



No New-Net

Fabian Isensee¹(✉), Philipp Kickingereder², Wolfgang Wick³,
Martin Bendszus², and Klaus H. Maier-Hein¹

- ¹ Division of Medical Image Computing, German Cancer Research Center (DKFZ),
Heidelberg, Germany
f.isensee@dkfz-heidelberg.de
- ² Department of Neuroradiology, University of Heidelberg Medical Center,
Heidelberg, Germany
- ³ Neurology Clinic, University of Heidelberg Medical Center, Heidelberg, Germany

Abstract. In this paper we demonstrate the effectiveness of a well trained U-Net in the context of the BraTS 2018 challenge. This endeavour is particularly interesting given that researchers are currently besting each other with architectural modifications that are intended to improve the segmentation performance. We instead focus on the training process arguing that a well trained U-Net is hard to beat. Our baseline U-Net, which has only minor modifications and is trained with a large patch size and a Dice loss function indeed achieved competitive Dice scores on the BraTS2018 validation data. By incorporating additional measures such as region based training, additional training data, a simple postprocessing technique and a combination of loss functions, we obtain Dice scores of 77.88, 87.81 and 80.62, and Hausdorff Distances (95th percentile) of 2.90, 6.03 and 5.08 for the enhancing tumor, whole tumor and tumor core, respectively on the test data. This setup achieved rank two in BraTS2018, with more than 60 teams participating in the challenge.

Keywords: CNN · Brain tumor · Glioblastoma · U-Net · Dice loss

1 Introduction

Quantitative assessment of brain tumors provides valuable information and therefore constitutes an essential part of diagnostic procedures. Automatic segmentation is attractive in this context, as it allows for faster, more objective and potentially more accurate description of relevant tumor parameters, such as the volume of its subregions. Due to the irregular nature of tumors, however, the development of algorithms capable of automatic segmentation remains challenging.

The brain tumor segmentation challenge (BraTS) [1] aims at encouraging the development of state of the art methods for tumor segmentation by providing a large dataset of annotated low grade gliomas (LGG) and high grade glioblastomas (HGG). The BraTS 2018 training dataset, which consists of 210 HGG and 75 LGG cases, was annotated manually by one to four raters and

all segmentations were approved by expert raters [2–4]. For each patient a T1 weighted, a post-contrast T1-weighted, a T2-weighted and a Fluid-Attenuated Inversion Recovery (FLAIR) MRI was provided. The MRI originate from 19 institutions and were acquired with different protocols, magnetic field strengths and MRI scanners. Each tumor was segmented into edema, necrosis and non-enhancing tumor and active/enhancing tumor. The segmentation performance of participating algorithms is measured based on the DICE coefficient, sensitivity, specificity and 95th percentile of Hausdorff distance.

It is unchallenged by now that convolutional neural networks (CNNs) dictate the state of the art in biomedical image segmentation [5–10]. As a consequence, all winning contributions to recent BraTS challenges were exclusively build around CNNs. One of the first notably successful neural network for brain tumor segmentation was DeepMedic, a 3D CNN introduced by Kamnitsas et al. [5]. It comprises a low and a high resolution pathway that capture semantic information at different scales and recombines them to predict a segmentation based on precise local as well as global image information. Kamnitsas et al. later enhanced their architectures with residual connections for BraTS 2016 [11]. With the success of encoder-decoder architectures for semantic segmentation, such as FCN [12, 13] and most notably the U-Net [14], it is unsurprising that these architectures are used in the context of brain tumor segmentation as well. In BraTS 2017, all winning contributions were at least partially based on encoder-decoder networks. Kamnitsas et al. [9], who were the clear winner of the challenge, created an ensemble by combining three different network architectures, namely 3D FCN [12], 3D U-Net [14, 15] and DeepMedic [5], trained with different loss functions (Dice loss [16, 17] and crossentropy) and different normalization schemes. Wang et al. [10] used a FCN inspired architecture, enhanced with dilated convolutions [13] and residual connections [18]. Instead of directly learning to predict the regions of interest, they trained a cascade of networks that would first segment the whole tumor, then given the whole tumor the tumor core and finally given the tumor core the enhancing tumor. Isensee et al. [6] employed a U-Net inspired architecture that was trained on large input patches to allow the network to capture as much contextual information as possible. This architecture made use of residual connections [18] in the encoder only, while keeping the decoder part of the network as simple as possible. The network was trained with a multiclass Dice loss and deep supervision to improve the gradient flow.

Recently, a growing number of architectural modifications to encoder-decoder networks have been proposed that are designed to improve the performance of the networks for their specific tasks [6, 7, 10, 17, 19–22]. Due to the sheer number of such variants, it becomes increasingly difficult for researchers to keep track of which modifications extend their usefulness over the few datasets they are typically demonstrated on. We have implemented a number of these variants and found that they provide no additional benefit if integrated into a well trained U-Net. In this context, our contribution to the BraTS 2018 challenge is intended to demonstrate that such a U-Net, without using significant architectural alterations, is capable of generating competitive state of the art segmentations.

2 Methods

In the following we present the network architecture and training schemes used for our submission. As hinted in the previous paragraph, we will use a 3D U-Net architecture that is very close to its original publication [15] and optimize the training procedure to maximize its performance on the BraTS 2018 training and validation data.

2.1 Preprocessing

With MRI intensity values being non standardized, normalization is critical to allow for data from different institutes, scanners and acquired with varying protocols to be processed by one single algorithm. This is particularly true for neural networks where imaging modalities are typically treated as color channels. Here we need to ensure that the value ranges match not only between patients but between the modalities as well in order to avoid initial biases of the network. We found the following workflow to work well. We normalize each modality of each patient independently by subtracting the mean and dividing by the standard deviation of the brain region. The region outside the brain is set to 0. As opposed to normalizing the entire image including the background, this strategy will yield comparative intensity values within the brain region irrespective of the size of the background region around it.

2.2 Network Architecture

U-Net [14] is a successful encoder-decoder network that has received a lot of attention in the recent years. Its encoder part works similarly to a traditional classification CNN in that it successively aggregates semantic information at the expense of reduced spatial information. Since in segmentation, both semantic as well as spatial information are crucial for the success of a network, the missing spatial information must somehow be recovered. U-Net does this through the decoder, which receives semantic information from the bottom of the ‘U’ (see Fig. 1) and recombines it with higher resolution feature maps obtained directly from the encoder through skip connections. Unlike other segmentation networks, such as FCN [12] and previous iterations of DeepLab [13] this allows U-Net to segment fine structures particularly well.

Our network architecture is an instantiation of the 3D U-Net [15] with minor modifications. Following our successful participation in 2017 [6], we stick with our design choice to process patches of size $128 \times 128 \times 128$ with a batch size of two. Due to the high memory consumption of 3D convolutions with large patch sizes, we implemented our network carefully to still allow for an adequate number of feature maps. By reducing the number of filters right before upsampling and by using inplace operations whenever possible, this results in a network with 30 feature channels at the highest resolution, which is nearly double the number we could train with in our previous model (using the same 12 GB NVIDIA Titan X GPU). Due to our choice of loss function, traditional ReLU activation functions




Fig. 1. We use a 3D U-Net architecture with minor modifications. It uses instance normalization [23] and leaky ReLU nonlinearities and reduces the number of feature maps before upsampling. Feature map dimensionality is noted next to the convolutional blocks, with the first number being the number of feature channels.

did not reliably produce the desired results, which is why we replaced them with leaky ReLUs (leakiness 10^{-2}) throughout the entire network. With a small batch size of 2, the exponential moving averages of mean and variance within a batch learned by batch normalization [24] are unstable and do not reflect the feature map activations at test time very well. We found instance normalization [23] to provide more consistent results and therefore used it to normalize all feature map activations (between convolution and nonlinearity). For an overview over our segmentation architecture, please refer to Fig. 1.

2.3 Training Procedure

Our network architecture is trained with randomly sampled patches of size $128 \times 128 \times 128$ voxels and batch size 2. We refer to an epoch as an iteration over 250 batches and train for a maximum of 500 epochs. The training is terminated early if the exponential moving average of the validation loss ($\alpha = 0.95$) has not improved within the last 60 epochs. Training is done using the ADAM optimizer with an initial learning rate $l_{\text{init}} = 1 \cdot 10^{-4}$, which is reduced by factor 5 whenever the above mentioned moving average of the validation loss has not improved in the last 30 epochs. We regularize with a l2 weight decay of 10^{-5} .

One of the main challenges with brain tumor segmentation is the class imbalance in the dataset. While networks will train with crossentropy loss function, the resulting segmentations may not be ideal in the sense of the Dice score they obtain. Since the Dice scores is one of the most important metrics based upon which contributions are ranked, it is imperative to optimize this metric. We achieve that by using a soft Dice loss for the training of our network. While several formulations of the Dice loss exist in the literature [16, 17, 25], we prefer to use a multi-class adaptation of [16] which has given us good results in

segmentation challenges in the past [6,8]. This multiclass Dice loss function is differentiable and can be easily integrated into deep learning frameworks:

$$\mathcal{L}_{\text{dc}} = -\frac{2}{|K|} \sum_{k \in K} \frac{\sum_i u_i^k v_i^k}{\sum_i u_i^k + \sum_i v_i^k} \quad (1)$$

where u is the softmax output of the network and v is a one hot encoding of the ground truth segmentation map. Both u and v have shape i by c with i being the number of pixels in the training patch and $k \in K$ being the classes.

When training large neural networks from limited training data, special care has to be taken to prevent overfitting. We address this problem by utilizing a large variety of data augmentation techniques. The following augmentation techniques were applied on the fly during training: random rotations, random scaling, random elastic deformations, gamma correction augmentation and mirroring. Data augmentation was done with our own in-house framework which is publically available at <https://github.com/MIC-DKFZ/batchgenerators>.

The fully convolutional nature of our network allows to process arbitrarily sized inputs. At test time we therefore segment an entire patient at once, alleviating problems that may arise when computing the segmentation in tiles with a network that has padded convolutions. We furthermore use test time data augmentation by mirroring the images and averaging the softmax outputs.

2.4 Region Based Prediction

Wang et al. [10] use a cascade of CNNs to segment first the whole tumor, then the tumor core and finally the enhancing tumor. We believe the cascade and their rather complicated network architecture to be of lesser importance, but the fact that they did not learn the labels (enhancing tumor, edema, necrosis) but instead optimized the regions that are finally evaluated in the challenge directly to be key to their good performance in last years challenge. For this reason we will also train a version of our model where we replace the final softmax with a sigmoid and optimize the three (overlapping) regions (whole tumor, tumor core and enhancing tumor) directly with the Dice loss.

2.5 Cotraining

285 training cases is a lot for medical image segmentation, but may still not be enough to prevent overfitting entirely. We therefore also experiment with cotraining on additional public and institutional data. For public data, we chose to use the BraTS data made available in the context of the Medical Segmentation Decathlon (<http://medicaldecathlon.com>). This dataset comprises 484 cases with ground truth segmentations collected from older BraTS challenges.

Cotraining is done for only two datasets at a time. Given that the label definitions between BraTS 2018 and the other datasets may differ, we use separate segmentation layers ($1 \times 1 \times 1$ convolution) at the end, which act as a supervised version of m heads [26]. During training, each segmentation layer only receives

gradients from examples of its corresponding dataset. The losses of both layers are averaged to obtain the total loss of a minibatch. The rest of the network weights are shared.

2.6 Postprocessing

One of the most challenging parts in the BraTS challenge data is distinguishing small blood vessels in the tumor core region (that must be labeled either as edema or as necrosis) from enhancing tumor. This is particularly detrimental for LGG patients that may have no enhancing tumor at all. The BraTS challenge awards a Dice score of 1 if a label is absent in both the ground truth and the prediction. Conversely, only a single false positive voxel in a patient where no enhancing tumor is present in the ground truth will result in a Dice score of 0. Therefore we replace all enhancing tumor voxels with necrosis if the total number of predicted enhancing tumor is less than some threshold. This threshold is chosen for each experiment independently by optimizing the mean Dice (using the above mentioned convention) on the BraTS 2018 training cases.

2.7 Dice and Cross-entropy

While being widely popular and providing state of the art results on many medical segmentation challenges, the Dice loss has some downsides, such as badly calibrated softmax probabilities (basically binary 0-1 predictions) and occasional convergence issues (if the true positive term is too small for rare classes) compared to the negative log-likelihood loss (also referred to as cross-entropy loss function). We therefore also experiment with using these losses in conjunction by using both a Dice as well as a negative log-likelihood term and simply adding them together to form the total loss (unweighted sum).

3 Experiments and Results

We designed our training scheme by running a five fold cross-validation on the 285 training cases of BraTS 2018. If additional data is used, the additional training cases are split into five folds as well and used for co-training. Training set results are summarized in Table 3, validation set results can be found in Table 2. Unless noted otherwise, validation set results were obtained by using the five networks from the training cross-validation as an ensemble. For consistency with other publications, all reported values were computed by the online evaluation platform (<https://ipp.cbica.upenn.edu/>).

Due to the relatively small size of the validation set (66 cases vs 285 training cases) we base our main analysis on the cross-validation results. We are confident that conclusions drawn from the training set are more robust and will generalize well to the test set.

Results on the BraTS2018 training data are summarized in Table 3. We refer to our basic U-Net that was trained on BraTS2018 training data with

Table 1. Results on BraTS 2018 training data (285 cases). All results were obtained by running a five fold cross-validation. Metrics were computed by the online evaluation platform.

	Dice			HD95		
	enh.	whole	core	enh.	whole	core
Isensee et al. (2017) [6]	70.69	89.51	82.76	6.24	6.04	6.95
baseline	73.43	89.76	82.17	4.88	5.86	7.11
baseline + reg	73.81	90.02	82.87	5.01	6.26	6.48
baseline + reg + cotr (dec)	75.94	91.33	85.28	4.29	4.82	5.05
baseline + reg + cotr (dec) + post	78.68	91.33	85.28	3.49	4.82	5.05
baseline + reg + cotr (dec) + post + DC&CE	78.62	91.75	85.69	2.84	4.88	5.11
baseline + reg + cotr (inst) + post + DC&CE	76.32	90.35	84.36	3.74	5.64	5.98
baseline + reg + post + DC&CE	76.78	90.30	83.55	3.66	5.36	6.03

Table 2. Results on BraTS2018 validation data (66 cases). Results were obtained by using the five models from the training set cross-validation as an ensemble. Metrics were computed by the online evaluation platform.

	Dice			HD95		
	enh.	whole	core	enh.	whole	core
baseline	79.59	90.80	84.32	3.12	4.79	8.16
baseline + reg + cotr (dec) + post + DC&CE (*)	80.46	91.21	85.77	2.52	4.38	6.73
baseline + reg + cotr (inst) + post + DC&CE (**)	80.95	91.15	86.6	2.44	5.02	6.73
baseline + reg + post + DC&CE	80.66	90.92	85.22	2.74	5.83	7.20
ensemble of (*) and (**)	80.87	91.26	86.34	2.41	4.27	6.52

large input patches and a Dice loss function as *baseline*. With Dice scores of 73.43/89.76/82.17 (enh/whole/core) on the training set this baseline model is by itself already very strong, especially when compared to the model of Isensee et al. [6] that achieved the third place in BraTS2017 (the training data for both challenges is identical, allowing a direct comparison of the models). Adding region based training (*reg*) improved the Dice scores of both the enhancing tumor as well as the tumor core. When training with decathlon data (*cotr (dec)*), we gain two Dice points in enhancing tumor and minor improvements for the tumor core. Our postprocessing, which is targeted at correcting false positive enhancing tumor predictions in LGG patients has a substantial impact on enhancing tumor Dice. On the training set it increases the mean enhancing tumor Dice by almost three points. Using the sum Dice and cross-entropy as a loss function yields yet another small improvement. Interestingly, using our institutional data for cotraining yields much worse results on the training set. In order to isolate the impact of additional training data we added the model *baseline + reg + post + DC&CE* to the table.

While the model that uses institutional data performed worse on the training set, it was slightly better on the validation set (see Table 2). We explain this discrepancy by the possibility that the Dice and Hausdorff distance scores obtained from the training set cross-validation may be overestimated when cotraining with decathlon data. Since any potential case correspondences between decathlon data and BraTS2018 is unknown due to the naming scheme of the decathlon cases, we cannot exclude the possibility that cases that are currently in the validation split for BraTS 2018 appear in the training split of the decathlon data (albeit with different ground truth segmentations). This uncertainty, along with the strong performance of the model cotrained with institutional data on the validation set led us to the decision to submit an ensemble of these two models. The ensemble achieves Dice scores of 80.87/91.26/86.34 (enh/whole/core) and Hausdorff distances of 2.41/4.27/6.52 on the validation set. For comparison, we also included the validation set result achieved with no additional training data.

Figure 2 shows a qualitative example segmentation. The patient shown is taken from the validation set (CBICA_AZA_1). As can be seen in the middle (t1ce), there are several blood vessels close to the enhancing tumor. Segmentation CNNs typically struggle to correctly differentiate between such vessels and actual enhancing tumor. This is most likely due to a) a difficulty in detecting tube-like structures b) few training cases where these vessels are an issue c) the use of Dice loss functions that does not sufficiently penalize false segmentations of vessels due to their relatively small size. In the case shown here, our model correctly segmented the vessels as background.

Test set results (as communicated by the organizers of the challenge) are presented in Table 3. We used an ensemble of the two models that were trained with institutional and decathlon data for our final submission. Each of these models is in turn an ensemble of five models resulting from the corresponding cross-validation, resulting in a total of 10 predictions for each test case. Our algorithm achieved the second place out of 64 participating teams. We compare



Fig. 2. Qualitative results. The case shown here is patient CBICA_AZA_1 from the validation set. Left: flair, middle: t1ce, right: our segmentation. Enhancing tumor is shown in yellow, necrosis in turquoise and edema in violet. (Color figure online)

Table 3. Test set results of NVDLMED, the winner of BraTS2018, and our method, which achieved the second place.

		Dice			Hausd. dist.		
		enh.	whole	core	enh.	whole	core
NVDLMED	Mean	76.64	88.39	81.54	3.77	5.90	4.81
	StdDev	25.57	11.83	24.99	8.61	10.01	7.52
	Median	84.41	92.06	91.67	1.73	3.16	2.45
MIC-DKFZ	Mean	77.88	87.81	80.62	2.90	6.03	5.08
	StdDev	23.93	12.89	25.02	3.85	9.98	8.09
	Median	84.94	91.79	90.72	1.73	3.16	2.83

our results to the winning contribution by Myronenko et al. (team NVDLMED). While our model had strong results for enhancing tumor, NVDLMED outperformed our approach in both tumor core and whole tumor. Please refer to [27] for a detailed summary of the challenge results.

4 Discussion

In this paper we demonstrated that a generic U-Net architecture that has only minor modifications can obtain very competitive segmentation, if trained correctly. While our base model is already quite strong, enhancing its training procedure by using region-based training, cotraining with additional training data, postprocessing to target false positive enhancing tumor detection as well as a combination of Dice and cross-entropy loss, increases its performance substantially. For our final submission we used an ensemble of a model cotrained with public and another cotrained with institutional data. Despite using only a generic U-Net architecture, our approach achieved the second place in the BraTS2018 challenge, underlining the impact a well designed framework can have on model training.

References

1. Menze, B.H., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE TMI* **34**(10), 1993–2024 (2015)
2. Bakas, S., et al.: Advancing The Cancer Genome Atlas glioma MRI collections with expert segmentation labels and radiomic features. *Nat. Sci. Data* **4**, 170117 (2017)
3. Bakas, S., et al.: Segmentation labels and radiomic features for the pre-operative scans of the TCGA-GBM collection. *TCIA* (2017)
4. Bakas, S., et al.: Segmentation labels and radiomic features for the pre-operative scans of the TCGA-LGG collection. *TCIA* (2017)
5. Kamnitsas, K., et al.: Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *MIA* **36**, 61–78 (2017)

6. Isensee, F., Kickingereder, P., Wick, W., Bendszus, M., Maier-Hein, K.H.: Brain tumor segmentation and radiomics survival prediction: contribution to the BRATS 2017 challenge. In: Crimi, A., Bakas, S., Kuijff, H., Menze, B., Reyes, M. (eds.) BrainLes 2017. LNCS, vol. 10670, pp. 287–297. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-75238-9_25
7. Li, X., Chen, H., Qi, X., Dou, Q., Fu, C.-W., Heng, P.A.: H-DenseUNet: Hybrid densely connected UNet for liver and liver tumor segmentation from CT volumes. arXiv preprint [arXiv:1709.07330](https://arxiv.org/abs/1709.07330) (2017)
8. Isensee, F., Jaeger, P.F., Full, P.M., Wolf, I., Engelhardt, S., Maier-Hein, K.H.: Automatic cardiac disease assessment on cine-MRI via time-series segmentation and domain specific features. In: Pop, M., et al. (eds.) STACOM 2017. LNCS, vol. 10663, pp. 120–129. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-75541-0_13
9. 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
10. Wang, G., Li, W., Ourselin, S., Vercauteren, T.: Automatic brain tumor segmentation using cascaded anisotropic convolutional neural networks. In: Crimi, A., Bakas, S., Kuijff, H., Menze, B., Reyes, M. (eds.) BrainLes 2017. LNCS, vol. 10670, pp. 178–190. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-75238-9_16
11. Kamnitsas, K., et al.: DeepMedic for brain tumor segmentation. In: Crimi, A., Menze, B., Maier, O., Reyes, M., Winzeck, S., Handels, H. (eds.) BrainLes 2016. LNCS, vol. 10154, pp. 138–149. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-55524-9_14
12. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3431–3440 (2015)
13. Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: DeepLab: semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**(4), 834–848 (2018)
14. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
15. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_49
16. Drozdal, M., Vorontsov, E., Chartrand, G., Kadoury, S., Pal, C.: The importance of skip connections in biomedical image segmentation. In: Carneiro, G., et al. (eds.) LABELS/DLMIA -2016. LNCS, vol. 10008, pp. 179–187. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46976-8_19
17. Milletari, F., Navab, N., Ahmadi, S.-A.: V-Net: Fully convolutional neural networks for volumetric medical image segmentation. In: International Conference on 3D Vision, pp. 565–571. IEEE (2016)
18. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9908, pp. 630–645. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46493-0_38

19. Jégou, S., Drozdal, M., Vazquez, D., Romero, A., Bengio, Y.: The one hundred layers tiramisu: fully convolutional densenets for semantic segmentation. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 1175–1183. IEEE (2017)
20. Oktay, O., et al.: Attention U-Net: learning where to look for the pancreas. arXiv preprint [arXiv:1804.03999](https://arxiv.org/abs/1804.03999) (2018)
21. Roy, A.G., Navab, N., Wachinger, C.: Concurrent spatial and channel squeeze & excitation in fully convolutional networks. arXiv preprint [arXiv:1803.02579](https://arxiv.org/abs/1803.02579) (2018)
22. Kayalibay, B., Jensen, G., van der Smagt, P.: CNN-based segmentation of medical imaging data. arXiv preprint [arXiv:1701.03056](https://arxiv.org/abs/1701.03056) (2017)
23. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Instance normalization: The missing ingredient for fast stylization. arXiv preprint [arXiv:1607.08022](https://arxiv.org/abs/1607.08022) (2016)
24. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint [arXiv:1502.03167](https://arxiv.org/abs/1502.03167) (2015)
25. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Jorge Cardoso, M.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: Cardoso, M.J., et al. (eds.) DLMIA/ML-CDS -2017. LNCS, vol. 10553, pp. 240–248. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67558-9_28
26. Lee, S., Purushwalkam, S., Cogswell, M., Crandall, D., Batra, D.: Why M heads are better than one: Training a diverse ensemble of deep networks. arXiv preprint [arXiv:1511.06314](https://arxiv.org/abs/1511.06314) (2015)
27. Bakas, S., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. arXiv preprint [arXiv:1811.02629](https://arxiv.org/abs/1811.02629) (2018)