Peer Reviewed Letter openaccess

spectrai: a deep learning framework for spectral data

Conor C. Horgan* and Mads S. Bergholt*

Centre for Craniofacial and Regenerative Biology, King's College London, London SE1 9RT, UK

Contacts

C.C. Horgan: conor.horgan@kcl.ac.uk
https://orcid.org/0000-0003-0495-1615

M.S. Bergholt: mads.bergholt@kcl.ac.uk
https://orcid.org/0000-0003-3986-8942

Spectroscopy and spectral imaging have widespread applications in many scientific fields. Deep learning techniques have achieved many successes in recent years across numerous domains. However, the application of deep learning to spectral data remains a complex task due to the need for tailored augmentation routines, specific architectures for spectral data and significant memory requirements. Here we present spectrai, a comprehensive open-source deep learning framework and Python/MATLAB package designed to facilitate the training of neural networks on spectral data. spectrai provides numerous built-in spectral data pre-processing and augmentation methods, neural networks for spectral data including spectral (image) denoising, spectral (image) classification, spectral image segmentation and spectral image super-resolution. spectrai includes both command line and graphical user interface (GUI) tools designed to assist users with model and hyperparameter decisions for a wide range of applications. We demonstrate three case studies of spectral denoising, spectral segmentation and super-resolution. By providing baseline implementations of these functions, spectrai enables wider use of deep learning in spectroscopy and spectral imaging.

Keywords: deep learning, spectroscopy, spectral imaging

Introduction

Spectroscopy and spectral imaging play an important role in many fields including machine vision, remote sensing and biomedical imaging. Spectroscopy and spectral imaging techniques produce information-rich 1D data (a spectrum, λ) or 3D data (a spectral hypercube, $x \times y \times \lambda$), with significant potential for a multitude of deep learning applications. Deep learning techniques have significantly advanced many fields of imaging, achieving state-of-the-art results across a variety of tasks including classification, segmentation, super-resolution and denoising. For example, in medical imaging alone, deep learning has enabled incredible results for breast cancer prediction from mammography images,[1] virtual histological staining of tissues,[2] automated real-time colorectal polyp detection and segmentation,[3] and deep learning-based super-resolution fluorescence microscopy.[4] The growing success of deep learning combined with the necessity for GPU computation and distributed training strategies to meet the demands of deep learning neural networks have led to the development

Correspondence
C.C. Horgan: conor.horgan@kcl.ac.uk
M.S. Bergholt: mads.bergholt@kcl.ac.uk

Received: 4 August 2022
Revised: 29 August 2022
Accepted: 6 September 2022
Publication: 21 September 2022
doi: 10.1255/jsi.2022.a7
ISSN: 2040-4565

Citation
C.C. Horgan and M.S. Bergholt, "spectrai: a deep learning framework for spectral data", *J. Spectral Imaging* **11,** a7 (2022).
https://doi.org/10.1255/jsi.2022.a7

of several deep learning libraries including cuDNN,[5] Keras,[6] Theano,[7] Caffe,[8] TensorFlow[9] and PyTorch.[10] Together, these libraries have helped to accelerate progress in deep learning, providing flexible and generalisable frameworks for building neural networks and the associated training and distribution pipelines. These libraries have made it easier to build and train neural networks across a wide variety of tasks from image classification to natural language processing and predictive modelling.[11–14] Due to their flexible nature, these deep learning libraries are somewhat domain agnostic, providing general purpose tools and functions that can be composed together to achieve domain-specific tasks. However, the successful application of deep learning to different domains still requires significant expertise, with task-specific design decisions for network architecture, loss function, learning rate and schedule, as well as a multitude of hyperparameters. Achieving state-of-the-art performance while preventing overfitting and undue model bias requires a careful understanding of the influence of each of these components, with tuning for a given task and dataset. For domains with non-standard data formats and processing requirements, significant domain-specific implementation is required.[15,16] This is particularly true for the application of deep learning to spectral data.

Existing deep learning frameworks and models for computer vision are largely oriented towards RGB images. Spectral data differs substantially from RGB images, however, and poses numerous requirements for which standard neural network architectures, data augmentations and hyperparameter defaults are often unsuitable. Application of existing deep learning models to spectral datasets, thus, requires careful modification and adaptation to make training on spectral data possible and effective. For example, while spectral augmentations (e.g., spectral flipping or shifting) may be applied, standard image augmentations (e.g., brightness or contrast changes) may introduce unwanted spectral distortions. Similarly, although 2D convolutional neural networks (CNNs) may be extended to multi-channel hyperspectral images, many spectral deep learning applications employ 1D or 3D CNNs, necessitating modification of existing 2D CNN architectures or the development of novel task-specific architectures. Lastly, the large size of spectral image hypercubes poses significant memory constraints which may require modification of network architectures and training hyperparameters (e.g., batch size, patch size, scaling, data augmentation) to enable effective single- or multi-GPU training.

Despite these additional considerations posed by spectral data, deep learning has seen increasing application across multiple spectral imaging domains, achieving improved pixel-wise classification across hyperspectral imaging, mass spectrometry and infrared spectroscopy,[17–21] spectral and hyperspectral image denoising/correction[22–25] and spectral image super-resolution[22,26,27] with results superior to classical machine learning algorithms typically applied to spectral data.[28,29] However, the limited support for domain-specific tasks (e.g., spectral imaging) in current deep learning libraries requires individual research groups to implement their own data processing, augmentation and training pipelines. This results in a substantial duplication of efforts and limits effective comparison between different methods.

Recently, several efforts have aimed at developing domain-specific deep learning tools to simplify deep learning application in different fields, reduce duplications of effort and enable robust comparisons. In medical imaging, for example, data formats, sizes and processing requirements differ significantly from those of RGB images, and a number of medical image-specific deep learning platforms have been developed including DLTK,[30] NiftyNet,[15] Eisen,[31] TorchIO,[32] MONAI[33] and pymia.[34] These platforms provide a host of functions suitable for the processing, augmentation and training of medical images and have seen great successes in the medical imaging community. Frameworks such as Selene,[35] pysster[36] and Kipoi[37] provide functionality for deep learning on biological sequence and genomic data. There have also been efforts to provide a repository of neural network implementations targeted towards pixel-wise classification of hyperspectral data.[38] However, to the best of our knowledge, no framework has been developed for general purpose application to spectral data. Here, we present spectrai, an open-source, general purpose deep learning framework designed specifically for spectral data.

Methods

spectrai overview

spectrai is built on the popular PyTorch deep learning library and includes baseline implementations of several networks for different applications including spectral (image) denoising, spectral (image) classification, spectral image segmentation and spectral image super-resolution (Figure 1). In addition, an easy-to-use MATLAB GUI interfaces the Python codebase, guiding users

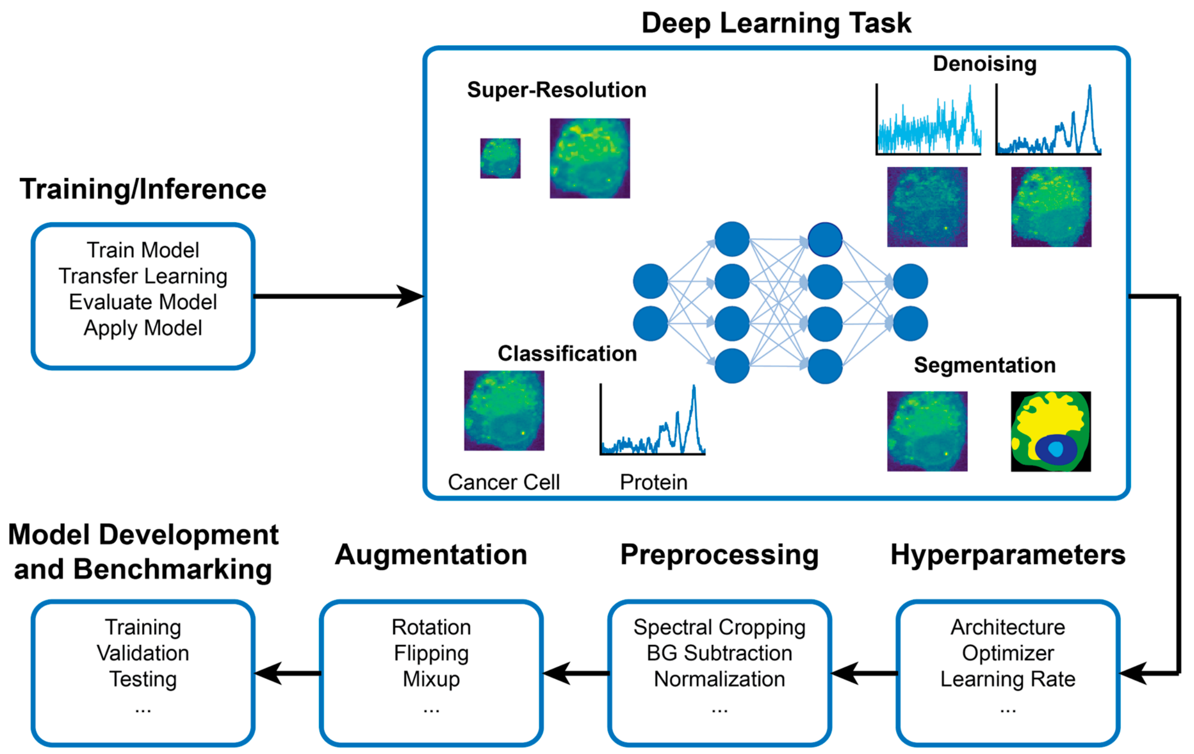

Figure 1. Overview of spectrai, a deep learning library designed for spectral data. spectrai provides the infrastructure for a variety of spectral deep learning tasks, with baseline implementations of numerous neural network architectures as well as augmentation and pre-processing methods for spectral data.

on model selection, suitable loss functions, suggested initial learning rates and default hyperparameters based on the selected task (e.g., spectral image segmentation). spectrai provides 1) an easy-to-use framework, with Python and MATLAB interfaces to guide users; and 2) baseline implementations of deep learning spectral infrastructure including data pre-processing routines, data augmentations, neural network architectures and training procedures to reduce duplication of effort across research groups and enable effective comparisons between different methods. Users can modify the existing code to implement additional data pre-processing and augmentation methods and extend the library of spectral neural network architectures. spectrai code and comprehensive documentation is available online and can be downloaded at https://github.com/conor-horgan/spectrai or installed as a Python package using pip install spectrai. spectrai is licensed under an open-source Apache 2.0 license.

spectrai's MATLAB GUI (see https://github.com/conor-horgan/spectrai) guides non-expert users through the deep learning training process outlined in Figure 1 in several ways. In the first instance, unsuitable or incompatible model architectures and hyperparameter options will be disabled given the user's selected deep learning task. On top of this, for a given deep learning task, spectrai suggests default model architectures and hyperparameter values to serve as reasonable baseline values. For example, when performing classification on a dataset with more than two classes, the spectrai MATLAB GUI will suggest a default cross-entropy loss function and disable the choice of loss functions such as binary cross-entropy and L1. Third, when performing a deep learning task with a multipurpose architecture such as a UNet, spectrai automatically modifies the final layer of the model to meet the given task requirements and data type. Example configuration files (https://github.com/conor-horgan/spectrai/tree/main/spectrai/configs) further provide default hyperparameter settings for a number of spectral deep learning tasks. spectrai's tools for visualisation of training and validation loss curves as well as model outputs during training then provide feedback to users on model performance, to help during experimentation to determine suitable model architectures and hyperparameters.

Case studies

spectrai addresses the lack of deep learning tools and frameworks designed for spectral data to lower the barriers to applying deep learning models to spectral data and is designed for both non-expert users and experienced deep learning practitioners. Here we demonstrate three example applications of spectrai for deep learning of very different types of spectral data to show the versatility: spectral image segmentation, spectral denoising and spectral image super-resolution. In each case, initial model architecture and hyperparameter selection began with the suggested spectrai default values before information provided by spectrai including visualisation of training and validation curves and intermediate model outputs were used to guide model architecture selection and tune hyperparameters.

Spectral image segmentation

Image segmentation is an important area of research with applications in fields as diverse as cell biology and remote sensing. Deep learning has demonstrated impressive results.[39,40] Together with the growing importance of spectral imaging across many fields, the development and application of neural networks for spectral image segmentation is essential.

To demonstrate the use of spectrai for spectral image segmentation, we used the recently published AeroRIT dataset.[41] This dataset consists of a single, large hyperspectral image (1973 × 3975 pixels) with reflectance data sampled every 10 nm from 400 nm to 900 nm (51 bands). (Note that the full dataset contains 372 bands between 387 nm and 1003 nm.) To train a neural network for spectral image segmentation using spectrai we extracted 64 × 64-pixel non-overlapping patches, randomly splitting these into training, validation and test sets (data split 85 : 10 : 5). Every pixel in the AeroRIT dataset images has been labelled as belonging to one of six classes [five classes (roads, buildings, vegetation, cars, water) plus one class for unspecified pixels]. Here we show segmentation results (Figure 2) achieved on the test set after training a UNet model with a batch size of 16 for 60 epochs, using the Adam optimiser[42] with a cross-entropy loss function and a constant learning rate of 1×10^{-4} [training time: 22 minutes, Titan V GPU (NVIDIA)]. These results demonstrates that spectrai can efficiently be used for image segmentation.

Spectral denoising

Spectral denoising is an important task for the processing of spectral data, aiming to achieve the removal of

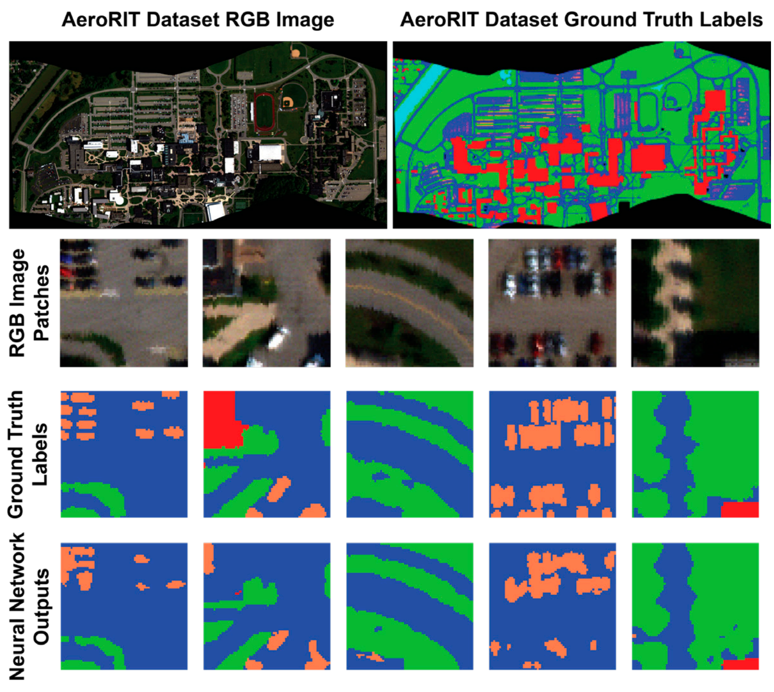

Figure 2. spectrai semantic segmentation of AeroRIT hyperspectral remote sensing dataset[41] using a UNet architecture.

unwanted noise while preserving important spectral information. Spectral denoising can, for example, be used to improve downstream data analysis and/or reduce data acquisition times. Here, we illustrate the use of spectrai for the spectral denoising of a dataset of Raman spectra of MDA-MB-231 human breast cancer cells, recently developed as part of DeepeR towards efforts to improve Raman spectral acquisition times.[22] This dataset consists of 172,312 pairs of low SNR (0.1 s spectral integration time) and high SNR (1 s spectral integration time) spectra from 11 MDA-MB-231 cells. A ResUNet model was trained for 500 epochs using the Adam optimiser with an L1 loss function and a one-cycle learning rate scheduler [training time: 26 hours, Titan V GPU (NVIDIA)], achieving results superior to Savitzky–Golay filtering that is one of the most common techniques in the biomedical Raman community for noise reduction (Figure 3).

Spectral image super-resolution

Image super-resolution is yet another important task with a multitude of applications across a wide variety of domains. Spectral image super-resolution poses additional challenges relative to super-resolution of RGB images as both spatial and spectral information must be preserved. For demonstration of spectral image super-resolution using spectrai, we focus on a dataset of intraoperative hyperspectral images of human brains for brain cancer detection.[43] The dataset consists of 36 hyperspectral images acquired intraoperatively during brain surgeries from 22 patients. The images are on average 439 × 400 pixels with 826 spectral bands between 400 nm and 1000 nm. A hyperspectral residual channel attention network (RCAN)[44] for 8× spatial super-resolution was trained on 64 × 64 pixel randomly cropped patches (bicubic downsampled 8× to produce 8 × 8 pixel inputs) from 33 of 36 images with spectral bands between 450 nm and 900 nm, with patches from one image reserved for the validation set and non-overlapping patches from a further image reserved for the test set (one hyperspectral image was removed from dataset as it was acquired from the same patient used for the test set). The RCAN model was trained with a batch size of 2 for 500 epochs using the Adam optimiser with an L1 loss function and a constant 1×10^{-4} learning rate [training time: 17 hours, Titan V GPU (NVIDIA)], achieving results superior to bicubic upsampling (Figure 4).

Conclusion

Recent advances in the application of deep learning to spectral data have shown significant potential. However, the extension of deep learning methods to spectral data is non-trivial, requiring substantial overheads relative to the applications of deep learning to RGB images. spectrai aims to minimise these overheads when applying deep learning to spectral data. By providing baseline implementations for spectral neural network architectures, spectral data augmentations and the necessary infrastructure to train neural networks on spectral data, spectrai aims to make it quicker and easier for researchers to apply deep learning to new spectral datasets, compare models and visualise results. The core spectrai platform is built using Python and PyTorch, with open-source code designed to enable experienced practitioners to extend spectrai to introduce additional neural network models, data augmentations and processing pipelines. Future development of spectrai aims to increase the range of applications, neural network architectures, and data augmentation and pre-processing methods available.

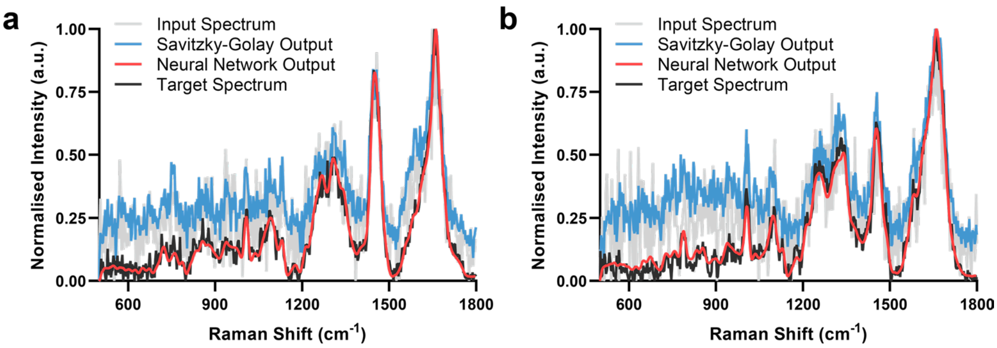

Figure 3. spectrai spectral denoising of Raman spectral dataset of MDA-MB-231 breast cancer cells.[22] Example spectra shown for a) a lipid-rich cell region and b) for the cell nucleus.

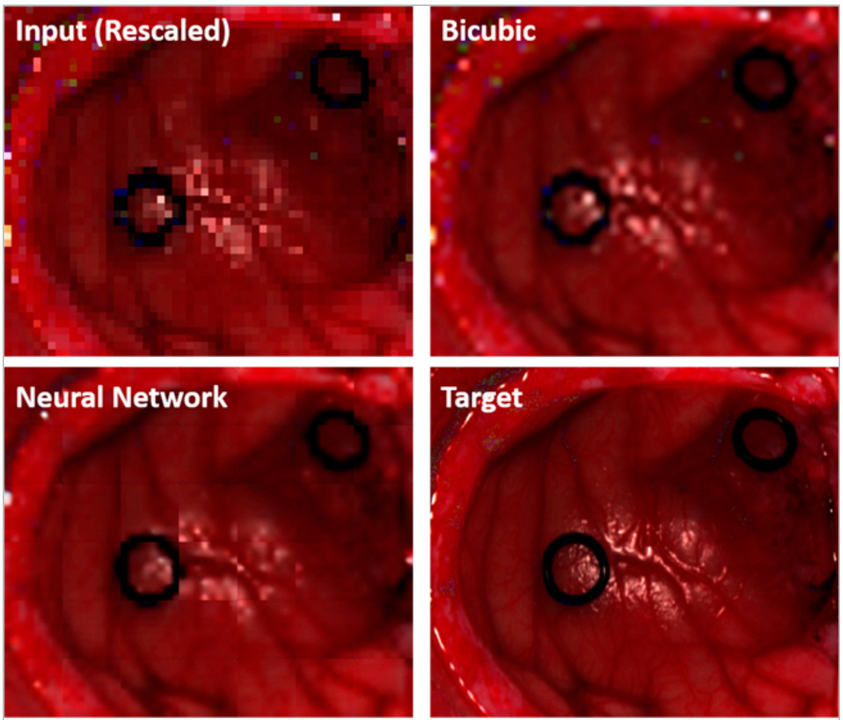

Figure 4. Spectrai 8× super-resolution of intraoperative hyperspectral brain image from the HELICoiD dataset[43] performed using a hyperspectral RCAN architecture.

Community feedback and contribution to the future development of spectrai is particularly welcomed. We hope to significantly expand the range of state-of-the-art spectral neural network architectures available to provide standard baselines for benchmarking and comparison purposes.

Acknowledgements

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 802778). The authors report no conflicts of interest.

References

1. S.M. McKinney, M. Sieniek, V. Godbole, J. Godwin, N. Antropova, H. Ashrafian, T. Back, M. Chesus, G.S. Corrado, A. Darzi, M. Etemadi, F. Garcia-Vicente, F.J. Gilbert, M. Halling-Brown, D. Hassabis, S. Jansen, A. Karthikesalingam, C.J. Kelly, D. King, J.R. Ledsam, D. Melnick, H. Mostofi, L. Peng, J.J. Reicher, B. Romera-Paredes, R. Sidebottom, M. Suleyman, D. Tse, K.C. Young, J. De Fauw and S. Shetty, "International evaluation of an AI system for breast cancer screening", *Nature* **577,** 89–94 (2020). https://doi.org/10.1038/s41586-019-1799-6
2. Y. Rivenson, H. Wang, Z. Wei, K. de Haan, Y. Zhang, Y. Wu, H. Günaydın, J.E. Zuckerman, T. Chong, A.E. Sisk, L.M. Westbrook, W.D. Wallace and A. Ozcan, "Virtual histological staining of unlabelled tissue-autofluorescence images via deep learning", *Nat. Biomed. Eng.* **3,** 466–477 (2019). https://doi.org/10.1038/s41551-019-0362-y
3. P. Wang, X. Xiao, J.R. Glissen Brown, T.M. Berzin, M. Tu, F. Xiong, X. Hu, P. Liu, Y. Song, D. Zhang, X. Yang, L. Li, J. He, X. Yi, J. Liu and X. Liu, "Development and validation of a deep-learning algorithm for the detection of polyps during colonoscopy", *Nat. Biomed. Eng.* **2,** 741–748 (2018). https://doi.org/10.1038/s41551-018-0301-3
4. H. Wang, Y. Rivenson, Y. Jin, Z. Wei, R. Gao, H. Günaydın, L.A. Bentolila, C. Kural and A. Ozcan, "Deep learning enables cross-modality super-resolution in fluorescence microscopy", *Nat. Methods* **16,** 103–110 (2019). https://doi.org/10.1038/s41592-018-0239-0

5. S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro and E. Shelhamercu, "DNN: efficient primitives for deep learning", *arXiv Preprint arXiv:1410.0759* (2014). https://doi.org/10.48550/arXiv.1410.0759
6. F. Chollet, *Keras*. https://keras.io (2015).
7. The Theano Development Team, "Theano: a Python framework for fast computation of mathematical expressions", *arXiv Preprint arXiv:1605.02688* (2016). https://doi.org/10.48550/arXiv.1605.02688
8. Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding", *Proceedings of the 22nd ACM International Conference on Multimedia* 675–678 (2014). https://doi.org/10.1145/2647868.2654889
9. M. Abadi *et al.*, "TensorFlow: A system for large-scale machine learning", *12th USENIX Symp. Oper. Syst. Des. Implement.* 265–283 (2016). https://www.usenix.org/conference/osdi16/technical-sessions/presentation/abadi
10. A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai and S. Chintala, "PyTorch: An imperative style, high-performance deep learning library", *arXiv Preprint arXiv:1912.01703* (2019). https://doi.org/10.48550/arXiv.1912.01703
11. E. Moen, D. Bannon, T. Kudo, W. Graf, M. Covert and D. Van Valen, "Deep learning for cellular image analysis", *Nat. Methods* **16,** 1233–1246 (2019). https://doi.org/10.1038/s41592-019-0403-1
12. A. Krizhevsky, I. Sutskever and G.E. Hinton, "ImageNet classification with deep convolutional neural networks", *Commun. ACM* **60(6),** 84–90 (2017). https://doi.org/10.1145/3065386
13. J. Devlin, M.-W. Chang, K. Lee and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding", *arXiv Preprint arXiv:1810.04805* (2018). https://doi.org/10.48550/arXiv.1810.04805
14. A.W. Senior, R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, T. Green, C. Qin, A. Žídek, A.W.R. Nelson, A. Bridgland, H. Penedones, S. Petersen, K. Simonyan, S. Crossan, P. Kohli, D.T. Jones, D. Silver, K. Kavukcuoglu and D. Hassabis, "Improved protein structure prediction using potentials from deep learning", *Nature* **577,** 706–710 (2020). https://doi.org/10.1038/s41586-019-1923-7
15. E. Gibson, W. Li, C. Sudre, L. Fidon, D.I. Shakir, G. Wang, Z. Eaton-Rosen, R. Gray, T. Doel, Y. Hu, T. Whyntie, P. Nachev, M. Modat, D.C. Barratt, S. Ourselin, M.J. Cardoso and T. Vercauteren, "NiftyNet: a deep-learning platform for medical imaging", *Comput. Meth. Prog. Biomed.* **158,** 113–122 (2018). https://doi.org/10.1016/j.cmpb.2018.01.025
16. F. Pérez-García, R. Sparks and S. Ourselin, "TorchIO: A Python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning", *Comput. Meth. Prog. Biomed.* **208,** 106236 (2021). https://doi.org/10.1016/j.cmpb.2021.106236
17. H. Fabelo, M. Halicek, S. Ortega, M. Shahedi, A. Szolna, J.F. Piñeiro, C. Sosa, A.J. O'Shanahan, S. Bisshopp, C. Espino, M. Márquez, M. Hernández, D. Carrera, J. Morera, G.M. Callico,R. Sarmiento and B. Fei, "Deep learning-based framework for *in vivo* identification of glioblastoma tumor using hyperspectral images of human brain", *Sensors* **19(4),** 920 (2019). https://doi.org/10.3390/s19040920
18. J. Behrmann, C. Etmann, T. Boskamp, R. Casadonte, J. Kriegsmann and P. Maaβ, "Deep learning for tumor classification in imaging mass spectrometry", *Bioinformatics* **34,** 1215–1223 (2018). https://doi.org/10.1093/bioinformatics/btx724
19. S. Berisha, M. Lotfollahi, J. Jahanipour, I. Gurcan, M. Walsh, R. Bhargava, H. Van Nguyen and D. Mayerich, "Deep learning for FTIR histology: leveraging spatial and spectral features with convolutional neural networks", *Analyst* **144,** 1642–1653 (2019). https://doi.org/10.1039/C8AN01495G
20. J.A. Fine, A.A. Rajasekar, K.P. Jethava and G. Chopra, "Spectral deep learning for prediction and prospective validation of functional groups", *Chem. Sci.* **11,** 4618–4630 (2020). https://doi.org/10.1039/C9SC06240H
21. W. Zhao and S. Du, "Spectral-spatial feature extraction for hyperspectral image classification: a dimension reduction and deep learning approach", *IEEE Trans. Geosci. Remote Sens.* **54(8),** 4544–4554 (2016). https://doi.org/10.1109/TGRS.2016.2543748
22. C.C. Horgan, M. Jensen, A. Nagelkerke, J.-P. St-Pierre, T. Vercauteren, M.M. Stevens and M.S. Bergholt, "High-throughput molecular imaging via deep learning enabled Raman spectroscopy", *Anal. Chem.* **93,** 15850–15860 (2021). https://doi.org/10.1021/acs.analchem.1c02178

23. S. Guo, T. Mayerhöfer, S. Pahlow, U. Hübner, J. Popp and T. Bocklitz, "Deep learning for 'artefact' removal in infrared spectroscopy", *Analyst* **145,** 5213–5220 (2020). https://doi.org/10.1039/D0AN00917B
24. B. Manifold, E. Thomas, A.T. Francis, A.H. Hill and D. Fu, "Denoising of stimulated Raman scattering microscopy images via deep learning", *Biomed. Opt. Express* **10(8),** 3860–3874 (2019). https://doi.org/10.1364/BOE.10.003860
25. Q. Yuan, Q. Zhang, J. Li, H. Shen and L. Zhang, "Hyperspectral image denoising employing a spatial-spectral deep residual convolutional neural network", *IEEE Trans. Geosci. Remote Sens.* **57,** 1205–1218 (2019). https://doi.org/10.1109/TGRS.2018.2865197
26. Y. Li, J. Hu, X. Zhao, W. Xie and J.J. Li, "Hyperspectral image super-resolution using deep convolutional neural network", *Neurocomputing* **266,** 29–41 (2017). https://doi.org/10.1016/j.neucom.2017.05.024
27. H. He, M. Xu, C. Zong, P. Zheng, L. Luo, L. Wang and B. Ren, "Speeding up the line-scan Raman imaging of living cells by deep convolutional neural network", *Anal. Chem.* **91(11),** 7070–7077 (2019). https://doi.org/10.1021/acs.analchem.8b05962
28. M. Hedegaard, C. Matthäus, S. Hassing, C. Krafft, M. Diem and J. Popp, "Spectral unmixing and clustering algorithms for assessment of single cells by Raman microscopic imaging", *Theor. Chem. Acc.* **130,** 1249–1260 (2011). https://doi.org/10.1007/s00214-011-0957-1
29. M. Sattlecker, N. Stone and C. Bessant, "Current trends in machine-learning methods applied to spectroscopic cancer diagnosis", *TrAC—Trends Anal. Chem.* **59,** 17–25 (2014). https://doi.org/10.1016/j.trac.2014.02.016
30. N. Pawlowski, S.I. Ktena, M.C.H. Lee, B. Kainz, D. Rueckert, B. Glocker and M. Rajchl, "DLTK: state of the art reference implementations for deep learning on medical images", *arXiv Preprint arXiv:1711.06853* (2017). https://doi.org/10.48550/arXiv.1711.06853
31. F. Mancolo, "Eisen: a python package for solid deep learning", *arXiv Preprint arXiv:2004.02747* (2020). https://doi.org/10.48550/arXiv.2004.02747
32. F. Pérez-García, R. Sparks and S. Ourselin, "TorchIO: a Python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning", *arXiv Preprint arXiv:2003.04696* (2020). https://doi.org/10.48550/arXiv.2003.04696
33. S. Aylward *et al.*, *MONAI*. (2020). https://monai.io
34. A. Jungo, O. Scheidegger, M. Reyes and F. Balsiger, "pymia: a Python package for data handling and evaluation in deep learning-based medical image analysis", *Comput. Meth. Prog. Biomed.* **198,** 105796 (2021). https://doi.org/10.1016/j.cmpb.2020.105796
35. K.M. Chen, E.M. Cofer, J. Zhou and O.G. Troyanskaya, "Selene: a PyTorch-based deep learning library for sequence data", *Nat. Methods* **16,** 315–318 (2019). https://doi.org/10.1038/s41592-019-0360-8
36. S. Budach and A. Marsico, "Pysster: classification of biological sequences by learning sequence and structure motifs with convolutional neural networks", *Bioinformatics* **34(17),** 3035–3037 (2018). https://doi.org/10.1093/bioinformatics/bty222
37. Ž. Avsec, R. Kreuzhuber, J. Israeli, N. Xu, J. Cheng, A. Shrikumar, A. Banerjee, D.S. Kim, L. Urban, A. Kundaje, O. Stegle and J. Gagneur, "Kipoi: accelerating the community exchange and reuse of predictive models for genomics", *bioRxiv* 375345 (2018). https://doi.org/10.1101/375345
38. N. Audebert, B. Le Saux and S. Lefevre, "Deep learning for classification of hyperspectral data: A comparative review", *IEEE Geosci. Remote Sens. Mag.* **7(2),** 159–173 (2019). https://doi.org/10.1109/MGRS.2019.2912563
39. O. Ronneberger, P. Fischer and T. Brox, "U-net: convolutional networks for biomedical image segmentation", *International Conference on Medical Image Computing and Computer-Aided Intervention* **9351,** 234–241 (2015). https://doi.org/10.1007/978-3-319-24574-4_28
40. V. Badrinarayanan, A. Kendall and R. Cipolla, "SegNet: a deep convolutional encoder-decoder architecture for image segmentation", *IEEE Trans. Pattern Anal. Mach. Intell.* **39(12),** 2481–2495 (2017). https://doi.org/10.1109/TPAMI.2016.2644615
41. A. Rangnekar, N. Mokashi, E.J. Ientilucci, C. Kanan and M.J. Hoffman, "AeroRIT: a new scene for hyperspectral image analysis", *IEEE Trans. Geosci. Remote Sens.* **58(11),** 8116–8124 (2020). https://doi.org/10.1109/TGRS.2020.2987199
42. D.P. Kingma and J.L. Ba, "Adam: a method for stochastic optimization", *3rd International Conference*

on Learning Representations, arXiv:1412.6980 (2015). https://doi.org/10.48550/arXiv.1412.6980

43. H. Fabelo, S. Ortega, A. Szolna, D. Bulters, J.F. Piñeiro, S. Kabwama, A. J-O'Shanahan, H. Bulstrode, S. Bisshopp, B.R. Kiran, D. Ravi, R. Lazcano, D. Madroñal, C. Sosa, C. Espino, M. Marquez, M. De La Luz Plaza, R. Camacho, D. Carrera, M. Hernández, G.M. Callicó, J.M. Molina, B. Stanciulescu, G.-Z. Yang, R. Salvador, E. Juárez, C. Sanz and R. Sarmiento, "In-vivo hyperspectral human brain image database for brain cancer detection", *IEEE Access* **7,** 39098–39116 (2019). https://doi.org/10.1109/ACCESS.2019.2904788
44. Y. Zhang, K. Li, K. Li, L. Wang, B. Zhong and Y. Fu, "Image super-resolution using very deep residual channel attention networks", *Proceedings of the European Conference on Computer Vision* 286–301 (2018). https://doi.org/10.1007/978-3-030-01234-2_18