doi.org/10.26044/ecr2020/C-05555>.
- [18] S. Kaliyugarasan, A. Lundervold, A. Lundervold, Pulmonary nodule classification in lung cancer from 3D thoracic CT scans using fastai and MONAI, *Interact. Multim. Artif. Intell.* (2021) <http://dx.doi.org/10.9781/ijimai.2021.05.002>.
- [19] E. Hodneland, J.A. Dybvik, K.S. Wagner-Larsen, V. Šoltészová, A.Z. Munthe-Kaas, K.E. Fasmer, C. Krakstad, A. Lundervold, A.S. Lundervold, Ø. Salvesen, et al., Automated segmentation of endometrial cancer on MR images using deep learning, *Sci. Rep.* 11 (1) (2021) 179, <http://dx.doi.org/10.1038/s41598-020-80068-9>.
- [20] S. Kaliyugarasan, M. Kocinski, A. Lundervold, A. Lundervold, 2D and 3D U-Nets for skull stripping in a large and heterogeneous set of head MRI using fastai, in: *Proceedings of the NIK2020*, 2020.
- [21] S. Budd, E.C. Robinson, B. Kainz, A survey on active learning and human-in-the-loop deep learning for medical image analysis, *Med. Image Anal.* 71 (2021) 102062, <http://dx.doi.org/10.1016/j.media.2021.102062>.
- [22] A. Diaz-Pinto, S. Alle, V. Nath, Y. Tang, A. Ihsani, M. Asad, F. Pérez-García, P. Mehta, W. Li, M. Flores, et al., MONAI label: A framework for AI-assisted interactive labeling of 3D medical images, 2022, <http://dx.doi.org/10.48550/arXiv.2203.12362>, arXiv preprint [arXiv:2203.12362](https://arxiv.org/abs/2203.12362).
- [23] H. Zhang, M. Cisse, Y.N. Dauphin, D. Lopez-Paz, Mixup: Beyond empirical risk minimization, 2017, <http://dx.doi.org/10.48550/arXiv.1710.09412>, arXiv preprint [arXiv:1710.09412](https://arxiv.org/abs/1710.09412).
- [24] S. Yun, D. Han, S.J. Oh, S. Chun, J. Choe, Y. Yoo, CutMix: Regularization strategy to train strong classifiers with localizable features, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 6023–6032, <http://dx.doi.org/10.1109/ICCV.2019.00612>.