
DLTK: State of the Art Reference Implementations for Deep Learning on Medical Images

Nick Pawlowski

Sofia Ira Ktena

Matthew C.H. Lee

Bernhard Kainz

Daniel Rueckert

Ben Glocker

Martin Rajchl

Biomedical Image Analysis Group

Department of Computing, Imperial College London, UK

np716@imperial.ac.uk

Abstract

We present DLTK, a toolkit providing baseline implementations for efficient experimentation with deep learning methods on biomedical images. It builds on top of TensorFlow and its high modularity and easy-to-use examples allow for a low-threshold access to state-of-the-art implementations for typical medical imaging problems. A comparison of DLTK’s reference implementations of popular network architectures for image segmentation demonstrates new top performance on the publicly available challenge data “Multi-Atlas Labeling Beyond the Cranial Vault”. The average test Dice similarity coefficient of 81.5 exceeds the previously best performing CNN (75.7) and the accuracy of the challenge winning method (79.0).

1 Introduction

The successful application of deep convolutional neural networks (CNNs) to the ImageNet challenge [12] by Krizhevsky et al. [7] has had a large impact on the field of computer vision. Biomedical image analysis is a particular area benefiting from latest advancements in computer vision methodology, but it only recently adapted deep learning techniques on a wider scale [9]. Adaptation of deep learning to biomedical imaging problems requires speciality operations and the lack of low-threshold access to validated reference implementations has led to slow progress. To address this, we present DLTK¹, a Deep Learning Toolkit for Medical Image Analysis. DLTK simplifies the application and experimentation with deep learning methods on medical imaging data by providing validated, high-performance reference implementations and required speciality operations.

2 Building DLTK


We believe that most of the deep learning research relevant to medical imaging will address at least one of these components: a) data reading and preprocessing; b) model definitions and network architectures; c) training and optimisation strategies; d) deployment of methods to new data. Therefore, DLTK enables a plug-and-play structure for those components based on TensorFlow’s [1] new high-level API [2]. This approach gives users the freedom to reuse existing components but also to rapidly integrate new functionality as subject of their research. We naturally interface within the TensorFlow framework and therefore benefit from its wide range of operations and community contributions. While other packages, such as NiftyNet [3], build on a predefined application structure, we emphasise an API level prioritising experimentation by exposing access to low-level operations.

¹<https://dltk.github.io/>

3 Experiments & Results

DLTK reference implementations of FCN [10] and U-Net [11] architectures using residual units [4] are tested on the publicly available challenge dataset from the MICCAI 2015 challenge “Multi-Atlas Labeling Beyond the Cranial Vault”. We test combinations of components, *a)* data reading with random and class-balanced sampling of patches, *b)* training with a Dice, cross-entropy and class-balanced cross-entropy loss for the two network architectures. All methods are trained with ADAM [6] with default parameters and tuned $\hat{\epsilon} = 10^{-5}$ to counteract loss spikes. Network inputs are patches of size 64^3 voxels. We found that the U-Net trained with cross-entropy loss and class-balanced sampling performed best out of all combinations. The comparative performance of all experiment dimensions is depicted in Figure 2 in the Appendix.

For external comparison, we submitted the best performing method to the challenge website and compare to the winning entry employing a multi-atlas approach [5] and the best performing CNN [8]. Our U-Net implementation achieves new state-of-the-art results of 81.5 in terms of DSC compared to the multi-atlas method by [5] with 79.0. The previously best performing CNN [8] achieves 75.7. Interestingly, we report a slightly lower validation performance of 78.9 compared to [8] with 79.0 which might indicate overfitting of the previous CNN [8] to the training data as we achieve better test set performance than validation performance. We further note that our architecture is far from its potential optimal performance as we did not fine-tune any of its hyperparameters or the data preprocessing and rather prefer to report out-of-the-box performance of the DLTK implementation.



(a) Axial view of an exemplary segmentation (b) 3D rendering of an exemplary segmentation

Figure 1: Prediction of the DLTK U-Net segmenting 13 organs on abdominal CT scans.

4 Conclusion

We present a new tool kit for performing deep learning experiments tailored to biomedical image analysis. DLTK offers basic baseline implementations for popular network architectures, sampling techniques and losses commonly used in medical image analysis. It further offers easy deployment including sliding-window inference invariant to input shapes. In future, we aim to extend the range of DLTK to include more additional components and latest developments and thus, enable low-threshold access to cutting-edge deep learning methods for medical imaging.

Acknowledgements

NP is supported by Microsoft Research PhD Scholarship and the EPSRC Centre for Doctoral Training in High Performance Embedded and Distributed Systems (HiPEDS, Grant Reference EP/L016796/1). MR is supported by an Imperial College Research Fellowship. We gratefully acknowledge the support of NVIDIA with the donation of one Titan X GPU for our research.

References

- [1] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Man, Rajat Monga, Sherry Moore, Derek Murray, Jon Shlens, Benoit Steiner, Ilya Sutskever, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Oriol Vinyals, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. *None*, 1(212):19, 2015.
- [2] Heng-Tze Cheng, Zakaria Haque, Lichan Hong, Mustafa Ispir, Clemens Mewald, Illia Polosukhin, Georgios Roumpos, D Sculley, Jamie Smith, David Soergel, et al. Tensorflow estimators: Managing simplicity vs. flexibility in high-level machine learning frameworks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1763–1771. ACM, 2017.
- [3] Eli Gibson, Wenqi Li, Carole Sudre, Lucas Fidon, Dzoshkun Shakir, Guotai Wang, Zach Eaton-Rosen, Robert Gray, Tom Doel, Yipeng Hu, Tom Whyntie, Parashkev Nachev, Dean C. Barratt, Sebastien Ourselin, M. Jorge Cardoso, and Tom Vercauteren. Niftynet: a deep-learning platform for medical imaging. 2017.
- [4] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *Arxiv.Org*, 7(3):171–180, 2015.
- [5] Mattias P Heinrich. Multi-organ segmentation using deeds, self-similarity context and joint fusion. 2015.
- [6] Diederik Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations*, pages 1–13, 2014.
- [7] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *Advances In Neural Information Processing Systems*, pages 1–9, 2012.
- [8] Måns Larsson, Yuhang Zhang, and Fredrik Kahl. *Robust Abdominal Organ Segmentation Using Regional Convolutional Neural Networks*, pages 41–52. Springer International Publishing, Cham, 2017.
- [9] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen A. W. M. van der Laak, Bram van Ginneken, and Clara I. Sánchez. A Survey on Deep Learning in Medical Image Analysis. (1995), 2017.
- [10] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation ppt. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
- [11] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. *MICCAI 2015*, 9351:234–241, may 2015.
- [12] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

A Experimental Comparison of Network Architectures




Figure 2: Box plot comparing the DSC scores, recall and precision for the U-Net (green) and FCN (blue). The white markers show the mean, the black line shows the median, the error bars indicate confidence intervals, and the additional black markers indicate outliers. The U-Net outperforms the FCN in almost every comparison. The FCN only outperforms the U-Net on the recall for the right adrenal gland.