



TriadNet: Sampling-free predictive intervals for lesional volume in 3D brain MR images

Benjamin Lambert, Florence Forbes, Senan Doyle, Michel Dojat

► To cite this version:

Benjamin Lambert, Florence Forbes, Senan Doyle, Michel Dojat. TriadNet: Sampling-free predictive intervals for lesional volume in 3D brain MR images. UNSURE 2023 - 5th International Workshop on Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, Oct 2023, Vancouver (BC), Canada. 10.48550/arXiv.2307.15638 . hal-04436218

HAL Id: hal-04436218

<https://hal.science/hal-04436218v1>

Submitted on 5 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

TriadNet: Sampling-free predictive intervals for lesional volume in 3D brain MR images

Benjamin Lambert^{1,2}, Florence Forbes³, Senan Doyle², and Michel Dojat¹

¹ Univ. Grenoble Alpes, Inserm, U1216, Grenoble Institut Neurosciences, 38000, FR

² Pixyl, Research and Development Laboratory, 38000 Grenoble, FR

³ Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, FR

Abstract. The volume of a brain lesion (e.g. infarct or tumor) is a powerful indicator of patient prognosis and can be used to guide the therapeutic strategy. Lesional volume estimation is usually performed by segmentation with deep convolutional neural networks (CNN), currently the state-of-the-art approach. However, to date, few work has been done to equip volume segmentation tools with adequate quantitative predictive intervals, which can hinder their usefulness and acceptance in clinical practice. In this work, we propose **TriadNet**, a segmentation approach relying on a multi-head CNN architecture, which provides both the lesion volumes and the associated predictive intervals simultaneously, in less than a second. We demonstrate its superiority over other solutions on BraTS 2021, a large-scale MRI glioblastoma image database.

Keywords: Brain MRI · Prediction Intervals · Uncertainty · Segmentation · Deep Learning

1 Introduction

The lesional volume is a powerful and commonly used biomarker in brain MRI analysis and interpretation. Such an imaging biomarker is a guide to predict the patient’s neurological outcome in Stroke [8] or to assess the grade of a Glioblastoma [3]. For Multiple Sclerosis (MS), the evolution of the lesional load between two patient’s visits helps to assess the progress of the disease and to personalize his/her treatment [14] and even to predict the disability [19]. For neurodegenerative diseases such as Alzheimer’s disease, the brain atrophy is quantified by estimating the volume of different anatomical regions (e.g. hippocampus or amygdala) compared to normative values [5].

Volume estimation is usually carried out through image segmentation, relying on Deep Convolutional Neural Networks (CNNs) trained on an annotated database, comprising both images and their corresponding manual delineations [10]. CNNs provide a mask, which is generally correct for easy detectable regions or lesions, but whose accuracy may be more uncertain when the zone to segment is disputable even for an expert. To help clinicians to focus on the more subtle regions, we propose to associate quantitative Predictive Intervals (PIs)

to volume estimation. Such PIs can straightforwardly be interpreted as uncertainty markers and facilitate the acceptance of advanced computerized tools by practitioners.

PI construction has been mainly studied in the context of 1D regression tasks [12,17,22] and applications in the context of medical image processing are very scarce. To compute PIs for lesion counting in 2D medical images, reference work proposes either a sampling approach or a regression model [6]. In the former, several plausible and diverse segmentation masks are generated for the same input image, forming a distribution over the quantity of interest (e.g lesion volume or number), from which the mean and the standard deviation can be extracted to define a PI. This Uncertainty Quantification (UQ) methodology offers several variants to generate the diverse set of predictions. Popular UQ methods regroup the Monte Carlo Dropout (MC) [7], Deep Ensemble [13], or Test Time Augmentation (TTA) [23]. Based on sampling, UQ methods are associated with an important computational burden to obtain the predictions. With the regression approach, a network is trained to directly predict the PI's components: the mean value as well as the lower and upper bounds from the data themselves. As no assumptions are made regarding the distribution of the regressed variable, this approach is referred to as Distribution-Free Uncertainty Quantification (DFUQ) [17]. In this direction, we introduce a sampling-free approach based on an original CNN architecture called TriadNet which exhibits the following assets:

- It enhances the 3D volume estimation with associated reliable PIs.
- It allows a fast and distribution-free estimation of PIs.
- The methodology is simple to implement and can be applied to any encoder-decoder segmentation architecture.

2 Problem Definition

We consider a 3D segmentation problem with N classes. Excluding the background class, we aim at estimating the true unknown volumes $Y \in \mathbb{R}^{N-1}$ of each foreground classes based on the predicted segmentation. In this context, for an estimation X of the volume, seen as a random variable, we define a predictive interval $\Gamma_\alpha(X)$ as a range of values that are conditioned to contain Y , the actual volume, with a certain degree of confidence $1 - \alpha$ (e.g 90% or 95%). That is, given a series of estimated volumes $X_1 \dots X_n$ and their associated ground truth volumes $Y_1 \dots Y_n$, $\Gamma_\alpha(\cdot)$ should be learned as to satisfy:

$$P(Y_{\text{test}} \in \Gamma_\alpha(X_{\text{test}})) \geq 1 - \alpha \quad (1)$$

for any $(Y_{\text{test}}, X_{\text{test}})$ following the same distribution as the (Y_i, X_i) 's. This property is called the *marginal coverage*, as the probability is marginal over the entire test dataset [1].

Sampling-based PI estimation methods rely on the hypothesis that X follows a normal distribution for each predicted class. Under this assumption, the mean value μ_X and standard deviation σ_X of the distribution are estimated by

sampling several distinct predictions for the same input, and PI are constructed as $I_\alpha(X) = [\mu_X - z\sigma_X, \mu_X + z\sigma_X]$ where z is the number of standard deviation, stipulating the degree of confidence of the interval. For instance, for a 90% confidence interval, z corresponds to 1.65. In contrast, direct PI estimation techniques (including regression approaches and our proposed TriadNet) directly output the mean value μ , the lower bound l_b and the upper bound u_b ($l_b \leq \mu \leq u_b$), without sampling.

3 Our solution: TriadNet

Overview: TriadNet corresponds to a CNN model modified in order to produce three outputs for each segmented class: the mean, lower bound and upper bound masks (see Figure 1). To obtain these distinct masks, we propose a multi-head architecture as well as a novel learning objective, the TriadLoss. The masks are then used to directly estimate the class-wise mean volume as well as the lower and upper bounds, by summing the segmented voxels.




Fig. 1. The Triadnet architecture. Each head yields a distinct mask for each class: lower bound, mean and upper bound masks. For ease of visualization, we only represent for a Glioblastoma application, the masks for the *edematous* class.

TriadNet: the architecture relies on the Attention Unet 3d (AttUNet3d) [16] as backbone. We modified it by duplicating the output convolutional block in order to obtain a total of 3 separate and identical heads. Each head generates a specific mask: respectively, one corresponding to the lower bound, one to the upper bound, and one for the mean value, by predicting a probabilistic distribution $p_{n,i}$ over the N classes and for each voxel i . This modification only slightly increase the complexity of the segmentation model, raising the number of parameters from 5 millions to 5.3 millions.

TriadLoss: the objective function is built on the observation that the lower bound mask should be more restrictive (*i.e.* higher precision and lower recall) than the mean mask. Similarly, the upper bound mask should be more permissive (*i.e.* higher recall and lower precision). To achieve this, we propose to rely on the Tversky loss [20], which provides a direct control on the trade-off between recall and precision. The Tversky loss $T_{\alpha,\beta}$ is an extension of the popular Dice loss [15], with 2 extra hyperparameters α and β which respectively control the weighting of False Positives (FP) and False Negatives (FN). With $\alpha = \beta = 0.5$, the Tversky loss is strictly equivalent to the standard Dice loss.

Writing p_{lower} , p_{mean} and p_{upper} the outputs of each head and y the ground-truth segmentation, we defined the Triad loss as:

$$\text{TriadLoss} = T_{1-\gamma,\gamma}(p_{lower,y}) + T_{0.5,0.5}(p_{mean,y}) + T_{\gamma,1-\gamma}(p_{upper,y}) \quad (2)$$

with γ an hyperparameter in the range $]0, 0.5[$ controlling the penalties applied to FP and FN during the training of the lower and upper bound heads. In other words, the mean decoder was trained with a standard Dice Loss. To obtain more restrictive masks (and lower volumes), the lower bound decoder was trained to minimize FP at the expense of a higher FN rate. Similarly, to obtain more permissive masks (and larger volumes), the upper bound decoder sought to minimize FN at the expense of a higher number of FP.

4 Material and Methods

4.1 Datasets

We illustrate our framework on a brain tumor segmentation task, using the open-source part of the BraTS 2021 dataset [2] containing 1251 patients. Four MRI sequences are available for each patient: FLAIR, T1, T2, and T1ce (T1w with contrast agent). The ground truth segmentation masks contain 4 classes: the background, the necrotic tumor core, the edematous, and the GD-enhancing (GDE) tumor. We randomly split the data into a training fold (651), a calibration fold (200), and a testing fold (400).

4.2 Comparison with known approaches

We compared TriadNet with 3 sampling-based approaches: Confidence Thresholding (CT), Monte Carlo dropout (MC), and Test Time Augmentation (TTA), as well as a sampling-free PI estimation framework based on the training of a regression CNN (RegCNN).

Confidence Thresholding (CT) is a simple approach to obtain PI's from the output probability estimates produced by a trained segmentation model. For each class, the probability map is binarized with progressively increasing thresholds. As the threshold increases, fewer voxels are segmented, thus the volume decreases. As this method relies on the calibration of the output probabilities, we perform Temperature Scaling [9] on the trained segmentation model before performing CT.

Monte Carlo Dropout (MC) is based on the Dropout technique [21] which consists of turning a subset of the model parameters off, to prevent overfitting. The MC dropout technique proposes to keep dropout activated during inference, meaning that T forward steps of the same image through the MC dropout model will lead to T different segmentations (and thus volume estimates), as the dropout mask is randomly sampled at each step.

Test Time Augmentation (TTA) consists in using data augmentation to generate alternative versions of the input images. Each augmented image is processed by the segmentation model, yielding to a distinct estimation of the volumes. By repeating this process, a distribution over volumes can be obtained, from which the PI is derived.

Regression CNN (RegCNN) proposes to train a regression neural network to directly predict the lower, mean and upper bounds of the target quantity from the data itself [1,6]. To achieve this, the Pinball loss P_t can be used to train the model to predict a desired quantile t . In our study, the regressor took as input the MRI sequences and automated segmentation produced by a segmentation model and was trained to predict three scores for each segmentation class, namely the $q_{\alpha/2}$, $q_{0.5}$ and $q_{1-\alpha/2}$ quantiles, allowing the construction of $(1-\alpha)\%$ confidence intervals. To do so, the regressor was trained with a compound loss $L = P_{\alpha/2} + P_{0.5} + P_{1-\alpha/2}$ to learn each quantile.

4.3 Post-hoc PI calibration

In practice, the predicted PIs may be inaccurate and not respect the desired *marginal coverage* property. To alleviate this, PI post-hoc calibration is usually performed using a set-aside calibration dataset [1]. This calibration step aims at finding the optimal corrective value q such that the calibrated PIs achieve the desired $(1-\alpha)\%$ coverage on the calibration dataset.

In the case of sampling-based PI, the corrective value takes the form of a multiplicative factor applied to the standard deviation (Equation 3). Alternatively, if the PI estimation is direct, q corresponds to an additive factor applied

to the lower and upper bounds (Equation 4) :

$$\Gamma_{\alpha, \text{cal}}(X) = [\mu_X - q\sigma_X, \mu_X + q\sigma_X] \quad (3)$$

$$\Gamma_{\alpha, \text{cal}}(X) = [l_b - q, u_b + q] \quad (4)$$

4.4 Evaluation

We performed all our experiments with $\alpha = 0.1$, meaning that we focussed on 90% PIs. Segmentation performance was assessed using the Dice score (DSC) between the predicted segmentation and the ground truth delineations (for TriadNet, the Dice was computed using the *mean* predicted mask). We also used the Mean Average Error (MAE) between the estimated mean volumes and the true volumes to assess the reliability of the volume prediction.

Useful PIs should have two properties. They should i) achieve the desired *marginal coverage* and ii) be the narrowest possible in order to be informative. To verify this, we computed two scores for PIs: the coverage error (Δf) and the interval width (W). Δf is defined as the distance between the empirical coverage and target coverage (90%). W is the average distance between the lower and upper bounds. Note that a successful PI calibration should ensure $\Delta f \geq 0$. However, as the width of intervals tend to augment with Δf , a value close to 0 is preferred. To estimate computational efficiency, we also reported the average time to produce a segmentation and PI for one input MRI volume.

To assess the impact of the choice of the γ hyper-parameter in the TriadLoss on PI quality, we trained Triadnet models with varying γ values, ranging from 0.1 to 0.4. To obtain robust statistics, each model is trained 5 times and we reported the average and standard deviation for each metrics.

4.5 Implementation Details

Three types of segmentation models are used in this study. First, *Baseline* AtUnet3d was trained to serve as a common basis for the implementation of CT, TTA and RegCNN approaches. For MC, we trained a dedicated *Dropout* AtUnet3ds by adding a dropout rate of 20% in each layer of the encoder and decoder. The last type of segmentation model was our proposed TriadNet. All models were trained with the ADAM optimizer [11], with a learning rate of $2e-4$, using the Dice loss for *Baseline* and *Dropout* models and the TriadLoss for TriadNet. For CT-based PIs, we used 20 different thresholds uniformly distributed in the range $[0.01, 0.99]$ to binarize the probability maps. For MC dropout, we performed $T = 20$ forward passes of the same input image with dropout activated to obtain the PIs. To implement the TTA baseline, we generated 20 random augmentations for each input MRI using flipping, rotation, translation and contrast augmentation with randomized parameters, implemented using the TorchIO Data Augmentation library [18]. Finally for RegCNN, we used an open-source regressor CNN implementation ⁴ [4].

⁴ https://docs.monai.io/en/stable/_modules/monai/networks/nets/regressor.html

5 Results and Discussion

Table 1 presents the performance of segmentation (DSC) and PIs for each approach and for all 3 segmented tumor tissues; and Table 2, the average computation time for each method. Finally, Figure 2 provides an illustration of PI computed by our proposed TriadNet on the test dataset.

Table 1. Performances for each tumor tissue for each method. Δf : coverage error, W : average interval width. Mean scores obtained over 5 runs. SD: standard deviation.

	Method	Δf (% $\pm$ SD)	$W \downarrow$ ($mL \pm SD$)	MAE $\downarrow$ ($mL \pm SD$)	DSC $\uparrow$ ($\pm SD$)
Necrotic	CT	5.6 ± 1.5	32.3 ± 5.3	3.4 ± 0.1	0.76 ± 0.00
	TTA	6.3 ± 1.1	25.0 ± 3.7	3.5 ± 0.1	0.76 ± 0.00
	RegCNN	6.1 ± 1.7	25.8 ± 5.7	6.3 ± 3.6	0.76 ± 0.00
	MC dropout	5.6 ± 0.6	20.6 ± 2.1	3.3 ± 0.1	0.76 ± 0.00
	TriadNet ($\gamma = 0.1$)	4.5 ± 0.8	14.5 ± 1.4	3.4 ± 0.1	0.75 ± 0.00
	TriadNet ($\gamma = 0.2$)	3.4 ± 0.6	13.7 ± 1.1	3.2 ± 0.1	0.76 ± 0.00
	TriadNet ($\gamma = 0.3$)	4.1 ± 0.9	15.0 ± 0.4	3.3 ± 0.1	0.76 ± 0.01
	TriadNet ($\gamma = 0.4$)	4.4 ± 0.4	16.6 ± 0.8	3.4 ± 0.1	0.75 ± 0.00
Edematous	CT	1.4 ± 1.2	54.4 ± 12.3	8.2 ± 0.9	0.85 ± 0.01
	TTA	-1.3 ± 2.3	34.9 ± 2.5	7.7 ± 0.2	0.85 ± 0.01
	RegCNN	1.7 ± 0.9	41.9 ± 1.7	9.2 ± 0.4	0.85 ± 0.01
	MC dropout	-0.01 ± 1.4	32.0 ± 2.0	7.5 ± 0.2	0.84 ± 0.01
	TriadNet ($\gamma = 0.1$)	0.9 ± 0.7	31.8 ± 0.8	7.4 ± 0.5	0.85 ± 0.00
	TriadNet ($\gamma = 0.2$)	1.6 ± 1.2	30.2 ± 1.7	7.2 ± 0.2	0.85 ± 0.00
	TriadNet ($\gamma = 0.3$)	3.2 ± 1.6	35.7 ± 4.3	7.1 ± 0.2	0.85 ± 0.00
	TriadNet ($\gamma = 0.4$)	1.4 ± 0.7	31.1 ± 2.2	7.5 ± 0.3	0.84 ± 0.01
GDE	CT	3.6 ± 1.4	17.8 ± 3.5	2.0 ± 0.0	0.85 ± 0.01
	TTA	3.0 ± 1.8	10.6 ± 1.0	2.0 ± 0.1	0.85 ± 0.01
	RegCNN	0.7 ± 0.5	22.2 ± 5.9	3.2 ± 0.3	0.85 ± 0.01
	MC dropout	3.5 ± 1.3	10.0 ± 0.5	1.9 ± 0.1	0.85 ± 0.01
	TriadNet ($\gamma = 0.1$)	3.7 ± 1.1	11.0 ± 0.6	1.8 ± 0.1	0.85 ± 0.00
	TriadNet ($\gamma = 0.2$)	4.2 ± 1.4	8.9 ± 0.5	1.8 ± 0.1	0.85 ± 0.00
	TriadNet ($\gamma = 0.3$)	4.1 ± 0.4	9.3 ± 0.6	1.8 ± 0.1	0.85 ± 0.00
	TriadNet ($\gamma = 0.4$)	4.0 ± 0.5	11.0 ± 0.4	1.9 ± 0.1	0.85 ± 0.00


Most methods provide PIs that, after calibration, achieve the target *marginal coverage* property ($\Delta f \geq 0$). In terms of interval width (W), the narrowest intervals are provided by our proposed TriadNet parameterized by $\gamma = 0.2$, while MC dropout ranks as the second best approach. To estimate the significance of this result, two-sided paired t-test between both methods were performed, showing that TriadNet’s PI are significantly narrower compared to MC dropout ones ($p < 0.05$ for each tumor class). The best volume estimation, computed using the MAE, is also obtained by TriadNet, while RegCNN estimation is systematically the worst. In terms of segmentation quality (DSC scores), all models achieve very similar performances. Finally, regarding computational efficiency

Table 2. Average prediction time to obtain a segmentation of a 3D MRI volume associated to predictive intervals on the volumes. SD=standard deviation

	CT	TTA	RegCNN	MC	TriadNet
Time (s $\pm$ SD) $\downarrow$	1.1 $\pm$ 0.1	14.3 $\pm$ 1.2	0.3 $\pm$ 0.1	6.4 $\pm$ 0.1	0.6 $\pm$ 0.1

(Table 2), RegCNN appears as the fastest approach, followed by TriadNet, both producing segmentation and associated PIs in less than one second for an input MRI volume. As expected, sampling approaches are much more time-consuming, with MC and TTA being respectively 10 and 24 times slower than our proposed TriadNet.

The choice of the γ parameter in the TriadLoss has proved to be important, with an optimal PI quality reached for $\gamma = 0.2$, equivalent to a weighting of 0.8 for FP and 0.2 for FN in the lower bound head; and 0.2 for FP and 0.8 for FN in the upper bound head. This setting allows the different masks (lower, mean and upper) to be different enough to allow a reliable PI estimation, which is not the case with higher γ values ($\gamma = 0.3$ and $\gamma = 0.4$). However when γ is lower ($\gamma = 0.1$), the penalty on FP and FN is too small, which yields to a larger amount of erroneous predictions, lowering PI quality.

**Fig. 2.** Predictive intervals generated by TriadNet ($\gamma = 0.2$) on the test dataset.

6 Conclusion

In this work, we addressed the problem of constructing PIs associated to 3D brain MR segmented volumes. Our proposed TriadNet provides narrower and thus more informative intervals in practice compared to competing methods, while preserving the desired *marginal coverage* property. Interestingly, it is also 10 times faster than the second best baseline, MC dropout, making it suitable for clinical routine. Finally, it only requires a minor modification of the segmentation architecture, which has no negative impact on segmentation quality. Future work

will investigate the robustness of TriadNet’s predictive intervals in the presence of domain shift, and evaluate how our approach behaves with respect to the size of the target region, ranking from very small targets (e.g the hippocampus region or MS lesions) to very large (e.g the overall grey matter volume).

References

1. Angelopoulos, A.N., Bates, S.: A gentle introduction to conformal prediction and distribution-free uncertainty quantification. arXiv preprint arXiv:2107.07511 (2021)
2. Baid, U., Ghodasara, S., et al.: The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. arXiv preprint arXiv:2107.02314 (2021)
3. Baris, M.M., Celik, A.O., et al.: Role of mass effect, tumor volume and peritumoral edema volume in the differential diagnosis of primary brain tumor and metastasis. *Clinical neurology and neurosurgery* **148**, 67–71 (2016)
4. Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., et al.: Monai: An open-source framework for deep learning in healthcare. arXiv preprint arXiv:2211.02701 (2022)
5. Contador, J., Pérez-Millán, A., et al.: Longitudinal brain atrophy and csf biomarkers in early-onset alzheimer’s disease. *NeuroImage: Clinical* **32**, 102804 (2021)
6. Eaton-Rosen, Z., Varsavsky, T., et al.: As easy as 1, 2... 4? uncertainty in counting tasks for medical imaging. *Medical Image Computing and Computer Assisted Intervention (MICCAI) 2019* pp. 356–364 (2019)
7. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. *International Conference on Machine Learning* **48**, 1050–1059 (2016)
8. Ghoneem, A., Osborne, M.T., et al.: Association of socioeconomic status and infarct volume with functional outcome in patients with ischemic stroke. *JAMA Network Open* **5**(4), e229178–e229178 (2022)
9. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: *International conference on machine learning*. pp. 1321–1330. PMLR (2017)
10. Hesamian, M.H., Jia, W., et al.: Deep learning techniques for medical image segmentation: achievements and challenges. *Journal of digital imaging* **32**, 582–596 (2019)
11. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015* (2015)
12. Kivaranovic, D., Johnson, K.D., Leeb, H.: Adaptive, distribution-free prediction intervals for deep networks. *International Conference on Artificial Intelligence and Statistics* pp. 4346–4356 (2020)
13. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems* **30** pp. 6402–6413 (2017)
14. Mattiesing, R.M., Gentile, G., et al.: The spatio-temporal relationship between white matter lesion volume changes and brain atrophy in clinically isolated syndrome and early multiple sclerosis. *NeuroImage: Clinical* **36**, 103220 (2022)
15. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. *2016 fourth international conference on 3D vision (3DV)* pp. 565–571 (2016)

16. Oktay, O., Schlemper, J., et al.: Attention u-net: Learning where to look for the pancreas. *Medical Imaging with Deep Learning (MIDL)* (2018)
17. Pearce, T., Brintrup, A., et al.: High-quality prediction intervals for deep learning: A distribution-free, ensembled approach. *International Conference on Machine Learning* pp. 4075–4084 (2018)
18. Pérez-García, F., Sparks, R., Ourselin, S.: Torchio: a python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning. *Computer Methods and Programs in Biomedicine* **208**, 106236 (2021)
19. Roca, P., Attie, A., et al.: Artificial intelligence to predict clinical disability in patients with multiple sclerosis using flair mri. *Diagnostic and Interventional Imaging* **101**(12), 795–802 (2020)
20. Salehi, S.S.M., et al.: Tversky loss function for image segmentation using 3d fully convolutional deep networks. *Machine Learning in Medical Imaging (MLMI) Workshop, Held in Conjunction with MICCAI 2017* pp. 379–387 (2017)
21. Srivastava, N., Hinton, G., et al.: Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* **15**(1), 1929–1958 (2014)
22. Tagasovska, N., Lopez-Paz, D.: Single-model uncertainties for deep learning. *Advances in Neural Information Processing Systems* **32** (2019)
23. Wang, G., Li, W., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T.: Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. *Neurocomputing* **338**, 34–45 (2019)