

NOVEL STRUCTURAL-SCALE UNCERTAINTY MEASURES AND ERROR RETENTION CURVES: APPLICATION TO MULTIPLE SCLEROSIS

Nataliia Molchanova^{1,2,3} Vatsal Raina^{2,4} Andrey Malinin⁵ Francesco La Rosa⁶
Henning Muller² Mark Gales⁴ Cristina Granziera⁷ Mara Graziani² Meritxell Bach Cuadra^{1,3}

¹ University of Lausanne and Lausanne University Hospital, Switzerland, ² University of Applied Sciences of Western Switzerland, Switzerland, ³ CIBM Center for Biomedical Imaging, Switzerland, ⁴ ALTA Institute, University of Cambridge, UK, ⁵ Shifts Project, Finland, ⁶ Icahn School of Medicine at Mount Sinai, USA, ⁷ University Hospital Basel, Switzerland.

ABSTRACT

This paper focuses on the uncertainty estimation for white matter lesions (WML) segmentation in magnetic resonance imaging (MRI). On one side, voxel-scale segmentation errors cause the erroneous delineation of the lesions; on the other side, lesion-scale detection errors lead to wrong lesion counts. Both of these factors are clinically relevant for the assessment of multiple sclerosis patients. This work aims to compare the ability of different voxel- and lesion-scale uncertainty measures to capture errors related to segmentation and lesion detection, respectively. Our main contributions are (i) proposing new measures of lesion-scale uncertainty that do not utilise voxel-scale uncertainties; (ii) extending an error retention curves analysis framework for the evaluation of lesion-scale uncertainty measures. Our results obtained on the multi-center testing set of 58 patients demonstrate that the proposed lesion-scale measure achieves the best performance among the analysed measures. All code implementations are provided at https://github.com/Medical-Image-Analysis-Laboratory/MS_WML_uncs.

Index Terms— Reliable AI, Structural-scale uncertainty, Multiple sclerosis, White matter multiple sclerosis lesions, Magnetic resonance imaging

1. INTRODUCTION

Magnetic resonance imaging (MRI) is a key imaging tool used for multiple sclerosis (MS) diagnosis and prognosis [1]. Brain white matter lesions (WML), visible in MRI, are included in the McDonald MS diagnostic criteria and used for the assessment of disease progression [2, 3, 4]. Both dissemination in space and time of WML in MRI are needed in clinical assessment of MS patients [2].

Given the variability of lesion sizes, locations and per-patient counts, manual WML annotation is a skill-demanding time-consuming task, yet prone to human errors. Automatic segmentation methods can speed up the annotation process

and reduce annotator bias. Deep learning models constitute the state of the art for MS lesion segmentation [5]. Uncertainty quantification has been previously explored for this task to provide clinicians with information about the reliability and confidence of predictions [6, 7, 8, 9, 10]. Various uncertainty measures of voxel-scale predictions were explored, such as entropy, variance, predicted confidence and mutual information [6, 9, 10]. The same studies suggested evaluating uncertainty measures on a structural scale, *i.e.* for each predicted lesion, by aggregating voxel-scale uncertainties. Some attempts were made to quantify the informativeness of uncertainty measures on voxel and lesion scales [6, 10]. For instance, [6] looked into uncertainty-based prediction filtering as a means to correlate uncertainty and false positive errors, and [10] used accuracy-confidence curves. However, these approaches do not explicitly assess the correspondence between the measures and model errors in segmentation or lesion detection, while such an analysis is essential to decide what uncertainty maps should be presented to clinicians.

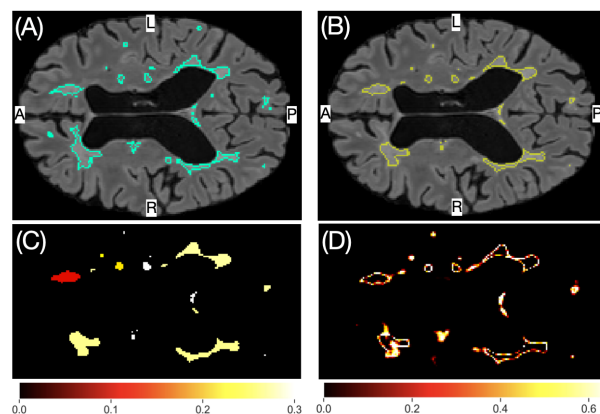


Fig. 1. Example of WML and uncertainty maps. (A) and (B) are ground truth (blue) and predicted (yellow) WML borders overlaid on a FLAIR scan, (C) and (D) lesion-scale (proposed DDU^{true}) and voxel-scale (EoE) uncertainty maps.

In this work, we promote the use of error retention curves (RC) [11] for assessing voxel- and lesion-scale uncertainty measures. We first compare the ability of six voxel-scale measures to capture errors related to the overall segmentation. Our first contribution is novel lesion-scale uncertainty measures that, in contrast to previously proposed ones [6, 9], are not based on voxel-scale uncertainties, but use structural information directly. The second contribution is extending the RC definition to the structural scale to quantify the correspondence between lesion-scale uncertainty measures and lesion detection errors. With this framework, we compare the proposed lesion-scale uncertainty measures to previously proposed ones. In particular, our results show a superior performance of the proposed lesion-scale measure in comparison to the measures based on the voxel-scale uncertainties.

2. METHODOLOGY

2.1. Voxel-scale uncertainty

In this work, we use deep ensembles for uncertainty quantification, which have shown to provide better uncertainty estimates for the WML segmentation task than a more common Monte Carlo dropout [9, 12]. In a classification tasks, several measures of predictive uncertainty can be approximated using predicted probability distributions over the classes from each of the ensemble members. Such uncertainty measures are categorised as data, knowledge and total uncertainty [11]. Data uncertainty measures the inherent noise within the source data distribution. Knowledge uncertainty captures the disagreement in the outputs of each model in the ensemble. This reflects the lack of knowledge by the model in certain regions of the input space. Finally, total uncertainty considers the contributions of both data and knowledge uncertainty.

Total uncertainty	
EoE	$-\sum_{\mathbf{y}} \frac{1}{K} \sum_{k=1}^K P_k(\mathbf{y}) \log \left[\frac{1}{K} \sum_{k=1}^K P_k(\mathbf{y}) \right]$
NC	$-\operatorname{argmax}_{\mathbf{y}} \frac{1}{K} \sum_{k=1}^K P_k(\mathbf{y})$
Data uncertainty	
ExE	$-\frac{1}{K} \sum_{k=1}^K \sum_{\mathbf{y}} P_k(\mathbf{y}) \log P_k(\mathbf{y})$
Knowledge uncertainty	
MI	$\text{EoE} - \text{ExE}$
EPKL	$-\frac{1}{K^2} \sum_{\mathbf{y}} \left[\sum_{k=1}^K P_k(\mathbf{y}) \sum_{k=1}^K \log P_k(\mathbf{y}) \right] - \text{ExE}$
RMI	$\text{EPKL} - \text{MI}$

Table 1. Definitions of voxel-scale uncertainty measures[13]. Notations: $P_k(\mathbf{y}) \equiv P(\mathbf{y}|\mathbf{x}, \theta_k)$ - a predictive posterior of the k^{th} model in the ensemble of size K , \mathbf{y} - vector of model’s outputs, \mathbf{x} - vector of inputs, θ_k - weights of the k^{th} model sampled from a posterior $q(\theta)$.

WML segmentation can be seen as distinct binary classification tasks for each voxel. Thus, voxel-scale uncertainty maps can be calculated using standard measures of data, knowledge and total uncertainty. We estimate the following uncertainty measures (definitions in Table 1): mutual information (MI), expected pair-wise KL divergence (EPKL) and reverse mutual information (RMI) for knowledge uncertainty; expected entropy (ExE) for data uncertainty; entropy of expected (EoE) and negated confidence (NC) for total uncertainty [13]. The entropy-based measures, MI and NC have previously been investigated in MS by [6, 9, 10] while EPKL and RMI are investigated for the first time.

2.2. Lesion-scale uncertainty

Previously considered lesion-scale uncertainty measures aggregate information from a voxel-scale uncertainty map $U \in \mathbb{R}^{H \times W \times D}$, by computing either mean [9] or log-sum [6] of U values within a predicted lesion region Ω , as $\frac{1}{|\Omega|} \sum_{i \in \Omega} U_i$ and $\sum_{i \in \Omega} \log U_i$, where U_i is an uncertainty of the i^{th} voxel of a lesion region Ω . Lesion regions Ω are determined as connected components on a predicted binary lesion map. The voxel-scale uncertainty map U can be obtained using different measures, we explore all the measures defined in Table 1.

Previous investigations showed that the highest uncertainty regions are usually concentrated on the borders of lesions [7, 8, 12] (see Figure 1 D). Thus, voxel-scale uncertainties are more likely to be informative for lesion delineation, and may not be so relevant for lesion detection. From this rationale, we propose two novel structural-scale measures that do not use the voxel-scale uncertainties, but rather look at the disagreement between different models in an ensemble about structural predictions. Let Ω be a lesion region predicted by an ensemble, Ω_k - a region of the same lesion predicted by the k^{th} ensemble member ($k = 1, 2, \dots, K$), *i.e.* a connected component on the binary mask predicted by the k^{th} model having the maximum intersection over union (IoU) with Ω . The proposed measure of the disagreement between the models in the ensemble on the detected area, further referred to as detection disagreement uncertainty (DDU), is defined as one minus agreement, where agreement is an average across ensemble members IoU between the lesion regions predicted by the ensemble and individual ensemble members:

$$\text{DDU} = 1 - \frac{1}{K} \sum_{k=1}^K \text{IoU}(\Omega, \Omega_k) \quad (1)$$

We tested two ways of obtaining binary segmentation maps from ensemble members, which yield different definitions of lesion regions Ω_k : (i) applying the probability threshold tuned for the ensemble to individual model predictions, (ii) tuning the probability thresholds separately for each model in the ensemble. The resulting measures will be referred to as DDU for (i) and DDU^{true} for (ii).

2.3. Error retention curves for uncertainty evaluation

Ideally, a high uncertainty score should highlight the predictions that are the most likely to be wrong. Hence, we expect a good quality uncertainty measure to reflect the increased likelihood of an erroneous prediction and thus correlate with model mistakes. The error RC [11, 12, 14] assesses the degree of this relation by only looking at the ranking of uncertainties. On the voxel scale, RCs analyse the correspondence between a measure of uncertainty and the model errors in segmentation (*e.g.* Dice Similarity Coefficient, DSC). The RC for a single scan is built by sequentially replacing a fraction τ of the most uncertain voxel predictions within the brain mask with the ground truth and re-computing the DSC. The area under the DSC retention curve (DSC-RC), further referred to as DSC-AUC, is a measure of correspondence between an uncertainty measure and segmentation errors that only takes into account the ranking of uncertainties in a particular scan. To summarise the performance on the whole dataset, an average across patients DSC-RC and area under it, *i.e.* $\widehat{\text{DSC-AUC}}$, are computed. It is also possible to estimate lower and upper bounds of performance by building *random* and *ideal* RCs. For a random RC, we assign random uncertainty values to each voxel of predictions. For the ideal one, a zero uncertainty is assigned to true positive and negative (TP and TN) voxels and while false positive and negative (FP and FN) voxels have an uncertainty of 1. To build the RCs, we use $\tau = 2.5 \cdot 10^{-3}$.

RCs are still unstudied at the scale of structural predictions, *i.e.* a lesion scale in our case. We propose to build structural-scale RCs to see the how lesion-scale uncertainty measures are associated with the lesion detection errors. We use a lesion F1-score [15] to measure the detection quality. It is defined as $\frac{2TP}{2TP+FN+FP}$, where TP is a count of predicted lesions that have maximum across the ground truth lesions IoU greater than γ ; inversely FP lesions have a maximum IoU lower than γ ; FN is the count of ground truth lesions not overlapping with predictions. Usually a few voxels overlap is considered enough for lesion detection [6, 16, 17], however we promote a more rigorous criteria and use $\gamma = 0.25$.

Novel analysis of uncertainty with lesion-scale F1-RC poses new challenges. First, compared to voxel-scale predictions, the amount of predicted lesions is much smaller and varies from patient to patient. Therefore, we sequentially reject *each* lesion in a scan based on its uncertainty ranking. Secondly, it is only possible to compute the uncertainties of predicted lesions, *i.e.* TP and FP, as FN lesion regions are defined from the ground truth and TN lesions do not exist. Thus, we suggest the following scheme of sequential lesion rejection: TP lesions remain TP, and FP lesions are just excluded. Note that FN lesions are not replaced in the F1-RC, and the FN count remains fixed in the F1-score formula. Thirdly, lesion F1-RC are computed for each patient in different sets of retention fractions due to varying lesion count across patients. To average subject-wise F1-RCs, we use a piece-wise linear

interpolation to interpolate the obtained curves to a similar set of retention fractions. It allows averaging the F1-RCs at each node to obtain an average across-dataset F1-RC. The lesion $\widehat{\text{F1-AUC}}$ is computed by averaging the F1-AUC of individual patients after interpolation. $\widehat{\text{F1-AUC}}$ quantifies the correspondence between a lesion-scale uncertainty measure and lesion detection errors in a particular dataset. In analogy to the DSC-RC, the ideal and random retention curves are built.

2.4. Dataset

We used a publicly available dataset provided by the Shifts project [12], designed for the study of uncertainty estimation across shifted domains (different locations, scanners, MS stages, *etc.*). It is composed of fluid-attenuated inversion recovery (FLAIR) [18] scans, which underwent denoising, skull stripping, bias field correction and interpolation to $1mm$ iso-voxel space, and their manual WML annotations used as the ground truth. Training and validation datasets contain data from four different centers with 33 and 7 scans, respectively. For testing we used a total of 58 subjects, which comprise both in-domain data sampled from the same distribution as the training one (33 cases) and out-of-domain data collected in a different center and scanner (25 cases). The combination of in- and out-of-domain data is beneficial for RC construction, as it helps to embrace a greater amount of model errors (errors made on both in-domain and out-of-domain data), thus a more reliable evaluation of the uncertainty measures is obtained.

2.5. Models

We use a state-of-the-art WML segmentation method based on 3D U-net architecture [17], which was previously used in uncertainty studies for the particular task [6, 9, 12, 10]. Our final model is an ensemble of five U-nets ($K = 5$), each trained with different random seed initialisation. Training of each model is performed for a maximum of 300 epochs with the best model selection based on validation performance. A combination of Dice and focal losses is used as an objective function optimised using an Adam algorithm with a constant learning rate of 10^{-5} . The architecture takes as an input $96 \times 96 \times 96$ sub-volumes of FLAIR, which are aggregated using Gaussian-weighted averaging at inference time. Binary WML segmentation masks are then obtained by thresholding the average across ensemble members probability maps at a value of 0.3 chosen on a validation set by optimising the DSC. For DDU^{true} measure (eq. 1) computation, binary segmentation maps are obtained from each ensemble member using the thresholds of 0.35, 0.25, 0.25, 0.25, 0.1 chosen separately for each ensemble member based on DSC.

3. RESULTS

The resulting average RCs are shown in Figure 2 and the corresponding AUCs are given in Table 2. At the voxel-scale, ExE and EoE measures show the highest $\widehat{DSC-AUC}$ at the level of 0.9853 ± 0.002 , meaning the best ability to capture errors in overall segmentation. Nevertheless, all uncertainty measures show comparable results that are closer to ideal performance than random. On the lesion scale, the performance of log-sum-based measures (the only ones highly dependent on lesion size) is the poorest in terms of $\widehat{F1-AUC}$. The best performance is given by the proposed DDU^{true} with $\widehat{F1-AUC}$ of 0.4265 ± 0.0294 , meaning the best ability to capture lesion detection errors. The performance of Mean EoE, DDU and Mean NC measures is close to DDU^{true} , but it is statistically significantly different at significance level of 0.05 according to one-sided paired Wilcoxon tests that show p-values of 0.007, 0.013 and 0.003, respectively.

4. DISCUSSION

In this work, we propose novel lesion-scale uncertainty measures that, compared to previously proposed ones, do not utilise voxel-scale uncertainties and are defined through the disagreement between ensemble members in structural predictions. Differently from previous works exploring uncertainty in MS, we used error retention curves to evaluate a range of uncertainty measures in their ability to capture model errors on both voxel and lesion scales. The comparison between uncertainty measures was performed in terms of their ability to capture overall segmentation or lesion detection errors via DSC-RC and lesion F1-RC, respectively.

We observe that a comparatively good performance of an uncertainty measure on the voxel scale does not mean that it will be useful for pinpointing lesion detection errors. For example, while ExE (data uncertainty) and EoE (total uncertainty) measures have a similar and superior performance in terms of $\widehat{DSC-AUC}$, the Mean ExE is the least informative for the lesion detection across mean-based lesion-scale measures. Additionally, we notice that a previously proposed log-sum aggregation strategy for lesion-scale measures construction yields one of the lowest $\widehat{F1-AUC}$ s. In contrast to lesion-scale measures based on voxel uncertainties, the proposed DDU^{true} shows significantly better performance in capturing lesion detection errors, measured by $\widehat{F1-AUC}$.

The proposed framework can be a step forward in integrating AI in a clinical practice considering a semi-automatic scenario where clinicians assess automated WML predictions and make necessary corrections. Although we identified uncertainty measures informative in terms of model errors of different kind, it is yet important to verify in practice if introducing uncertainty maps to clinicians can speed up or simplify the correction process of predicted WML masks.

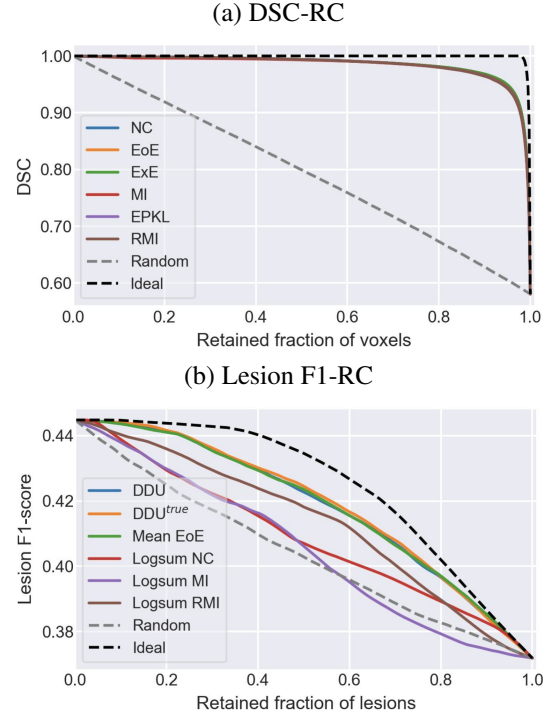


Fig. 2. Average across the test set DSC (a) and lesion F1 (b) retention curves for evaluation of uncertainty measures on voxel and lesion scales, respectively. Lesion F1-RC are made for the three best (Mean EoE, DDU, Mean NC) and the three worst (Logsum RMI, NC, MI) measures in terms of $\widehat{F1-AUC}$.

(a) $\widehat{DSC-AUC} * 100(\uparrow)$		(b) Lesion $\widehat{F1-AUC} * 100(\uparrow)$	
Ideal	99.89 ± 0.01	Ideal	43.22 ± 2.99
ExE	98.53 ± 0.20	DDU^{true} (our)	42.65 ± 2.94
EoE	98.53 ± 0.20	Mean EoE	42.60 ± 2.94
NC	98.50 ± 0.20	DDU (our)	42.59 ± 2.94
EPKL	98.44 ± 0.21	Mean NC	42.59 ± 2.94
RMI	98.41 ± 0.21	Mean MI	42.43 ± 2.92
MI	98.37 ± 0.21	Mean EPKL	42.42 ± 2.92
Random	79.66 ± 1.46	Logsum EPKL	42.40 ± 2.92
		Mean RMI	42.38 ± 2.92
		Logsum EoE	42.37 ± 2.92
		Mean ExE	42.25 ± 2.90
		Logsum ExE	42.22 ± 2.90
		Logsum RMI	42.10 ± 2.89
		Logsum NC	41.64 ± 2.86
		Logsum MI	41.26 ± 2.79
		Random	41.13 ± 2.78

Table 2. Average across the test set areas under DSC-RC (a) and Lesion F1-RC (b) measuring the correspondence between uncertainty measures and model errors in segmentation and lesion detection, respectively. Standard errors are computed using bootstrapping with the sample size of 50 patients for 10^4 repetitions.

5. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using human subject data made available in open access by the Shifts project [12]. Ethical approval was not required as confirmed by the license attached with the open data.

6. ACKNOWLEDGMENTS

This work was supported by the Hasler Foundation Responsible AI programme (MSxplain) and the EU Horizon 2020 project AI4Media (grant 951911). We acknowledge access to the facilities and expertise of the CIBM Center for Biomedical Imaging, a Swiss research center of excellence founded and supported by Lausanne University Hospital (CHUV), University of Lausanne (UNIL), École polytechnique fédérale de Lausanne (EPFL), University of Geneva (UNIGE) and Geneva University Hospitals (HUG).

7. REFERENCES

- [1] C.C. Hemond and R. Bakshi, “Magnetic resonance imaging in multiple sclerosis,” *Cold Spring Harbor Perspectives in Medicine*, vol. 8, no. 5, May 2018.
- [2] A. et al. Thompson, “Diagnosis of multiple sclerosis: 2017 revisions of the mcdonald criteria,” *The Lancet Neurology*, vol. 17, no. 2, pp. 162–173, Feb 2018.
- [3] K. Bendfeldt et al., “Association of regional gray matter volume loss and progression of white matter lesions in multiple sclerosis - A longitudinal voxel-based morphometry study,” *NeuroImage*, vol. 45, no. 1, pp. 60–67, Mar. 2009.
- [4] Fisniku, L. K. et al., “Disability and t2 mri lesions: a 20-year follow-up of patients with relapse onset of multiple sclerosis,” *Brain*, vol. 131, no. 3, pp. 808–817, Feb. 2008.
- [5] Chenyi Zeng, Lin Gu, Zhenzhong Liu, and Shen Zhao, “Review of deep learning approaches for the segmentation of multiple sclerosis lesions on brain mri,” *Frontiers in Neuroinformatics*, vol. 14, 2020.
- [6] T. Nair, D. Precup, D.L. Arnold, and T. Arbel, “Exploring uncertainty measures in deep networks for multiple sclerosis lesion detection and segmentation,” *Medical Image Analysis*, vol. 59, pp. 101557, January 2020.
- [7] R. McKinley et al., “Automatic detection of lesion load change in multiple sclerosis using convolutional neural networks with segmentation confidence,” *NeuroImage: Clinical*, vol. 25, pp. 102104, 2020.
- [8] B. Lambert, F. Forbes, S. Doyle, A. Tucholka, and M. Dojat, “Fast Uncertainty Quantification for Deep Learning-based MR Brain Segmentation,” in *EGC 2022 - Conférence francophone pour l’Extraction et la Gestion des Connaissances*, Blois, France, Jan. 2022, pp. 1–12.
- [9] B. Lambert, F. Forbes, A. Tucholka, S. Doyle, and M. Dojat, “Multi-Scale Evaluation of Uncertainty Quantification Techniques for Deep Learning based MRI Segmentation,” in *ISMRM-ESMRMB & ISMRT 2022 - 31st Joint Annual Meeting International Society for Magnetic Resonance in Medicine*, London, United Kingdom, May 2022, pp. 1–3.
- [10] B. Lambert, F. Forbes, S. Doyle, A. Tucholka, and M. Dojat, “Beyond voxel prediction uncertainty: Identifying brain lesions you can trust,” *International Workshop on Interpretability of Machine Intelligence in Medical Image Computing*, pp. 61–70, 2022.
- [11] A. Malinin, *Uncertainty estimation in deep learning with application to spoken language assessment*, Ph.D. thesis, University of Cambridge, United Kingdom, 2019.
- [12] A. Malinin et al., “Shifts 2.0: Extending The Dataset of Real Distributional Shifts,” *arXiv:2206.15407*, 2022.
- [13] A. Malinin and M. Gales, “Uncertainty estimation in autoregressive structured prediction,” in *International Conference on Learning Representations*, 2021.
- [14] R. Mehta et al., “Qu-brats: Miccai brats 2020 challenge on quantifying uncertainty in brain tumor segmentation – analysis of ranking scores and benchmarking results,” *Machine Learning for Biomedical Imaging*, vol. 1, 2022.
- [15] O. Commowick et al., “Objective Evaluation of Multiple Sclerosis Lesion Segmentation using a Data Management and Processing Infrastructure,” *Scientific Reports*, vol. 8, no. 1, pp. 13650, Sept. 2018.
- [16] A. Carass et al., “Longitudinal multiple sclerosis lesion segmentation: Resource and challenge,” *NeuroImage*, vol. 148, pp. 77–102, 2017.
- [17] F. La Rosa et al., “Multiple sclerosis cortical and wm lesion segmentation at 3t mri: a deep learning method based on flair and mp2rage,” *NeuroImage: Clinical*, vol. 27, pp. 102335, 2020.
- [18] C. Gramsch et al., “Diagnostic value of 3d fluid attenuated inversion recovery sequence in multiple sclerosis,” *Acta Radiologica*, vol. 56, no. 5, pp. 622–627, may 2015.