Knowledge-driven deep neural network models for brain tumour segmentation

Knowledge-driven deep neural network models for brain tumour segmentation

F Colecchia[1], J K Ruffle[1,2], G C Pombo[1], R Gray[1], H Hyare[1,2] and P Nachev[1,2]
[1]Queen Square Institute of Neurology, University College London, UK
[2]University College London Hospitals NHS Foundation Trust

E-mail: federico.colecchia@brunel.ac.uk

Abstract. Image segmentation is a computer vision task aiming to establish a probabilistic mapping between individual pixels (2D) or voxels (3D) in an input image and a set of predefined semantic categories with reference to domain-specific knowledge. When applied to medical images, e.g. Magnetic Resonance Imaging (MRI), it allows delineation between healthy and abnormal tissue. Despite challenges due to lesion morphological heterogeneity, segmentation of brain tumours has the potential to streamline otherwise time-consuming manual annotation. Whereas brain tumour segmentation has continually advanced incorporating innovative deep learning methods, heuristics normally employed by radiologists have often been neglected. The focus of nearly all tumour segmentation articles thus far on 3D isotropic research-grade scans has also led to results of unknown generalisability to hospital-quality data. In order to address these gaps, this study has coalesced modern deep learning methods and clinical-driven priors into an optimised segmentation pipeline evaluated on clinical data at a large neurology and neurosurgery tertiary centre.

1. Introduction

Image segmentation is a computer vision task that consists in assigning each pixel (2D) or voxel (3D) in an input image probabilistic meaning with reference to domain-specific knowledge. A significant body of knowledge has been developed regarding segmentation of medical images, e.g. Magnetic Resonance Imaging (MRI), that has the potential to enable automatic delineation between healthy and abnormal tissue [1-6]. Despite challenges due to lesion morphological heterogeneity, segmentation of brain tumours, typically relying on classical machine learning or deep learning approaches, has the potential to streamline otherwise time-consuming manual annotation workflows.

2. Materials and Methods

Multimodal deep neural network models were developed using data from the 2018 ‘Multimodal Brain Tumor Segmentation Challenge’ (BraTS), consisting of FLAIR, T1, Gadolinium contrast-enhanced T1 (T1CE), and T2-weighted MRI sequences of 484 patients. The data includes manually-annotated labels for tumour (non-enhancing and enhancing) and peritumoural oedema. A set of models based on the U-Net architecture were developed predicting voxels associated with each label. The architecture employed, detailed in Figure 1, relies on 3D convolutional layers and is designed to combine information from different scales [7]. Depth 4 U-Net models were used, corresponding to approximately 2 million trainable parameters.

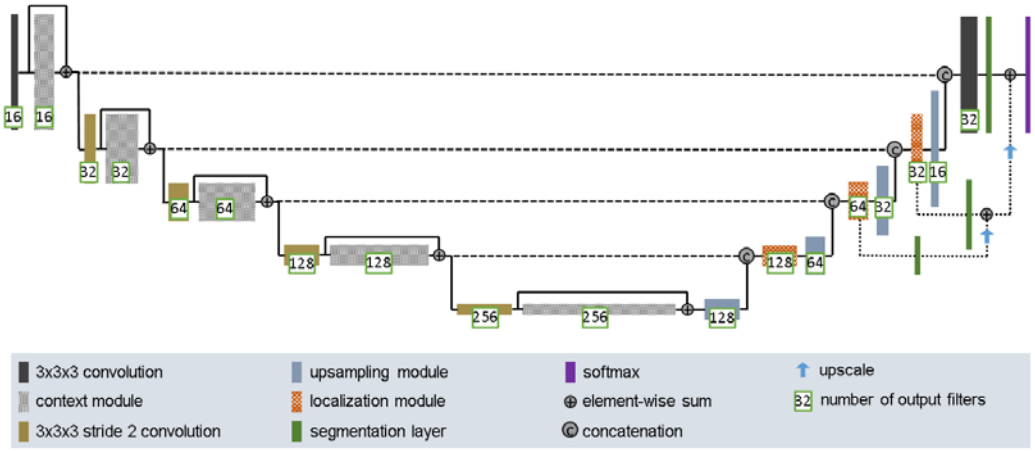

Figure 1. Network architecture [7] employed for the models used in this investigation.

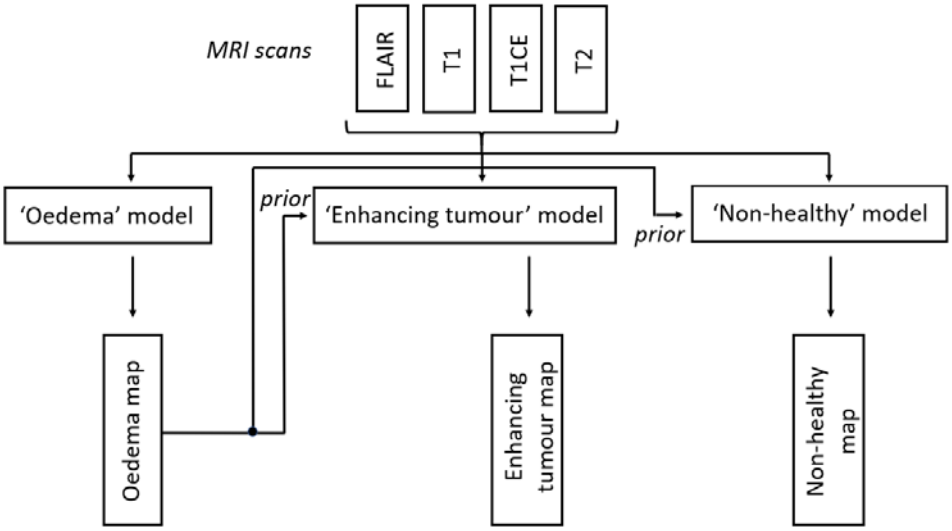

Figure 2. High-level overview of the clinically-driven approach proposed in this article.

Separate models were developed for peritumoural oedema, enhancing tumour, and non-healthy tissue (oedema, enhancing and non-enhancing tumour combined). This decision, driven by the objective of optimising each model for prediction accuracy, was also inspired by recent trends relying on a multiplicity of approaches for enhanced tumour segmentation performance [8].

The models were deployed on 493 MRI scans of patients with glioblastoma multiforme from multiple MRI scanners with different acquisition parameters, made available by the University College London Hospitals NHS Foundation Trust (UCLH). Irrevocably anonymised images acquired during routine clinical care were used with research ethics approval. The scans were super-resolved and de-noised into 1 mm^3 isotropic space using a generative model developed in-house [9]. Images were registered into common space (Montreal Neurological Institute, MNI), skull-stripped, manually reviewed, and intensity clamped using kernel density estimation.

The models were trained on BraTS data within a clinical knowledge-informed framework in a way that mimics the workflows typically employed by radiologists. A high-level overview of the approach followed for building the suite of models used for this study is given in Figure 2. For each patient, the oedema model was trained first, and tumours were subsequently segmented with the inclusion of an attention mechanism relying on a probabilistic oedema prior. The prior was embedded in the enhancing tumour model via voxel-wise multiplication between the model activation and the

oedema prediction map, in line with the approach developed in [10]. This focussed the models on enhancing tumour in the vicinity of oedema, in line with common spatial patterns. Blood vessels were vetoed in the enhancing tumour models using a statistical atlas of cerebral arteries [11], in order to reduce the impact of blood vessels which have a potentially-confounding signature on T1CE. All models were deployed on clinical data and the best-performing ones were selected by a trained radiologist based on prediction quality.

3. Results

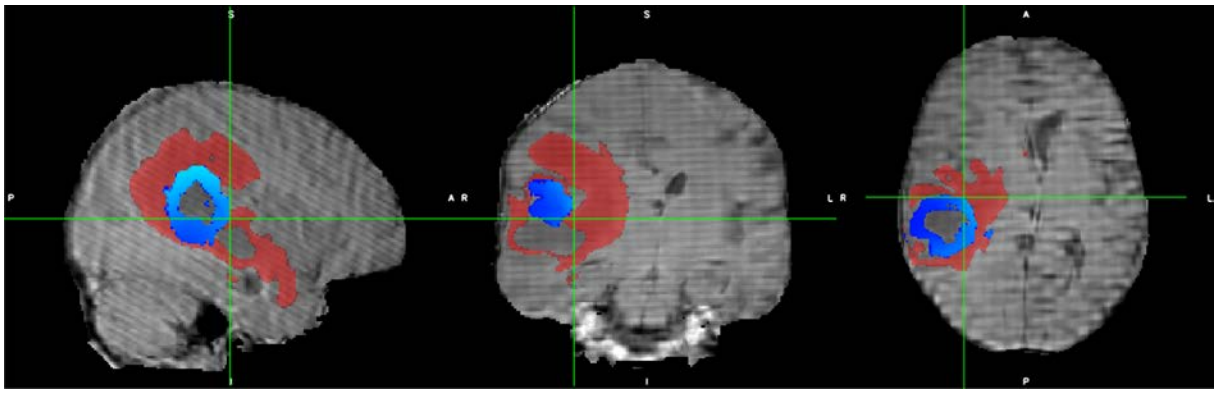

Figure 3. Sagittal, coronal, and axial sections from an artefact-degraded clinical MRI scan: peritumoural oedema and enhancing tumour predictions are shown in red and blue, respectively.

Model performance is illustrated in Figure 3, where sagittal, coronal, and axial sections from a clinical scan (FLAIR) are shown with oedema (red) and enhancing tumour (blue) predictions overlaid. Morphology is adequately predicted, including the presence of oedema protrusions into healthy tissue. Accuracy figures on the clinical dataset analysed are given in Table 1.

Table 1. Dice coefficients (average and standard deviation) on the clinical dataset analysed.

	Oedema	Enhancing tumour	Non-healthy tissue
Average	0.71	0.55	0.75
Std	0.19	0.26	0.17

4. Discussion

The novelty of the present approach to brain tumour segmentation is two-fold. Whereas several recent advancements that have been proposed in relation to brain tumour segmentation rely on novel deep learning methods [12-15], the possibility of taking advantage of those heuristics that are normally employed by radiologists when reviewing clinical scans has received comparatively little attention. Secondly, the majority of tumour segmentation studies have so far focussed on 3D isotropic research-grade scans as opposed to hospital-quality data. This has in turn resulted in an unknown degree of generalisability of the results to data collected within standard clinical workflows. This article documents a preliminary investigation towards addressing both gaps.

5. Conclusions and Outlook

Results suggest the feasibility of combining modern deep learning methods and clinical-driven priors into optimised segmentation pipelines trained on research-grade data and yet capable of performing on a par with state-of-the-art techniques on hospital-quality scans. Ongoing work focusses on non-

enhancing tumour segmentation, which is a harder task when compared to peritumoural oedema and enhancing tumour due to heavier reliance on morphological features for classification. This study is a stepping stone towards streamlining the processing of large volumes of hospital-grade patient imaging data, as part of a broader effort towards numerical modelling of brain tumour tissue and anomaly detection in the brain.

6. Acknowledgements

Funding is acknowledged from the Wellcome Trust and from the UCLH NIHR Biomedical Research Centre.

7. References

[1] Balafar M A *et al* 2010 *Artificial Intelligence Review* **33** 261-74
[2] Dupont C *et al* 2016 *IRBM* **37** 131-43
[3] Kermi A *et al* 2018 *IET Image Processing* **12** 1964-71
[4] Zhou C *et al* 2016 *Br. J. Radiol.* **89** 20151054
[5] Pinto A *et al* 2018 *Pattern Recognition* **82** 105-17
[6] Soltaninejad M *et al* 2018 *Computer Methods and Programs in Biomedicine* **157** 69-84
[7] Isensee F *et al* 2018 *arXiv 1802.10508*
[8] Kamnitsas K *et al* 2018 In: Crimi A, Bakas S, Kuijf H, Menze B, Reyes M (eds) Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries. BrainLes 2017. Lecture Notes in Computer Science 10670 Springer Cham
[9] Brudfors M *et al*. 2019 *arXiv:1909.01140*
[10] Petit O *et al* 2019 *Medical Imaging with Deep Learning* London UK
[11] Mouches P and Forkert N D 2019 *Sci Data* **6** 29
[12] Isın A *et al* 2016 *Procedia Computer Science* **102** 317-24
[13] Pereira S *et al* 2016 *IEEE Transactions on Medical Imaging* **35** 1240-51
[14] Dong H *et al* 2017 *Annual Conference on Medical Image Understanding and Analysis* Edinburgh UK
[15] Akkus Z *et al* 2017 *Journal of Digital Imaging* **30** 449-59