



Segmenting Brain Tumors from MRI Using Cascaded Multi-modal U-Nets

Michał Marcinkiewicz¹, Jakub Nalepa^{1,2}(✉), Pablo Ribalta Lorenzo²,
Wojciech Dudzik^{1,2}, and Grzegorz Mrukwa^{1,2}

¹ Future Processing, Gliwice, Poland
jnalepa@ieee.org

² Silesian University of Technology, Gliwice, Poland

Abstract. Gliomas are the most common primary brain tumors, and their accurate manual delineation is a time-consuming and very user-dependent process. Therefore, developing automated techniques for reproducible detection and segmentation of brain tumors from magnetic resonance imaging is a vital research topic. In this paper, we present a deep learning-powered approach for brain tumor segmentation which exploits multiple magnetic-resonance modalities and processes them in two cascaded stages. In both stages, we use multi-modal fully-convolutional neural nets inspired by U-Nets. The first stage detects regions of interests, whereas the second stage performs the multi-class classification. Our experimental study, performed over the newest release of the BraTS dataset (BraTS 2018) showed that our method delivers accurate brain-tumor delineation and offers very fast processing—the total time required to segment one study using our approach amounts to around 18 s.

Keywords: Brain tumor · Segmentation · Deep learning · CNN

1 Introduction

Gliomas are the most common primary brain tumors in humans. They are characterized by different levels of aggressiveness which directly influences prognosis. Due to the gliomas' heterogeneity (in terms of shape and appearance) manifested in multi-modal magnetic resonance imaging (MRI), their accurate delineation is an important yet challenging medical image analysis task. However, manual segmentation of such brain tumors is very time-consuming and prone to human errors. It also lacks reproducibility which adversely affects the effectiveness of patient's monitoring, and can ultimately lead to inefficient treatment.

Therefore, automatic brain tumor *detection* (i.e., which pixels in an input image are tumorous) and *classification* (what is a type of a tumor and/or which

We applied the *sequence-determines-credit* approach for the sequence of authors, and M. Marcinkiewicz, J. Nalepa are joint first authors.

part of a tumor, e.g., edema, non-enhancing solid core, or enhancing structures a given pixel belongs to; see examples in Fig. 1) from MRI are vital research topics in the pattern recognition and medical image analysis fields. A very wide practical applicability of such techniques encompasses computer-aided diagnosis, prognosis, staging, and monitoring of a patient. In this paper, we propose a deep learning technique to detect and segment gliomas from MRI in a cascaded processing pipeline. These gliomas are further segmented into the enhancing tumor (ET), tumor core (TC), and the whole tumor (WT).




Fig. 1. Different parts of a brain tumor (*detection* is presented in the second row—green parts show the agreement with a human reader) *segmented* using the proposed method (third row) alongside original images (first row): red—peritumoral edema, yellow—necrotic and non-enhancing tumor core, green—GD-enhancing tumor. (Color figure online)

1.1 Contribution

The contribution of this work is multi-fold:

- We propose a deep learning technique for detection and segmentation of brain tumors from MRI. Our deep neural networks (DNNs) are inspired by the U-Nets [28] with considerable changes to the architecture, and they are cascaded—the first DNN performs detection, whereas the second segments a tumor into the enhancing tumor, tumor core, and the whole tumor (Fig. 1).

- To improve generalization capabilities of our segmentation models, we build an ensemble of DNNs trained over different folds of a training set, and average the responses of the base classifiers.
- We show that our approach can be seamlessly applied to the multi-modal MRI analysis, and allows for introducing separate processing pathways for each modality.
- We validate our techniques over the newest release of the Brain Tumor Segmentation dataset (BraTS 2018), and show that they provide high-quality detection and segmentation, and offer instant segmentation.

1.2 Paper Structure

This paper is organized as follows. In Sect. 2, we discuss the current state of the art in brain-tumor delineation. The proposed deep learning-based techniques are presented in Sect. 3. The results of our experiments are analyzed in Sect. 4. Section 5 concludes the paper and highlights the directions of our future work.

2 Related Literature

Approaches for automated brain-tumor delineation can be divided into *atlas-based*, *unsupervised*, *supervised*, and *hybrid* techniques (Fig. 2). In the *atlas-based* algorithms, manually segmented images (referred to as *atlases*) are used to segment incoming (previously unseen) scans [25]. These atlases model the anatomical variability of the brain tissue [22]. Atlas images are extrapolated to new frames by warping and applying non-rigid registration techniques. An important drawback of such techniques is the necessity of creating large (and representative) annotated reference sets. It is time-consuming and error prone in practice, and may lead to atlases which cannot be applied to other tumors because they do not encompass certain types of brain tumors [1, 6].




Fig. 2. Automated delineation of brain tumors from MRI—a taxonomy.

Unsupervised algorithms search for hidden structures within unlabeled data [9, 19]. In various meta-heuristic approaches, e.g., in evolutionary algorithms [33], brain segmentation is understood as an optimization problem, in which pixels (or voxels) of similar characteristics are searched. It is tackled in a biologically-inspired manner, in which a population of candidate solutions (being the pixel or voxel labels) evolves in time [7]. Other unsupervised algorithms

encompass clustering-based techniques [14, 29, 35], and Gaussian modeling [30]. In *supervised* techniques, manually segmented image sets are utilized to train a model. Such algorithms include, among others, decision forests [10, 39], conditional random fields [36], support vector machines [17], and extremely randomized trees [24].

Deep neural networks, which established the state of the art in a plethora of image-processing and image-recognition tasks, have been successful in segmentation of different kinds of brain tissue as well [12, 16, 21] (they very often require computationally intensive data pre-processing). Holistically nested neural nets for MRI were introduced in [38]. White matter was segmented in [11], and convolutional neural networks were applied to segment tumors in [13]. Interestingly, the winning BraTS’17 algorithm used deep neural nets ensembles [15]. However, the authors reported neither training nor inference times of their algorithm which may prevent from using it in clinical practice. *Hybrid* algorithms couple together methods from other categories [26, 31, 37].

We address the aforementioned issues and propose a deep learning algorithm for automated brain tumor segmentation which exploits a new multi-modal fully-convolutional neural network based on U-Nets. The experimental evidence (presented in Sect. 4) obtained over the newest release of the BraTS dataset (BraTS 2018) shows that it can effectively deal with multi-class classification, and it delivers high-quality tumor segmentation in real time.

3 Methods

In this work, we propose an algorithm which utilizes cascaded U-Net-based deep neural networks for detecting and segmenting brain tumors. Our approach for this task is driven by an assumption that the most salient features of a lesion are not contained in a single image modality.

There are multiple ways to exploit all the modalities in deep learning-based engines. One way is to store three (or four) modalities as channels of a single image, like RGB (RGBA), and process it as a standard color image. Although this approach has a significant downside—only the first layer (which extracts the most basic features) has access to the modalities as separate inputs, it can be successfully applied to easier computer-vision and image-processing tasks. Consecutive layers in the network process the outputs of the previous layers—a mix of features from all the modalities.

Hu and Xia processed each modality separately, and merged them at the very end of the processing chain to produce the final segmentation mask, to fully benefit from information manifested in each modality [8]. In this work, we combine both techniques—we use merged modalities for brain-tumor detection, and separate processing pathways for further segmentation of a tumor.

3.1 Detection of Brain Tumors from MRI

The first stage of our image analysis approach involves taking the whole image as an input (i.e., different modalities are stacked together as the channels of an image), and producing a binary mask of the region of interest (therefore, it performs *detection* of a tumor). This binary mask is used to select the voxels of all modalities from the original images (rendering remaining pixels as background). This region is passed to the segmentation unit by the U-Net in the second stage for the final multi-class segmentation.

The architecture of our DNN used for detection is visualized in Fig. 3 (note that we present multiple processing pathways which are exploited in segmentation; for detection, only one pathway is used, and the sigmoid function was applied as the non-linearity). The DNN prediction is binarized using a threshold \mathbf{T}_b . The binary mask is post-processed using the 3D connected components analysis—the size of connected components is calculated, and the one with the largest volumes remains. If the next (second) largest connected component is at least \mathbf{T}_{cc} (in %) of the volume of the largest, it is kept as well. The binary masks resulting from the first stage are used to produce input to the second stage. More details on the architecture of our deep network itself are presented in the following subsection.

3.2 Segmentation of Detected Brain Tumors

Our DNN for brain tumor *segmentation* separates processing pathways and merges them at the very bottom of the network, where the feature space is compacted the most, and at each bridged connection (Fig. 3). By doing that, we assure that the low- and high-level features are extracted separately for all modalities in the contracting path. Those features can “interact” with each other in the expanding path, producing high-quality segmentations. Our preliminary experiments showed that the pre-contrast T1 modality carries the smallest amount of information, therefore in order to reduce the amount of segmentation time and resources (to make our method easily applicable in a real-life clinical setting), we did not use that modality in our pipeline. However, the proposed U-Net-based architecture is fairly flexible and allows for using any number of input modalities.

Our models are based on a well-known U-Net [28], with considerable changes to the architecture. First, there are separate pathways for each modality, effectively making three contracting paths. In the original architecture the number of filters was doubled at each down-block, whereas in our model it is constant everywhere, except in the very bottom part of the network (where the concatenation and merging of the paths takes place) where it is doubled. The down-block in our model consists of three convolutional layers (48 filters of the size 3×3 each, with stride 1). The second alteration to the original U-Net are the bridged connections, which join (concatenate) activations from each pathway of the contracting paths with their corresponding activations from the expanding path, where they become merged. This procedure allows the DNN to extract high-level features while preserving the context stored earlier. The expanding path is

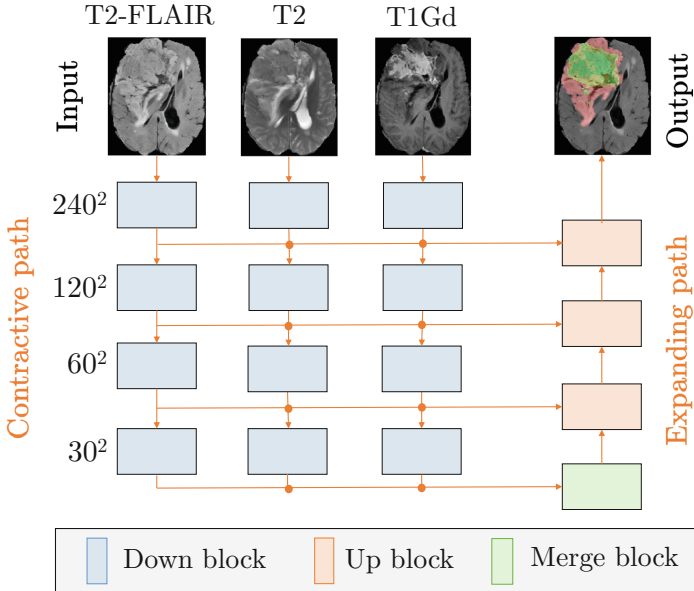


Fig. 3. The proposed deep neural network architecture. Three separate pathways (e.g., for FLAIR, T1c, and T2) are shown as a part of the contractive path. At each level (each set of down blocks) the output is concatenated and sent to a corresponding up block. At the bottom, there is a merging block, where all the features are merged before entering the expanding path. The output layer is a 1×1 convolution with one filter for the first stage (detection), and three filters for the second stage (segmentation).

standard—each up-block doubles the size of an activation map by the upsampling procedure, which is followed by two convolutions (48 filters of 3×3 size each, with stride 1). In the last layer, there is a 1×1 convolution with 1 filter in the detection, and 3 filters in the multi-class classification stages, respectively.

The output of the second stage is an activation map of the size $I_w \times I_h \times 3$, where the last dimension represents the number of classes, and I_w and I_h are the image width and height, respectively. The activation is then passed through a softmax operation, which performs the final multi-class classification.

4 Experimental Validation

4.1 Data

The Brain Tumor Segmentation (BraTS) dataset [2–5, 20] encompasses MRI-DCE data of 285 patients with diagnosed gliomas—210 high-grade glioblastomas (HGG), and 75 low-grade gliomas (LGG). Each study was manually annotated by one to four experienced readers. The data comes in four co-registered modalities: native pre-contrast (T1), post-contrast T1-weighted (T1c), T2-weighted

(T2), and T2 Fluid Attenuated Inversion Recovery (FLAIR). All the pixels have one of four labels attached: healthy tissue, Gd-enhancing tumor (ET), peritumoral edema (ED), the necrotic and non-enhancing tumor core (NCR/NET).

The data was acquired with different clinical protocols, various scanners, and at 19 institutions, therefore the pixel intensity distribution may vary significantly. The studies were interpolated to the same shape ($240 \times 240 \times 155$, hence 155 images of 240×240 size, with voxel size 1 mm^3), and they were pre-processed (skull-stripping was applied). Overall, there are 285 patients in the training set \mathbf{T} (210 HGG, 75 LGG), 66 patients in the validation set \mathbf{V} (without ground-truth data provided by the BraTS 2018 organizers), and 191 in the test set Ψ (unseen data used for the final verification of the trained models).

4.2 Experimental Setup

The DNN models were implemented using `Python3` with the `Keras` library over `CUDA 9.0` and `CuDNN 5.1`. The experiments were run on a machine equipped with an Intel i7-6850K (15 MB Cache, 3.80 GHz) CPU with 32 GB RAM and `NVIDIA GTX Titan X GPU` with 12 GB VRAM. The training metric was the DICE score for both stages (detection and segmentation), which is calculated as

$$\text{DICE}(A, B) = \frac{2 \cdot |A \cap B|}{|A| + |B|}, \quad (1)$$

where A and B are two segmentations, i.e., manual and automated. DICE ranges from zero to one (one is the perfect score). The optimizer was Nadam (Adam with Nesterov momentum) with the initial learning rate 10^{-5} , and the optimizer parameters: $\beta_1 = 0.9$, $\beta_2 = 0.999$. The training ran until DICE over the validation set did not increase by at least 0.002 in 10 epochs. The training time for one epoch is around 10 min (similar for both stages). The networks converges in around 20–30 epochs (the complete training for each fold takes 7–8 h). For detection, we used the manually-tuned thresholds: $\mathbf{T}_b = 0.5$, and $\mathbf{T}_{cc} = 20\%$.

Both networks are relatively small, which directly translates to the low computational requirements during inference—one volume can be processed and classified end-to-end within around 5 s. To exploit the training set completely, and still be able to use validation subset to avoid over-fitting, the final prediction was performed with an ensemble of five models trained on different folds of the training set (we followed the 5-fold cross-validation setting over the training set). Using an ensemble of five models (and averaging their outputs to elaborate the final prediction) was shown to improve the performance, while extending the inference time to around 18 s per full volume.

4.3 Experimental Results

In Table 1, we gather the results (DICE) obtained over the training and validation BraTS 2018 datasets (in the 5-fold cross-validation setting). The whole tumor class represents the performance of the first stage of our classification

system (evaluated on all the classes merged into one—exactly as the first stage model is trained). Here, we report the average DICE for 5 non-overlapping folds of the training set \mathbf{T} , and the final DICE for validation \mathbf{V} , and test Ψ sets obtained using the ensembles of 5 base deep classifiers learned using training examples belonging to different folds. Note that the ground-truth data (pixel labels) for \mathbf{V} and Ψ were not known to the participants during the BraTS 2018 challenge, hence they could not be exploited to improve the models.

Table 1. Segmentation performance (DICE) over the BraTS 2018 validation set obtained using our DNNs trained with T1c, T2, and FLAIR images. The scores are presented for whole tumor (WT), tumor core (TC), and enhancing tumor (ET) classes. For the training set, we report the average across 5 non-overlapping folds, whereas for the validation set—the results reported automatically by the BraTS competition server (for validation, we used an ensemble of 5 DNNs trained over different folds).

Dataset	Label	DICE	Sensitivity	Specificity
Training	ET	0.7365	0.8483	0.9981
	WT	0.9268	0.9239	0.9956
	TC	0.8779	0.8891	0.9973
Validation	ET	0.7519	0.8373	0.9972
	WT	0.8980	0.9096	0.9935
	TC	0.8118	0.8142	0.9974

The results show that an ensemble of DNNs manifests fairly good generalization capabilities over the unseen data, and it consistently obtains high-quality classification. Interestingly, we did *not* use any data augmentation techniques in our approach (which can be perceived as an implicit regularization), and even without increasing the size and heterogeneity of the training data, the ensembles were able to accurately delineate brain tumors in unseen scans. It also indicates that data augmentation could potentially further improve the capabilities (both detection and segmentation) of our deep models by providing a large number of artificially created (but visually plausible and anatomically correct) training examples generated using the original \mathbf{T} .

In Table 2, we gather the results obtained over the unseen test set Ψ —we report not only DICE, but also the Hausdorff distance (HD) given as

$$\text{HD}(A, B) = \max(h(A, B), h(B, A)), \quad (2)$$

where $h(A, B)$ is the directed Hausdorff distance:

$$h(A, B) = \max_{a \in A} \min_{b \in B} \|a - b\|, \quad (3)$$

and $\|\cdot\|$ is a norm operation (e.g., Euclidean distance) [32]. It can be noted that this metric is quite sensitive to outliers (the lower HD, the higher quality

Table 2. The results reported for the unseen test set Ψ show that our models can generalize fairly well over the unseen data (however, the results are worse when compared to the validation set). We report DICE alongside the Hausdorff distance (HD).

Measure	DICE (ET)	DICE (WT)	DICE (TC)	HD (ET)	HD (WT)	HD (TC)
Mean	0.6493	0.8590	0.7342	7.3117	9.5128	11.5729
Std dev.	0.3021	0.1272	0.2604	12.8380	16.1665	15.1306
Median	0.7770	0.8983	0.8332	3.0000	4.1231	7.3485
25 quantile	0.6023	0.8376	0.6877	2.0000	2.4495	3.4783
75 quantile	0.8423	0.9299	0.9010	5.0500	7.8717	13.9277

segmentation we have in terms of contour similarity). The results show that our deep-network ensemble can generalize quite well over the unseen data, however the DICE values are slightly lowered when compared to the validation set. We can attribute it to the heterogeneity of the testing data (as mentioned earlier, we did not apply any data augmentation to increase the representativeness of the training set). Interestingly, the whole-tumor segmentation remained at the very same level (see DICE in Table 2), and our method delivered high-quality whole-tumor delineation (we can observe the highest decrease of accuracy for the enhancing part of a tumor, and it amounts to more than 0.08 DICE on average). It also leads us to the conclusion that for tumor segmentation (differentiating between different parts of a lesion), the deep models require larger and more diverse sets (perhaps due to subtle tissue differences which cannot be learnt from a limited number of brain-tumor examples) and potentially better regularization.

5 Conclusions

In this paper, we presented an approach for effective detection and segmentation (into different parts of a tumor) of brain lesions from magnetic resonance images which exploits cascaded multi-modal fully-convolutional neural networks inspired by the U-Net architecture. The first deep network in our pipeline performs tumor detection, whereas the second—multi-class tumor segmentation. We cross-validated the proposed technique (in the 5-fold cross-validation setting) over the newest release of the BraTS dataset (BraTS 2018), and the experimental results showed that:

- Our cascaded multi-modal U-Nets deliver accurate segmentation, and ensembling the models (and averaging the response of base classifiers) trained across separate folds allows us to build the final model which generalizes well over the unseen testing data.
- We showed that our networks can be trained fairly fast (7–8 h using 1 GPU), and deliver real-time inference (around 18 s per volume).
- We showed that our models can be seamlessly applied to both two- and multi-class classification (i.e., tumor detection and segmentation, respectively).

Our current research is focused on applying our techniques to different organs and modalities (e.g., lung PET/CT imaging [23]), and developing data augmentation approaches for medical images. Such algorithms (which ideally generate artificial but visually plausible and realistic images) can be perceived as implicit regularizers which help improve the performance of models over the unseen data by introducing new examples into a training set [18, 27, 34].

Acknowledgments. This research was supported by the National Centre for Research and Development under the Innomed Research and Development Grant No. POIR.01.02.00-00-0030/15.

References

1. Aljabar, P., Heckemann, R., Hammers, A., Hajnal, J., Rueckert, D.: Multi-atlas based segmentation of brain images: atlas selection and its effect on accuracy. *NeuroImage* **46**(3), 726–738 (2009)
2. Bakas, S., et al.: Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. *Nat. Sci. Data* **4**, 1–13 (2017). <https://doi.org/10.1038/sdata.2017.117>
3. Bakas, S., et al.: Segmentation labels and radiomic features for the pre-operative scans of the TCGA-GBM collection, the Cancer Imaging Archive (2017). <https://doi.org/10.7937/K9/TCIA.2017.KLXWJJ1Q>
4. Bakas, S., et al.: Segmentation labels and radiomic features for the pre-operative scans of the TCGA-LGG collection, the Cancer Imaging Archive (2017). <https://doi.org/10.7937/K9/TCIA.2017.GJQ7R0EF>
5. Bakas, S., Reyes, M., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. *CoRR abs/1811.02629* (2018). <http://arxiv.org/abs/1811.02629>
6. Bauer, S., Seiler, C., Bardyn, T., Buechler, P., Reyes, M.: Atlas-based segmentation of brain tumor images using a markov random field-based tumor growth model and non-rigid registration. In: *Proceedings of IEEE EMBC*, pp. 4080–4083 (2010). <https://doi.org/10.1109/IEMBS.2010.5627302>
7. Chander, A., Chatterjee, A., Siarry, P.: A new social and momentum component adaptive PSO algorithm for image segmentation. *Expert Syst. Appl.* **38**(5), 4998–5004 (2011)
8. Dong, H., Yang, G., Liu, F., Mo, Y., Guo, Y.: Automatic brain tumor detection and segmentation using U-Net based fully convolutional networks. *CoRR abs/1705.03820* (2017). <http://arxiv.org/abs/1705.03820>
9. Fan, X., Yang, J., Zheng, Y., Cheng, L., Zhu, Y.: A novel unsupervised segmentation method for MR brain images based on fuzzy methods. In: Liu, Y., Jiang, T., Zhang, C. (eds.) *CVBIA 2005*. LNCS, vol. 3765, pp. 160–169. Springer, Heidelberg (2005). https://doi.org/10.1007/11569541_17
10. Geremia, E., Clatz, O., Menze, B.H., Konukoglu, E., Criminisi, A., Ayache, N.: Spatial decision forests for MS lesion segmentation in multi-channel magnetic resonance images. *NeuroImage* **57**(2), 378–390 (2011)
11. Ghafoorian, M., et al.: Location sensitive deep convolutional neural networks for segmentation of white matter hyperintensities. *CoRR abs/1610.04834* (2016). <http://arxiv.org/abs/1610.04834>

12. Ghafoorian, M., et al.: Transfer learning for domain adaptation in MRI: application in brain lesion segmentation. In: Proceedings of MICCAI, pp. 516–524 (2017)
13. Havaei, M., Dutil, F., Pal, C., Larochelle, H., Jodoin, P.-M.: A convolutional neural network approach to brain tumor segmentation. In: Crimi, A., Menze, B., Maier, O., Reyes, M., Handels, H. (eds.) BrainLes 2015. LNCS, vol. 9556, pp. 195–208. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-30858-6_17
14. Ji, S., Wei, B., Yu, Z., Yang, G., Yin, Y.: A new multistage medical segmentation method based on superpixel and fuzzy clustering. *Comp. Math. Meth. Med.* **2014**, 747549:1–747549:13 (2014)
15. Kamnitsas, K., et al.: Ensembles of multiple models and architectures for robust brain tumour segmentation. In: Crimi, A., Bakas, S., Kuijff, H., Menze, B., Reyes, M. (eds.) BrainLes 2017. LNCS, vol. 10670, pp. 450–462. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-75238-9_38
16. Korfiatis, P., Kline, T.L., Erickson, B.J.: Automated segmentation of hyperintense regions in FLAIR MRI using deep learning. *Tomogr. J. Imaging Res.* **2**(4), 334–340 (2016). <https://doi.org/10.18383/j.tom.2016.00166>
17. Ladgham, A., Torkhani, G., Sakly, A., Mtibaa, A.: Modified support vector machines for MR brain images recognition. In: Proceedings of CoDIT, pp. 032–035 (2013). <https://doi.org/10.1109/CoDIT.2013.6689515>
18. Lorenzo, P.R., Nalepa, J.: Memetic evolution of deep neural networks. In: Proceedings of the Genetic and Evolutionary Computation Conference GECCO 2018, pp. 505–512. ACM, New York (2018)
19. Mei, P.A., de Carvalho Carneiro, C., Fraser, S.J., Min, L.L., Reis, F.: Analysis of neoplastic lesions in magnetic resonance imaging using self-organizing maps. *J. Neurol. Sci.* **359**(1–2), 78–83 (2015)
20. Menze, B.H., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Trans. Med. Imaging* **34**(10), 1993–2024 (2015). <https://doi.org/10.1109/TMI.2014.2377694>
21. Moeskops, P., Viergever, M.A., Mendrik, A.M., de Vries, L.S., Benders, M.J.N.L., Isgum, I.: Automatic segmentation of MR brain images with a convolutional neural network. *IEEE Trans. Med. Imaging* **35**(5), 1252–1261 (2016). <https://doi.org/10.1109/TMI.2016.2548501>
22. Park, M.T.M., et al.: Derivation of high-resolution MRI atlases of the human cerebellum at 3T and segmentation using multiple automatically generated templates. *NeuroImage* **95**, 217–231 (2014)
23. Pawelczyk, K., et al.: Towards detecting high-uptake lesions from lung CT scans using deep learning. In: Battiato, S., Gallo, G., Schettini, R., Stanco, F. (eds.) ICIAP 2017. LNCS, vol. 10485, pp. 310–320. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68548-9_29
24. Pinto, A., Pereira, S., Correia, H., Oliveira, J., Rasteiro, D.M.L.D., Silva, C.A.: Brain tumour segmentation based on extremely rand. forest with high-level features. In: Proceedings of IEEE EMBC, pp. 3037–3040 (2015). <https://doi.org/10.1109/EMBC.2015.7319032>
25. Pipitone, J., et al.: Multi-atlas segmentation of the whole hippocampus and subfields using multiple automatically generated templates. *NeuroImage* **101**, 494–512 (2014)
26. Rajendran, A., Dhanasekaran, R.: Fuzzy clustering and deformable model for tumor segmentation on MRI brain image: a combined approach. *Procedia Eng.* **30**, 327–333 (2012). <https://doi.org/10.1016/j.proeng.2012.01.868>
27. Rezaei, M., et al.: Conditional adversarial network for semantic segmentation of brain tumor. *CoRR abs/1708.05227*, pp. 1–10 (2017)

28. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. CoRR abs/1505.04597 (2015)
29. Saha, S., Bandyopadhyay, S.: MRI brain image segmentation by fuzzy symmetry based genetic clustering technique. In: Proceedings of IEEE CEC, pp. 4417–4424 (2007)
30. Simi, V., Joseph, J.: Segmentation of glioblastoma multiforme from MR images - a comprehensive review. Egypt. J. Radiol. Nucl. Med. **46**(4), 1105–1110 (2015)
31. Soltaninejad, M., et al.: Automated brain tumour detection and segmentation using superpixel-based extremely randomized trees in FLAIR MRI. Int. J. of Comp. Assist. Radiol. Surg. **12**(2), 183–203 (2017)
32. Taha, A.A., Hanbury, A.: Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool. BMC Med. Imaging **15**(1), 29 (2015)
33. Taherdangkoo, M., Bagheri, M.H., Yazdi, M., Andriole, K.P.: An effective method for segmentation of MR brain images using the ant colony optimization algorithm. J. Digit. Imaging **26**(6), 1116–1123 (2013)
34. Varghese, A., Mohammed, S., Sai, C., Ganapathy, K.: Generative adversarial networks for brain lesion detection. In: Proceedings of SPIE, vol. 10133, p. 10133 (2017)
35. Verma, N., Cowperthwaite, M.C., Markey, M.K.: Superpixels in brain MR image analysis. In: Proceedings of IEEE EMBC, pp. 1077–1080 (2013). <https://doi.org/10.1109/EMBC.2013.6609691>
36. Wu, W., Chen, A.Y.C., Zhao, L., Corso, J.J.: Brain tumor detection and segmentation in a CRF (conditional random fields) framework with pixel-pairwise affinity and superpixel-level features. Int. J. of Comp. Assist. Radiol. Surg. **9**(2), 241–253 (2014)
37. Zhao, X., Wu, Y., Song, G., Li, Z., Zhang, Y., Fan, Y.: A deep learning model integrating FCNNs and CRFs for brain tumor segmentation. CoRR abs/1702.04528 (2017)
38. Zhuge, Y., et al.: Brain tumor segmentation using holistically nested neural networks in MRI images. Med. Phys. **44**(10), 5234–5243 (2017). <https://doi.org/10.1002/mp.12481>
39. Zikic, D., et al.: Decision forests for tissue-specific segmentation of high-grade gliomas in multi-channel MR. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012. LNCS, vol. 7512, pp. 369–376. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33454-2_46