

Segmentation of Traumatic Brain Injuries with Convolutional Neural Networks

Konstantinos Kamnitsas¹, Christian Ledig¹, Virginia F. J. Newcombe^{2,3}, Joanna P. Simpson², Andrew D. Kane², David K. Menon^{2,3}, Daniel Rueckert¹, Ben Glocker¹

¹ Biomedical Image Analysis Group, Imperial College London, UK

² University Division of Anaesthesia, Department of Medicine, Cambridge University, UK

³ Wolfson Brain Imaging Centre, Cambridge University, UK

Objective: In this work we investigate the capabilities of Convolutional Neural Networks (CNN) for the segmentation of brain lesions. We present a fully automatic segmentation system, the design of which was driven by the most recent advances in the field of Deep Learning, adapted for the accurate segmentation of brain lesions in multi-modal MRI scans. We demonstrate that Deep Convolutional Neural Networks with carefully designed architecture can be successfully trained even on biomedical datasets of limited size and show very promising performance on the challenging task of TBI lesion segmentation. Our network was evaluated on two datasets comprising MRI scans of patients with severe TBI, where it surpassed the performance of previously developed methods.


Methods: We have developed a 3D CNN for the segmentation of MRI brain scans, updated with the state-of-the-art techniques from the field of Computer Vision and further adaptation of its architecture to the specific task of brain lesion segmentation. Each voxel in a brain scan is classified as injured or healthy by processing the contents of an image patch around it. During training, the network automatically learns to detect meaningful data-driven features, which allow accurate classification of tissue in new unseen images without the necessity for extended prior knowledge about the task, something necessary for designing the hand engineered features of previous approaches. We propose an efficient way for utilizing greater context for the segmentation of each voxel without proportional computational increase, by processing the brain scan at multiple scales. Additionally, we show that deep neural networks can be successful without the use of max-pooling layers, which were considered necessary for other tasks such as object recognition, but the spatial invariance and sub-sampling they introduce may hinder segmentation performance. To avoid overfitting the biomedical datasets, which are usually of limited size, we employ state-of-the-art techniques from the field of Computer Vision for the regularization of our network, such as batch-normalization and dropout. Finally, by implementing fully connected layers as convolutional layers, the network is able to perform dense classification of multiple voxels simultaneously, enabling the segmentation of a volumetric brain scan in only few minutes on a standard PC with a modern graphics processing unit.

Results: The performance of our developed system was evaluated on two datasets, which comprise multi-modal MRI scans, acquired in the acute stage from patients that suffered severe traumatic brain injuries. On both datasets, our network outperforms previously developed techniques. Our system achieves accurate segmentation of TBI lesions of varying sizes, both of focal contusions and smaller diffuse injuries. Representative cases are presented in this work for qualitative assessment. The developed network can segment a multimodal

MRI brain scan in 3 minutes when run on a modern GPU, allowing its use in a variety of research and clinical settings.

Conclusion: The results achieved by our system suggest the promising capabilities of CNNs for segmentation of brain lesions. A carefully engineered network can be successfully trained on biomedical datasets of quite limited size and achieve delicate, state-of-the-art segmentation even on the challenging case of the vastly heterogeneous traumatic brain injuries. Regardless its overall promising performance, our network was found to sometimes miss lesions of particularly small size, as well as exhibit diffuse false segmentation in some problematic cases. Regardless the size of these false predictions, they are still of great clinical importance as even an axonal injury of small size may cause significant functional deficits. Performance could be improved by the use of networks of greater size, which in turn would require bigger manually annotated datasets. Further investigation of different network architectures and ways to incorporate prior knowledge about structural anatomy of the brain could also help alleviate the aforementioned problems and lead to an accurate automated system that could significantly benefit further research on the mechanisms of TBI, its clinical diagnosis and treatment.

Acknowledgements: The first author is funded by an Imperial College PhD Scholarship.



(a) FLAIR (b) Manual (c) Our CNN (d) RForest

Figure: (a) FLAIR scans of four severe TBI cases. (b) Manual lesion segmentation by expert. (c) The probability maps for the lesion class, as predicted by our system, DeepMedic, and (d) by our Random Forest. Our network provides delicate segmentation of lesions with significantly reduced false positives.