

Journal Pre-proof

A comparative study of deep learning for cortical lesion MRI segmentation with explainability analysis in multiple sclerosis

Nataliia Molchanova, Alessandro Cagol, Mario Ocampo–Pineda, Po–Jui Lu, Matthias Weigel, Xinjie Chen, Erin S. Beck, Charidimos Tsagkas, Daniel S. Reich, Colin Vanden Bulcke, Anna Stölting, Serena Borrelli, Pietro Maggi, Sebastian Baez Lugo, Delphine Ribes, Adrien Depeursinge, Cristina Granziera, Henning Müller, Pedro M. Gordaliza, Meritxell Bach Cuadra



PII: S2213-1582(26)00066-5
DOI: <https://doi.org/10.1016/j.nicl.2026.104007>
Reference: YNICL 104007

To appear in: *NeuroImage: Clinical*

Received date: 18 December 2025

Revised date: 12 May 2026

Accepted date: 17 May 2026

Please cite this article as: N. Molchanova, A. Cagol, M. Ocampo–Pineda et al., A comparative study of deep learning for cortical lesion MRI segmentation with explainability analysis in multiple sclerosis. *NeuroImage: Clinical* (2026), doi: <https://doi.org/10.1016/j.nicl.2026.104007>.

This is a PDF of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability. This version will undergo additional copyediting, typesetting and review before it is published in its final form. As such, this version is no longer the Accepted Manuscript, but it is not yet the definitive Version of Record; we are providing this early version to give early visibility of the article. Please note that Elsevier's sharing policy for the Published Journal Article applies to this version, see: <https://www.elsevier.com/about/policies-and-standards/sharing#4-published-journal-article>. Please also note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2026 Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Highlights

A Comparative Study of Deep Learning for Cortical Lesion MRI Segmentation with Explainability Analysis in Multiple Sclerosis

Nataliia Molchanova, Alessandro Cagol, Mario Ocampo-Pineda, Po-Jui Lu, Matthias Weigel, Xinjie Chen, Erin S. Beck, Charidimos Tsagkas, Daniel S. Reich, Colin Vanden Bulcke, Anna Stölting, Serena Borrelli, Pietro Maggi, Sebastian Baez Lugo, Delphine Ribes, Adrien Depeursinge, Cristina Granziera, Henning Müller, Pedro M. Gordaliza, Meritxell Bach Cuadra

- nnU-Net enables promising cortical lesion segmentation across a large multi-center dataset.
- Comparative study evaluates nnU-Net configurations; improvements are modest and mostly non-significant.
- Expert review validates clinical relevance of false-positive predicted cortical lesions.
- Bottleneck feature analysis explores U-Net generalization limits across clinical centers.
- Study highlights practical barriers hindering clinical integration of CL segmentation.

A Comparative Study of Deep Learning for Cortical Lesion MRI Segmentation with Explainability Analysis in Multiple Sclerosis

Nataliia Molchanova^{a,b,c,d,*}, Alessandro Cagol^{f,g,h,j}, Mario Ocampo-Pineda^{f,g,h}, Po-Jui Lu^{f,g,h}, Matthias Weigel^{f,g,h,i}, Xinjie Chen^{f,g,h}, Erin S. Beck^{k,l}, Charidimos Tsagkas^l, Daniel S. Reich^l, Colin Vanden Bulcke^m, Anna Stölting^m, Serena Borrelli^{m,1}, Pietro Maggi^m, Sebastian Baez Lugo^o, Delphine Ribes^o, Adrien Depeursinge^c, Cristina Granziera^{f,g,h}, Henning Müller^{c,e}, Pedro M. Gordaliza^{d,b,a,**}, Meritxell Bach Cuadra^{d,b,a,**}

^a*Faculty of Biology and Medicine, University of Lausanne (UNIL), Lausanne, Switzerland*

^b*Radiology Department, Lausanne University Hospital (CHUV), Lausanne, Switzerland*

^c*MedGIFT, Institute of Informatics, School of Management, HES-SO Valais-Wallis University of Applied Sciences and Arts Western Switzerland, Sierre, Switzerland*

^d*CIBM Center for Biomedical Imaging, Lausanne, Switzerland*

^e*Department of Radiology and Medical Informatics, University of Geneva, Geneva, Switzerland*

^f*Translational Imaging in Neurology (ThINK) Basel, Department of Medicine and Biomedical Engineering, University Hospital Basel and University of Basel, Basel, Switzerland*

^g*Multiple Sclerosis Center, Department of Neurology, University Hospital Basel, Basel, Switzerland*

^h*Research Center for Clinical Neuroimmunology and Neuroscience Basel (RC2NB), University Hospital Basel and University of Basel, Basel, Switzerland*

ⁱ*Division of Radiological Physics, Department of Radiology, University Hospital Basel, Basel, Switzerland*

^j*Dipartimento di Scienze della Salute, Università degli Studi di Genova, Genova, Italy*

^k*Department of Neurology, Icahn School of Medicine at Mount Sinai, New York City, USA*

^l*Translational Neuroradiology Section, National Institute of Neurological Disorders and Stroke, National Institutes of Health, Bethesda, USA*

^m*Neuroinflammation Imaging Lab (NIL), Université catholique de Louvain, Brussels, Belgium*

ⁿ*Department of Neurology, Hôpital Erasme, Hôpital Universitaire de Bruxelles, Université libre de Bruxelles, Brussels, Belgium*

*Corresponding author.

**Equal contributions.

Email address: nataliia[dot]molchanova[at]unil[dot]ch (Nataliia Molchanova)

*°EPFL+ECAL Lab, École polytechnique fédérale de Lausanne
(EPFL), Lausanne, Switzerland*

Abstract

Cortical lesions (CLs) have emerged as valuable biomarkers in multiple sclerosis (MS), offering high diagnostic specificity and prognostic relevance. However, their routine clinical integration remains limited due to subtle magnetic resonance imaging (MRI) appearance, challenges in expert annotation, and a lack of standardized automated methods. We present a multi-centric comparative study of CL detection and segmentation in MRI. A total of 656 MRI scans, including clinical trial and research data from four institutions, were acquired at 3T and 7T using MP2RAGE and MPRAGE sequences with expert-consensus annotations. We rely on the self-configuring nnU-Net framework, designed for medical imaging segmentation, and propose adaptations tailored to the improved CL detection. We evaluated model generalization through out-of-distribution testing, demonstrating promising lesion detection capabilities with an F1-score of 0.64 and 0.5 in and out of the domain, respectively. We also analyze internal model features and model errors for a better understanding of AI decision-making. Our study examines how data variability, lesion ambiguity, and protocol differences impact model performance, offering future recommendations to address these barriers to clinical adoption. Furthermore, we designed and implemented a medical expert questionnaire for better assessment of clinical value of the model predictions. To reinforce the reproducibility, the implementation and models will be publicly accessible and ready to use at [GitHub](#) and [Zenodo](#).

Keywords:

Multiple sclerosis, Cortical lesions, Segmentation, Detection, Magnetic Resonance Imaging, Brain, Deep learning, Trustworthy AI

1. Introduction

Multiple sclerosis (MS) is a chronic inflammatory and neurodegenerative disease of the central nervous system, affecting over 2 million people worldwide [1, 2]. Magnetic resonance imaging (MRI) is central to the diagnosis and monitoring of MS, enabling the detection of characteristic brain

lesions [3, 4, 5, 6]. White matter lesions (WMLs) have served as the primary radiological hallmark, guiding diagnosis and treatment decisions [3]. In recent years, additional imaging biomarkers—such as cortical lesions (CLs)[7, 8], paramagnetic rim lesions [9, 10], and the central vein sign [11, 8] — have gained recognition for their ability to support the diagnosis or to capture disease activity more comprehensively. CLs are of clinical interest due to their association with cognitive impairment, progressive disability, and early-stage onset [7, 12, 8]. Indeed, CLs have diagnostic value and were incorporated in the 2017 revision of the *McDonald criteria* due to their high specificity for MS [13]. Unlike WMLs that appear in numerous neuroinflammatory and neurovascular conditions, the presence of a CLs enhances diagnostic certainty for MS and helps differentiate it from clinical conditions mimicking MS [13, 8]. This distinction is crucial for preventing misdiagnosis and inappropriate treatment, particularly as the McDonald criteria generally favor sensitivity over specificity. Yet, despite their clinical relevance, CLs (as distinct from juxtacortical lesions) remain underused in routine practice due to several technical barriers: their small size (e.g., several mm intracortical lesions), low contrast on standard FLAIR or T1w MRI sequences, and the need for expert radiological interpretation [6]. Different MS lesion types are illustrated in Figure 1.




Figure 1: Types of MS lesions categorized by their location with respect to the cortex. Lesions appear on MP2RAGE MRI as hypointense regions (white arrow pointers). CL types are underlined, the rest are WMLs. Leukocortical (CL) and juxtacortical (WML) are sometimes difficult to distinguish due to unclear cortical involvement. CSF - cerebrospinal fluid, GM - gray matter, WM - white matter.

Automated segmentation methods offer a promising avenue to facilitate the adoption of CL assessment into clinical workflows. They can reduce the

burden of manual annotation, increase reproducibility, and enable broader access to quantitative lesion metrics. However, the development and deployment of such tools are hindered by several caveats: limited annotated datasets, scanner and protocol variability, and poor generalization across sites. CL segmentation, in particular, is further challenged by *high class imbalance* and substantial heterogeneity in lesion appearance. These limitations have slowed the translation of automated CL assessment tools into clinical use.

Deep learning (DL) approaches for CL segmentation have been limited by the scarcity of adequately annotated training data. Consequently, initial DL approaches employed joint segmentation strategies, where CLs and WMLs were treated as a single lesion class without distinction. Concretely, in [14], authors proposed a shallow 3D U-Net model [15, 16] for this joint segmentation using fluid-attenuated inversion recovery (FLAIR) and magnetization prepared 2 rapid gradient echo (MP2RAGE) [17] sequences. While this approach reported improved detection compared to previous non-DL methods [18], the joint classification inherently inflated CL detection performance, as predictions in juxtacortical white matter regions—which would constitute false positives in CL-specific segmentation—were correctly classified as WMLs, thereby avoiding penalization despite anatomical imprecision.

Dedicated CL segmentation with DL was first explored using ultra-high-field (7T) MRI, which enhances the visibility of CLs [19]. The CLAIMS framework [20] demonstrated substantial improvements in CL-specific detection, particularly for leukocortical and subpial lesions. Crucially, this approach required explicit CL identification, meaning any detections in white matter regions (including juxtacortical areas) were counted as false positives rather than alternative lesion types. When accounting for this methodological difference, the CL detection performance was comparable to previous joint segmentation approaches, but with significantly improved specificity for cortical tissue. Despite its effectiveness, the application remains limited due to the low availability of 7T scanners in typical clinical settings and modest inter-rater reliability even at these high fields.

Further advances aimed at leveraging enhanced contrast visibility on conventional scanners include combining multiple advanced contrasts or synthetic sequences. Reference [21] investigated the impact of fluid and white matter suppression (FLAWS) alongside MP2RAGE using the 3D nnU-Net [22] framework. This CL-specific approach at 3T achieved detection results comparable to earlier methods while maintaining the anatomical precision re-

quired for clinical CL assessment. However, its clinical translation is hindered by the requirement for FLAWS sequences not routinely implemented in clinical protocols, potentially limiting widespread adoption. Similarly, [23] evaluated retrospective segmentation methods leveraging legacy clinical trial data, introducing novel contrast combinations such as FLAIR², T1/T2 ratio, and AI-derived synthetic double inversion recovery (DIR) contrasts. Their findings highlighted improved CL detection when using synthetic MRI-derived contrast maps and emphasized the persistent challenges related to variability across scanner platforms and clinical settings.

While these studies demonstrate significant progress in automated CLs segmentation, a comprehensive evaluation of different DL architectures and their comparative performance remains unexplored. Most works have focused on optimal input contrast combinations, leaving questions about model architecture selection, generalizability across clinical settings, and failure case analysis largely unaddressed.

Furthermore, clinical deployment of DL models requires robust performance across diverse imaging protocols, scanner vendors, and field strengths that differ significantly from controlled research environments. Previous generalizability assessments have been limited in scope, with studies testing on relatively small out-of-domain cohorts (36 subjects in [14], 20 subjects in [20]) using similar scanner types and protocols. Such evaluations may not adequately reflect the heterogeneity encountered in real-world clinical settings, where scanner vendors, acquisition parameters, and patient populations vary substantially. Additionally, the complex nature of CL detection, where subtle anatomical distinctions determine diagnostic relevance, necessitates comprehensive error analysis to understand model behavior and failure modes—an aspect that has not been systematically addressed in prior CL segmentation literature.

2. Contributions

We present a multi-centric comparative study of CL segmentation using a large multi-center dataset comprising 656 scans from 4 medical centers across different field strengths (3T/7T) and acquisition protocols. Our framework leverages commonly acquired MP2RAGE and MPRAGE sequences to ensure clinical applicability. Specifically, our contributions include:

- A systematic architectural comparison identifying effective model configurations for CL segmentation across diverse clinical settings,

- Out-of-distribution evaluation on 224 subjects across different scanner vendors, acquisition protocols, and patient cohorts—one of the most heterogeneous OOD assessments in CL segmentation to date,
- Systematic error analysis in CL segmentation, identifying failure modes and providing actionable insights for model improvement and clinical interpretation,
- Public release of model weights and implementation framework to enable widespread adoption in clinical and research environments.

3. Materials and methods

3.1. Data

Four medical centers shared their longitudinal and cross-sectional data for this study, including A) Basel University Hospital (USB), Switzerland, B) University Hospital of Lausanne (CHUV), Switzerland, C) National Institutes of Health, USA, and D) Catholic University of Louvain (UCLouvain), Belgium. The data information is provided in Table 1.

3.1.1. Manual annotation

The scans were manually annotated using local resources and annotation guidelines. In sites A-D, all lesions involving the cortex were included in the CL masks, including intracortical, leukocortical, or subpial, typically extending for at least 3mm in the longest diameter. In A, CLs were segmented on 1mm isotropic MP2RAGE through an independent assessment of two raters (neurologist and MD with five years of neuroimaging experience). In B, CLs were detected using $1 \times 1 \times 1.2\text{mm}^3$ 3D FLAIR, 3D DIR, and MP2RAGE acquired at 3T by two independent raters (an experienced neurologist and radiologist), followed by a consensus review; then, a trained technician manually delineated the consensus lesions. In C, CLs were identified and segmented on 3 and 7T images independently. 7T visible CLs were segmented on 0.5mm isotropic MP2RAGE (median of four acquisitions in the same scanning sessions) and 0.5mm multi-echo T2*w GRE by two independent raters (neurologists with two and >10 years of MS neuroimaging experience), followed by a consensus review. 3T CLs were independently segmented on 1mm isotropic 3D FLAIR and MP2RAGE by two raters (neurologists with two and four years of MS neuroimaging experience), followed by a consensus

review. In D, the detection of CLs was performed using 0.7mm isotropic 3D DIR and MP2RAGE by three independent raters (a neuroscientist and two neurologists with three, five, and 10 years of MS neuroimaging experience, respectively), followed by a consensus review; the trained neuroscientist rater performed the manual delineation of consensus lesions.

All scans from site B and 163 scans from site A (the first time point) had information about the lesions subtypes (either from MP2RAGE or FLAIR; see Figure 1) and included FLAIR-based WML segmentation masks. These WML masks were used exclusively for post-hoc error characterization (analysis of overlaps between model errors and WML regions), and were not included in model training or evaluation. Subjects from site C 7T data had information about the lesion types and segmentation of *potential lesions*. Potential lesions were the ones where expert raters could not conclude whether lesions were CL or not, and were not used during training or evaluation. The rest of the test data did not contain WML masks or lesion subtypes.

Table 1: Data description. RR - relapsing-remitting, PP - primary-progressive, SP - secondary-progressive, CIS - clinically isolated syndrome, Q2 - second quartile, IQR - interquartile range.

Medical center	A	B	C	D
Demographics				
# patients	163	43	35	112
Phenotype	RR (97), PP (22), SP (44)	RR	RR (24), PP (1), SP (9)	CIS (3), RR (44), PP (8), SP (17), MS-mimic (40)
% female	60%	58%	63%	70%
EDSS Q2 (IQR)	3 (1.5-4.5)	1.5 (1.5-4)	2.5 (1-4)	2.5 (2-4.5)
Timepoints	2	2 (2 year distance)	3	1
MRI				
Scanner	Magnetom Prisma, Siemens Healthineers	Magnetom Trio, Healthineers	3T Skyra and 7T whole-body research system, Siemens Healthcare	SIGNA™, GE Health Care
Field strength	3T	3T	3T, 7T	3T
MPRAGE acquisition parameters (TR/TE/TI, ms)	-	1.0 × 1.0 × 1.2, 2300/2.84/900	-	1 × 1 × 1, 2186/3/900
MP2RAGE acquisition parameters (TR/TE/TI1/TI2, ms)	1.0 × 1.0 × 1.0, 5000/2.98/700/2500	1.0 × 1.0 × 1.0, 5000/2.84/700/2500	1.2, for 3T: 0.8 × 0.8 × 0.8, 5000/2.9/700/2500, for 7T: 0.5 × 0.5 × 0.5, 6000/5,800/2700	1 × 1 × 1, 5000/3/700/2500
Processing				
Registration	-	-	-	Registration to 3D EPI with 0.67x0.67x0.67mm
Skull-stripping	HD-BET [24] to FLAIR + morphological dilation (2 iterations, 3x3x3 structural element)	HD-BET [24] to FLAIR + morphological dilation (2 iterations, 3x3x3 structural element)	SynthSeg [25] to MP2RAGE + manual correction for 7T	HD-BET to MPRAGE + morphological dilation

3.1.2. Data splitting

The data was split into training (Train), in-domain pure testing (Test-in), and out-of-domain (OOD) testing (Test-out). The data from site D was used for the OOD testing, due to the difference in the scanner manufacturer, compared to the other centers. The medical center possesses the clinically used MPRAGE modality (in addition to the DIR and MMP2RAGE used for lesion annotation) and data for both MS and clinical conditions mimicking MS (MS-mimic) patients. Thus, we can test the applicability of the proposed framework to support MS differential diagnosis. The rest of the data A-C were separated on Train and Test-in by using a custom stratified split approach that preserved both subject-level grouping and data distribution characteristics. The method first aggregated data at the subject-site level and then applied stratification techniques while ensuring subjects from the same site remained together in the same split. The implementation maintained similar distributions of site, lesion count, and total lesion volume between splits by creating composite stratification categories. The continuous variables (number of lesions and total volume) were discretized into 5 quantile-based bins, and single-sample stratification groups were merged with their nearest neighbors based on similarity in lesion metrics while respecting site constraints. This approach allocated 80% of the in-domain data to training and 20% to testing, with reproducibility ensured through fixed random seeding. Distribution statistics were verified through a comprehensive comparative analysis that examined the percentage distributions of categorical variables across splits, assessed descriptive statistics (mean, standard deviation, quartiles) of the continuous variables in each partition, and confirmed the preservation of subject-site integrity by counting unique subject-site combinations in the resulting datasets.

The preprocessing pipelines differ across sites, reflecting the practical constraints of a real-world multi-center study: data sharing agreements, local ethics requirements, and anonymization procedures vary by institution, meaning that available sequences and pre-applied processing steps differ. Where skull-stripping was not already applied, we used the tool most appropriate for each acquisition type: HD-BET [24] for standard 3T data at sites A, B, and D, and SynthSeg [25] for the 7T MP2RAGE data at site C, where the intensity distribution differs substantially from 3T. Site D additionally required registration to a 3D EPI reference space to co-register its multi-contrast acquisitions. Overall, the differences between A-D sites

are driven by varying acquisition and annotation protocols, as well as patient cohorts. The differences in the distributions of lesions across different sites and domain are shown in Figure 2. Sites A, B, and C have patients with different total lesion volumes and numbers of lesions per patient; however, this distribution is skewed towards smaller lesion volumes and counts. Site C, containing mostly ultra-high-field 7T data, has higher total lesion volumes and counts per patient since the visibility of CLs on 7T increased sensitivity for CL detection [26]. Moreover, the annotations performed on 7T MP2RAGE and T2-star-weighted (T2*w) GRE images resulted in the detection of many subpial lesions.

Table 2: Details of the data split. Modality 1 - MPRAGE, 2 - MP2RAGE.

Set name	Partition			Site (Number of Scans)			
	Modality	7T	MS-mimic	A	B	C	D
Train	2			191	60	5	0
	2	✓		0	0	31	0
	1			0	60	0	0
Test-in	2			50	14	1	0
	2	✓		0	0	6	0
	1			0	14	0	0
Test-out	2			0	0	0	72
	2		✓	0	0	0	40
	1			0	0	0	72
	1		✓	0	0	0	40
Total				241	148	43	224

3.2. Comparative study

We employed the nnU-Net framework [24, 27], a fully automated, self-configuring DL method widely regarded as a state-of-the-art approach for biomedical image segmentation tasks. While nnU-Net provides automated optimization for key segmentation pipeline components, it allows for user modifications at critical stages. nnU-Net automatically optimizes the entire segmentation pipeline, including preprocessing, U-Net network architecture configuration, training procedure, inference, and post-processing, based solely on dataset characteristics. The framework’s default process includes intensity normalization, resampling to a heuristically configured target voxel spacing (resolution adjustment), cropping or padding to standardized patch sizes, and data augmentation (e.g., rotation, scaling, and elastic deformations) to enhance robustness.

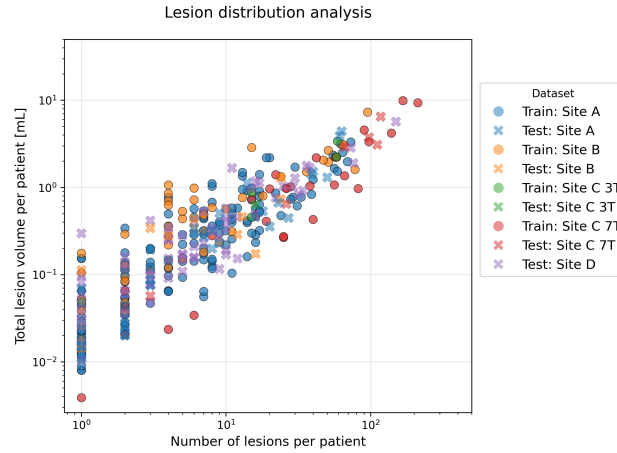


Figure 2: Distribution of total lesion volume in milliliters and number of lesions per patient across different medical sites.

Overall, the self-adaptive nature aims at nnU-Net’s high segmentation accuracy and robust generalization across diverse clinical imaging applications. Moreover, previous work on CL segmentation used 3D U-Net variations [14, 20] or nnU-Net [21]. The default configuration of the framework employs a dynamically configured 3D U-Net architecture with an encoder-decoder structure and skip connections. Training typically uses a combined Dice and binary cross-entropy (BCE) loss function, stochastic gradient descent with momentum, and polynomial learning rate scheduling during 1000 epochs. The framework implements five-fold cross-validation, providing predictions as an ensemble of all folds. nnU-Net also offers automated post-processing that removes small isolated components below a threshold volume.

In our systematic comparative experiments, we modified several key aspects of the default nnU-Net implementation to address the specific challenges of CL segmentation, like high class imbalance and the varied sizes of instances:

Optimization Strategy. We substituted the default stochastic gradient descent (SGD) with the Adam optimizer [28] using a learning rate of 3×10^{-4} . CL segmentation presents a strong class-imbalance problem: true CL voxels

constitute a tiny fraction of the image volume, resulting in sparse and weak gradient signals. SGD applies a uniform learning rate to all parameters, which causes the useful gradient signals from rare CL voxels to be dominated by the abundant background, making convergence extremely slow or unstable for this task. We confirmed this empirically: SGD-based training showed no meaningful reduction in loss across epochs. Adam’s per-parameter adaptive moment estimates address this directly by tracking gradient variance independently for each parameter, allowing larger effective steps along sparse gradient directions while dampening noisy ones. The default nnU-Net SGD configuration is calibrated for large, well-represented structures such as organs, and is not well suited to the highly imbalanced CL segmentation task. Tuning SGD momentum and schedule for this task would require an expensive grid search (training takes approximately 2 days per run). The learning rate of 3×10^{-4} is a widely adopted default for Adam across a broad range of deep learning tasks; Adam’s adaptive scaling means that values within roughly an order of magnitude perform similarly, and our experiments confirmed stable convergence at this value.

Network Architecture. We systematically evaluated multiple architectural variants to address the unique challenges of CL segmentation. Beyond the baseline vanilla U-Net configuration provided by nnU-Net, we tested the recently introduced residual encoder variant that enhances gradient flow through deeper network layers. This residual encoder architecture is available in three configurations (designated "M", "L", and "XL") with progressively larger patch and batch sizes, allowing us to examine the trade-off between computational requirements and the ability to capture broader contextual information relevant to CL identification [27]. Furthermore, we explored the application of the U-Mamba model layers [29] to the segmentation architecture. This approach was motivated by U-Mamba’s theoretical advantages in modeling long-range spatial dependencies, which we hypothesized would better capture the relationship between lesion appearance and surrounding cortical anatomy. We tested two Mamba integration strategies: a conservative approach with Mamba layers only at the bottleneck, and a more extensive implementation with Mamba layers throughout the encoder pathway. However, both U-Mamba configurations encountered significant convergence failures during training that persisted across multiple parameter settings, and were therefore excluded from the comparative analysis.

Loss Function. Beyond the default combination of Dice and cross-entropy losses, we evaluated its blob loss variant [30] to better capture the morphological characteristics of small CLs.

Resolution Standardization. As an experimental condition (Experiment 3), we investigated whether upsampling all images to match the native resolution of the 7T site C data (0.5mm^3) prior to nnU-Net preprocessing could improve lesion detection. Images were upsampled using *SimpleITK* with *b-spline* interpolation before being passed to the nnU-Net pipeline, so that its heuristic-based target resolution selection would converge on 0.5mm^3 . The rationale was to test whether preserving the fine structural detail available in the site C data could benefit detection of subtle cortical boundaries across all sites. It is important to note that this resampling strategy was evaluated in isolation as a single experimental ablation; it was not part of the standard pipeline used in all other experiments. Resampling to 0.5mm^3 isotropic resolution increases the number of voxels by a factor of ~ 8 compared to 1mm^3 isotropic data, substantially increasing GPU memory requirements and training time.

Postprocessing. We found that the default postprocessing, removing small instances, is suboptimal for segmenting CL, as the small size of many CLs (e.g., a few voxels) meant that true positive lesions were being inappropriately removed.

We systematically evaluated framework modifications, including architectures, losses, and preprocessing, to quantify the improvement for CL detection on a dataset with diverse acquisition protocols and field strengths.

3.3. Evaluation

3.3.1. Metrics

We use a standard [31, 21] evaluation pipeline, including the assessment of overall segmentation employing Dice similarity score (DSC), normalized DSC (nDSC) [32], and lesion detection quality (F1-score, Precision, and Recall).

Let TP, FP, FN be the number of true positive, false positive, and false negative voxel predictions. Then, the segmentation quality measures are defined as follows:

$$DSC = \frac{2TP}{2TP + FP + FN},$$

$$nDSC = \frac{2TP}{2TP + \kappa \cdot FP + FN}, \quad \kappa = h(r^{-1} - 1),$$

where h is the ratio between the positive and the negative classes in the ground truth scan segmentation, $r \in (0, 1)$ is the reference value set to the mean fraction of the positive class. The nDSC accounts for the difference in the positive class load across different scans, addressing the class-load bias in the DSC.

Let TPL , FPL , and FNL be the number of true-positive, false-positive, and false-negative lesions. The lesions were classified on TPL, FPL, and FNL through the connected components analysis with 26 connectivity. A predicted connected component is a TPL if it has a one-voxel overlap with the ground truth; FPL if no overlap. A ground truth connected component is an FNL if it has no overlap with the prediction. Then, the detection quality measures are defined as follows:

$$F_1\text{-score} = \frac{2TPL}{2TPL + FPL + FNL},$$

$$\text{Precision} = \frac{TPL}{TPL + FPL}, \quad \text{Recall} = \frac{TPL}{TPL + FNL}.$$

To assess the statistical significance of the differences in the medians of quality metrics for different models, two-sided paired Wilcoxon tests on ranks were used. The p-values from multiple Wilcoxon tests were corrected using a Benjamini-Hochberg procedure for controlling the false discovery rate (FDR) with a 0.05 error rate.

3.3.2. Detailed model performance analysis

During the comparative study we used the aforementioned metrics to select the optimal model. We then performed a detailed analysis to explain the model's decisions.

Characterization of lesion-detection errors. Sites A, B, and C manual annotations included different lesion subtypes (leukocortical, subpial, WML, etc.). We used this information to better characterize the model errors. For this, we analyzed the overlaps between TPL, FPL, and FNL masks with labeled cortical and WM masks to obtain lesion classification. Additionally, we examined the relationship between lesion volume and errors.

Bottleneck features analysis. To gain insight into the internal representations learned by the network and to assess whether the model captures clinically or technically relevant variability, we performed an analysis of internal model representations on the bottleneck features extracted from the trained segmentation model. These features, located at the network’s deepest layer, encode compressed information before upsampling and prediction. This analysis enables the exploration of the influence of data characteristics such as acquisition site, imaging modality, time point, or lesion burden on the learned feature space. Understanding these relationships can help uncover potential biases in the model and provide interpretability for its behavior across diverse data sources. Such information might be crucial given the diversity of the training data used in this study.

Specifically, we focused on the best-performing model and extracted its bottleneck features from the training dataset during inference. Since the model is an ensemble, the extracted nnU-Net bottleneck features were concatenated across folds. Then, the bottleneck tensors were five-dimensional, shaped as (batch, channel, height, width, depth), where the batch dimension corresponds to augmented versions of the same input scan generated through test-time augmentation (TTA). The corresponding bottleneck features are un-flipped along the input flip axes prior to averaging, so that all augmented versions are realigned to the canonical orientation; this mirrors the un-flip applied to the segmentation predictions. An alternative aggregation strategy would be to apply global average pooling to the features before averaging across TTA versions, yielding a fully orientation-invariant global descriptor. To obtain the required (samples, features) format for dimensionality reduction, we averaged across the batch dimension and flattened the remaining spatial and channel dimensions, yielding a compact feature vector for each input scan.

We then applied dimensionality reduction to project these feature vectors into a three-dimensional space for visualization and further analysis. Two complementary techniques were employed: principal component analysis (PCA) and uniform manifold approximation and projection (UMAP). PCA is a linear method that identifies orthogonal directions of maximal variance in the data, allowing for efficient compression while preserving global structure [33]. In contrast, UMAP is a nonlinear technique that constructs a low-dimensional embedding by preserving local neighborhood relationships, often capturing subtle nonlinear structures in the data [34]. The use of both methods allowed us to contrast global variance-driven patterns with local

manifold structures. The dimensionality reduction models were fitted exclusively on the training data and subsequently used to project the test data into the same space. This was done primarily to understand the specifics of the trained features and analyze their relationship to the testing data.

The resulting embeddings were colored and grouped according to categorical variables (such as medical center, MRI modality, and field strength). These visualizations provided a qualitative assessment of how strongly the model’s internal representation space reflects these underlying characteristics, indicating the model’s capacity to differentiate or generalize across different domains.

Medical expert feedback. For the OOD data from Site D, the model lesion predictions were reviewed by two trained annotators. One neuroscientist and one neurologist with 3 and 5 years of experience in MS neuroimaging, respectively.

The goal of this experiment was multifaceted and included testing several hypotheses:

- H1** The main objective to review FPLs and identify if they are initially missed MS lesions (either CL or WML).
- H2** Assess if MS-mimic patients have focal lesions, which can be confused with WML or CL.
- H3** Check if medical experts are less confident on FPLs and FNLs compared to TPLs and why.
- H4** Check if the perceived delineation quality differs across TPL, FPL, and FNL groups, to assess the validity of CL delineation quality metrics.

For this, we first selected all FPLs from MS patients (N=109) predicted by the best model on MP2RAGE scans – modality used for CL annotation in Site D. Then, we sampled size-matched FPLs (n=50) and TPLs (n=50) from the ground truth masks. We added FPLs predicted in MS-mimic patients (n=18). For each region of interest (ROIs), the annotators needed to review the scan with a lesion overlay using ITK-SNAP [35]. For each ROI, they followed the guidelines and responded to the questionnaire summarized in Table 3.

Table 3: The structure of the medical expert questionnaire: questions asked for each ROI. Participants could also leave an optional comment for each ROI.

Question	Description	Test
1. ROI classification (multiple selections possible)	Classify the ROI based on the visual assessment of the specified MRI modality. If there is a chance that ROI belongs to several classes, select all secondary classes in the order of decreasing certainty. Possible classes: Intracortical MS lesion, Leukocortical MS lesion, Unclear leuko- or juxtacortical MS lesion, Juxtacortical MS lesion, Deep white matter MS lesion, Vascular lesion, Vessel, Other (custom), Unknown.	H1,H2
2. Confidence level (single selection)	Assess your confidence when selecting the ROI classification. Possible confidence levels (5-level Likert scale): Not at all confident, Slightly confident, Moderately confident, Very confident, Extremely confident.	H3
3. Reasons for decreased confidence (multiple selection possible)	Select all the reasons decreasing your confidence in the first response. Possible responses: None, Small lesion size, Unclear cortical involvement, Low image quality, Low ROI visibility, Poor segmentation quality, Other (custom).	H3
4. Delineation needs correction (yes/no)	Do you think the lesion delineation needs to be revised?	H4

Before the experiment, an introductory session was conducted with the medical experts. The experts were provided with the written guidelines and an online table form to submit their responses. All the patient IDs were anonymized to minimize the possibility of recognition of specific patients and lesions. The experts were prohibited from consulting other contrasts.

4. Results

4.1. Comparative study

4.1.1. Experiment 1: architectures comparison

The model performance for different architectures is shown in Figure 3, and the values of the metrics with standard errors are shown in Table 4.

U-Mamba architectures could not finally be included in the comparative analysis as we encountered significant convergence challenges during training that prevented reliable model fitting despite multiple parameter configurations and optimization attempts. These convergence issues persisted across both the bottleneck-only and full-encoder implementations of the U-Mamba layers, aligning with known training instabilities in state-space models [27].

Among the successfully trained architectures, the performance differences were minimal. As shown in Tables 4, the Vanilla architecture achieved the

highest normalized DSC (0.499 ± 0.038) and recall (0.632 ± 0.038) in the *in-domain* (*Test-in*) dataset. At the same time, ResEncUNet variants had insignificantly fewer false-positive lesions on the *out-of-domain* (*Test-out*) dataset. Since the differences between the models were not significant, the choice of more computationally expensive larger architectures was not justified. Thus, we used the Vanilla architecture for further experiments.

Figure 3: Radial plots comparing different architectures using quality metric means and 90% CIs. P-values from the paired Wilcoxon statistical tests with FDR correction, comparing model X with the baseline, i.e., Default resampling (BCE + Dice Loss): * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ compared to the Vanilla baseline model.

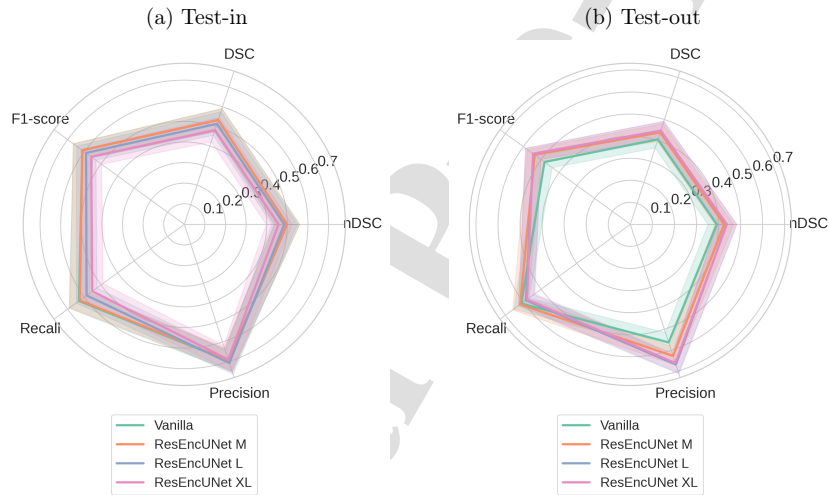


Table 4: Performance comparison of model architectures across i) Test-in and ii) Test-out domains. Values represent mean performance with standard error as subscripts. Bold values indicate the highest performance for each metric.

(a) Test-in					
Model	nDSC	DSC	F1-score	Recall	Precision
Vanilla	0.499 _{0.038}	0.532 _{0.036}	0.612 _{0.035}	0.632 _{0.038}	0.698 _{0.036}
ResEncUNet M	0.498 _{0.038}	0.533 _{0.036}	0.612 _{0.036}	0.627 _{0.038}	0.693 _{0.036}
ResEncUNet L	0.485 _{0.038}	0.513 _{0.037}	0.590 _{0.035}	0.586 _{0.038}	0.708 _{0.037}
ResEncUNet XL	0.457 _{0.039}	0.481 _{0.038}	0.559 _{0.038}	0.552 _{0.039}	0.689 _{0.041}

(b) Test-out					
Model	nDSC	DSC	F1-score	Recall	Precision
Vanilla	0.392 _{0.025}	0.405 _{0.025}	0.481 _{0.026}	0.611 _{0.026}	0.563 _{0.029}
ResEncUNet M	0.425 _{0.026}	0.439 _{0.026}	0.536 _{0.026}	0.618 _{0.025}	0.627 _{0.028}
ResEncUNet L	0.436 _{0.027}	0.447 _{0.026}	0.544 _{0.026}	0.585 _{0.026}	0.668 _{0.029}
ResEncUNet XL	0.438 _{0.027}	0.447 _{0.027}	0.545 _{0.026}	0.579 _{0.027}	0.661 _{0.028}

4.1.2. Experiment 2: Blob loss effect

The Vanilla model performance for different loss functions is illustrated in Figure 4, and the values of the metrics with standard errors are shown in Tables 5.

Blob BCE + Dice Loss showed numerically higher performance across most metrics compared to standard BCE + Dice Loss. For test-in domain, improvements included nDSC (0.505 vs 0.499), DSC (0.539 vs 0.532), F1-score (0.641 vs 0.612), and Precision (0.715 vs 0.698). Similar trends were observed out-of-domain, with the most notable gain in Precision (0.592 vs 0.563). Statistical significance was achieved for out-of-domain Precision only ($p < 0.05$); no other metric reached significance after FDR correction.

4.1.3. Experiment 3: effect of image upsampling

The model performance for different image resampling strategies is illustrated in Figure 4, and the values of the metrics with standard errors are shown in Tables 5.

Upsampling significantly improved lesion detection, increasing Recall from 0.632 to 0.748 ($p < 0.001$) for test-in domain. However, this came at the cost of substantially increased FPs, with Precision dropping from 0.698 to 0.296.

The poor precision-recall trade-off resulted in overall performance degradation, with F1-scores decreasing by approximately 45% (from 0.641 to 0.339).

Figure 4: Radial plots comparing different losses and resampling strategies using model performance metrics means and 90% CIs. P-values from the paired Wilcoxon statistical tests with FDR correction, comparing model X with the baseline, i.e., Vanilla: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ compared to baseline model.

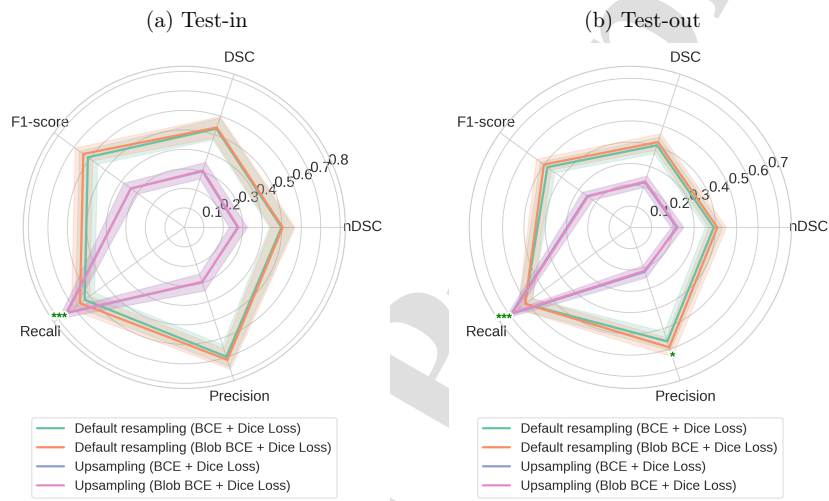


Table 5: Performance comparison of loss functions with and without upsampling across i) Test-in and ii) Test-out domains. Values represent mean performance with standard error as subscripts. Bold values indicate the highest performance for each metric.

(a) Test-in						
Loss function	Up-sampling	nDSC	DSC	F1-score	Recall	Precision
BCE + Dice	✓	0.499 _{0.038} 0.274 _{0.029}	0.532 _{0.036} 0.303 _{0.028}	0.612 _{0.035} 0.340 _{0.028}	0.632 _{0.038} 0.748 _{0.033}	0.698 _{0.036} 0.296 _{0.028}
Blob BCE + Dice	✓	0.505 _{0.038} 0.277 _{0.028}	0.539 _{0.036} 0.307 _{0.028}	0.641 _{0.034} 0.339 _{0.029}	0.661 _{0.036} 0.743 _{0.034}	0.715 _{0.034} 0.294 _{0.028}
(b) Test-out						
Loss function	Up-sampling	nDSC	DSC	F1-score	Recall	Precision
BCE + Dice	✓	0.392 _{0.025} 0.218 _{0.017}	0.405 _{0.025} 0.223 _{0.017}	0.481 _{0.026} 0.250 _{0.019}	0.611 _{0.026} 0.692 _{0.023}	0.563 _{0.029} 0.219 _{0.017}
Blob BCE + Dice	✓	0.408 _{0.026} 0.223 _{0.018}	0.421 _{0.025} 0.227 _{0.017}	0.502 _{0.026} 0.246 _{0.019}	0.611 _{0.026} 0.683 _{0.024}	0.592 _{0.029} 0.215 _{0.017}

4.2. Best model analysis

We selected the Vanilla nnU-net with Blob loss model for a more detailed performance evaluation and further explainability analysis.

4.2.1. Details on the model performance

The model performance in terms of lesion segmentation and detection across different sites, modalities, and diseases is shown in Table 6. According to the Kruskal-Wallis H-test with FDR correction, the differences between sites and sites \times modality are significant for all metrics, except for the precision: adjusted p-value=0.384 for precision and < 0.008 for other metrics. Among the MS patients, site A has the highest DSC and detection F1. Around a 30% drop in DSC and F1 was observed for site B compared to site A. Within site B, the difference between MPRAGE and MP2RAGE is marginally small and is not significant according to the U-test (adjusted p-value = 0.877 for DSC and 0.933 for F1). An additional 50% drop in DSC and F1 was observed for the site C 7T data; however, it only had 6 testing cases. FNL should largely explain this drop in performance, since the number is outstandingly high. OOD lesion detection performance was affected by a larger number of false negatives: #FNL in sites A and B was around 3

lesions per scan, in site D, 6.5 lesions per scan. We compared the OOD site D with the in-domain sites (A, B, C) performance using the Mann-Whitney U-test with FDR correction of p-values, and detected a significant difference in the nDSC, DSC, and F1 with adjusted p-values of 0.015, 0.004, and 0.007, respectively. However, the precision, recall, and FPL and FNL counts are not significantly different.

Table 6: Segmentation and lesion detection quality for the best model (Vanilla nnU-Net with Blob loss) per site and modality (1 - MPRAGE, 2 - MP2RAGE; 7T if ultra-high field, else 3T). Site D has a separation on MS and MS-mimics. Mean average and standard error were computed across subjects from Test-in (A, B, C) and Test-out (D).

(a) Segmentation quality

Site	MS-mimic	Modality	FS	n	nDSC	DSC
A		2	3T	50	0.604 _{0.046}	0.627 _{0.044}
B		1	3T	14	0.445 _{0.090}	0.473 _{0.092}
		2	3T	14	0.420 _{0.089}	0.431 _{0.090}
C		2	3T	1	0.215 _{0.000}	0.404 _{0.000}
		2	7T	6	0.068 _{0.012}	0.233 _{0.059}
D		1	3T	72	0.298 _{0.028}	0.319 _{0.029}
		2	3T	72	0.292 _{0.030}	0.311 _{0.029}
	✓	1	3T	40	0.600 _{0.078}	0.600 _{0.078}
	✓	2	3T	40	0.625 _{0.078}	0.625 _{0.078}

(b) Lesion detection quality

Site	MS-mimic	Modality	FS	n	F1-score	Recall	Precision	#FPL	#FNL
A		2	3T	50	0.742 _{0.038}	0.783 _{0.039}	0.741 _{0.041}	1.9 _{0.3}	2.7 _{0.6}
B		1	3T	14	0.536 _{0.095}	0.554 _{0.097}	0.677 _{0.103}	1.1 _{0.4}	3.7 _{1.0}
		2	3T	14	0.532 _{0.092}	0.538 _{0.090}	0.679 _{0.107}	1.0 _{0.4}	3.6 _{0.9}
C		2	3T	1	0.593 _{0.000}	0.471 _{0.000}	0.800 _{0.000}	2.0 _{0.0}	9.0 _{0.0}
		2	7T	6	0.302 _{0.062}	0.216 _{0.051}	0.645 _{0.060}	7.0 _{3.1}	55.5 _{15.9}
D		1	3T	72	0.446 _{0.038}	0.395 _{0.038}	0.586 _{0.046}	1.8 _{0.4}	6.5 _{1.3}
		2	3T	72	0.435 _{0.037}	0.395 _{0.038}	0.577 _{0.047}	1.7 _{0.4}	6.7 _{1.4}
	✓	1	3T	40	0.600 _{0.078}	-	0.600 _{0.078}	1.5 _{0.7}	-
	✓	2	3T	40	0.625 _{0.078}	-	0.625 _{0.078}	0.8 _{0.3}	-

4.2.2. Lesion-detection error analysis

The distribution of the lesion types and sizes across the TPL, FPL, and FNL groups is shown in Figure 5.

For both site A and B, we observed that the leukocortical lesions are more commonly undetected, compared to other types of lesions (65% of all FNLs from site A and 75% from site B). Also, we see that some FPLs had intersections with WML masks or leuko-/intra-cortical lesions annotated on FLAIR. Particularly, for site A timepoint 1, where the FLAIR annotations were available, 27 of 83 FPLs (33%) fell into the category of WML in WM, 6 FPLs (7%) were leukocortical under FLAIR, and 2 FPLs (2%) into CL in GM under FLAIR (i.e., intracortical). For site C (7T MP2RAGE), out of 333 FNLs, 196 were subpial lesions (59 % of all FNLs) and 95 were leukocortical lesions (29%). Out of 42 FPLs, 18 fall into the category of uncertain leuko-/juxta-cortical lesions. For all the sites, the majority of TPL are leukocortical lesions.

The distribution of lesion volumes communicates that the majority of missed lesions are small. Across sites A-C, the median FNL volume is 15mL, which is 6-18 mL less than the median TPL volume.

For a qualitative assessment, we visualized subjects with low (DSC < second quartile - Q2) versus high (DSC > third quartile - Q3) segmentation quality for different sites in Figure 6 and 7.

4.2.3. Bottleneck features analysis

The results of the nnU-Net bottleneck feature analysis are shown in Figure 8 for PCA and UMAP methods. For PCA projection, the clusters formed partially correspond to the site \times modality. The biggest cluster includes site A and site B MP2RAGE. The testing data from these sites is slightly shifted from the training data. For both PCA and UMAP, MPRAGE scans from site B and MP2RAGE scans from site C are grouped, which is due to the intensity distribution, given the relatively large distance between the MP2RAGE scans from site B of the same patients. For both PCA and UMAP, the data from site C form a distinct cluster, indicating its separation from the rest of the data. Most of the OOD test data, projected into the space of training bottleneck features, is located close to the largest cluster (sites A and B, MP2RAGE). For PCA, OOD test data has a central location for the first three components. Additional results quantifying the relationship between different continuous features, like volumetric errors, DSC, and uncertainty, with the bottleneck features are in the [Appendix A](#).

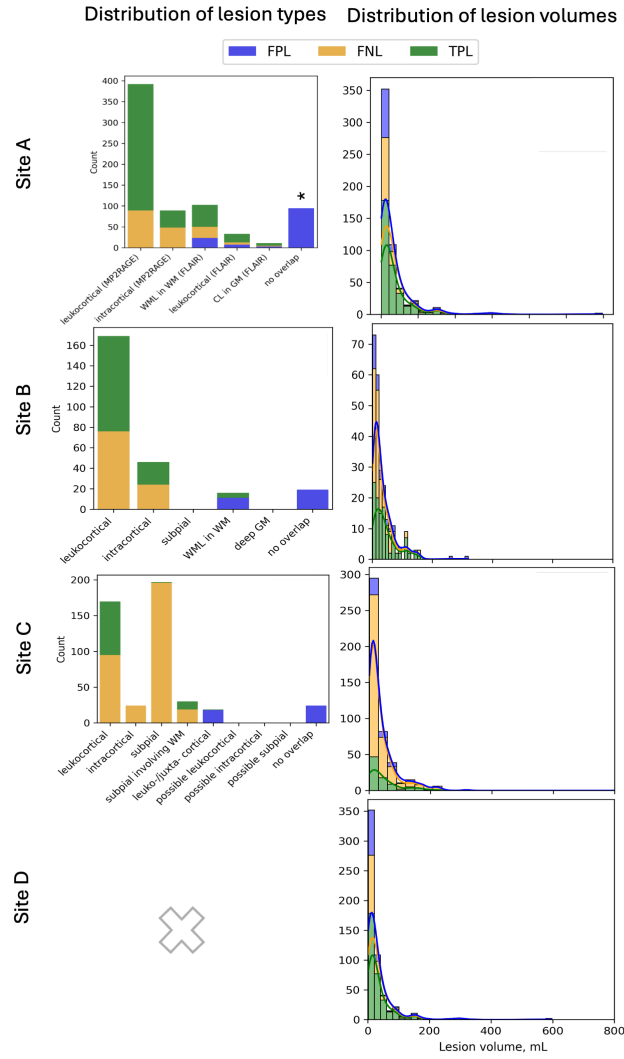


Figure 5: Distributions of lesion types and volumes across FPL, FNL, and TPL categories from sites A-D (rows). *For site A WML segmentation was only available for timepoint 1, limiting FPL typization.




Figure 6: Visualization of the manual and automatic segmentation per site for scans with low segmentation quality ($DSC < Q2$) and median lesion count per site. Images from left to right show MRI axial slices with i) no segmentation, ii) ground truth mask overlay (blue), iii) predicted mask overlay (pink).

4.2.4. Medical expert feedback

The results of the medical expert feedback are presented below, following the order of our initial hypotheses (H1 to H4). The results are reported for rater 1 and rater 2, separated by a slash (“/”) by rater order.

H1: MS lesions in FPLs. Among MS patients, medical experts classified FPLs as MS lesions in 95.4/88.1% of cases (26.6/49.5% are CLs, 61.5/22.0% are WMLs, 7.3/16.5% are unclear CL or WML). Other selected classes include unknown objects, vessels, vascular lesions, or class combinations when unclear. Figure 10 presents a bar plot with ROI classification by both raters.

H2: Similarity of MS and MS-mimic lesions. Among MS-mimic patients, 83.3/33.3% of FPLs were classified as lesions; the rest of the selected classes were vessels, vascular lesions, or a mix. Other (custom) class was selected 1/3 time(s), often suggesting sulcus presence. The inter-rater discrepancy in this category reflects differences in clinical experience with MS-mimic pathology between the two raters, combined with the small absolute number of MS-mimic FPLs ($n=18$), where a few differing classifications produce a large percentage difference.

H3: Differences in the confidence distributions. We compared the distribution of confidence of experts between different groups. Figure 9 shows the mean and standard deviation of the confidence levels across ROI types. **FNLs** showed significantly lower confidence than **TPLs** (one-sided Mann-Whitney, rater 1: $p = 0.0015$, rater 2: $p = 0.02$). The Anderson–Darling (AD) test indicated a significant overall distribution difference (P-value ≈ 0.001 , two-sided), suggesting that the groups diverge not only in central tendency but also in distribution shape. **MS FPLs & TPL** confidence levels are significantly different for rater 1 (P-value = 0.008), but not for rater 2 (P-value = 0.16) on a one-sided Mann-Whitney U-test; AD test P-value are 0.005 and 0.25, respectively. The reasons for decreased confidence are often attributed to small lesion size, unclear cortical involvement, and the need for an additional modality. Figure B.13 shows a more fine-grained analysis of the reasons for decreased confidence.

H4: delineation quality usefulness. Rater 1 showed no statistically significant differences in the proportion of lesions needing corrections across FN, TP, MS FP, and MS-mimic FP lesion groups ($\chi^2 = 5.18$, P-values = 0.16), and no pairwise comparisons remained significant after FDR correction. Rater 2 has a significant overall group effect ($\chi^2 = 8.09$, p-value =

0.044); however, no specific group pair reached significance after multiple-testing adjustment. This global signal was primarily driven by the MS-mimic FPL group, which reported zero positive cases; however, the evidence remains weak at the pairwise level.

5. Discussion

5.1. Model comparison

Our experiments demonstrated that simpler neural network architectures, specifically the vanilla nnU-Net, yielded comparable or superior CL segmentation performance compared to larger architectures, which suffered from convergence difficulties. This observation aligns with prior studies advocating for lightweight models in medical image segmentation tasks. For instance, LV-UNet and Mini-Net have shown that smaller architectures can achieve competitive performance while being more efficient and easier to train [36, 37].

The customized Blob BCE + Dice Loss modestly improved precision, particularly for OOD data, reinforcing the importance of lesion-specific loss functions in optimizing segmentation tasks involving small focal lesions. A similar outcome was previously demonstrated for the WML segmentation on FLAIR MRI [30]. However, image resampling significantly improved lesion recall yet markedly increased false-positive detections, drastically reducing overall segmentation precision. B-spline MRI upsampling might lead to an increased number of artifacts, leading to higher false-positive rates. Given a promise of high recall, one could explore cascade architectures or two-step approaches to refine such segmentations with high false positive rates.

These findings collectively suggest that effective CL segmentation relies more on appropriate training strategies and task-specific loss functions than on architectural sophistication. The limited benefit of model complexity, combined with modest improvements achievable through specialized loss functions, has important implications for clinical deployment: simpler models offer substantial practical advantages, including reduced computational requirements, lower operational costs, and decreased energy consumption—critical factors for sustainable healthcare AI implementation [38, 39].

5.2. Performance differences across sites

Substantial CL segmentation performance differences were observed across imaging sites and MRI modalities. Segmentation accuracy was highest for

MP2RAGE at 3T (sites A and B; DSC 0.47-0.63, detection F1 0.53-0.73), with slightly reduced performance for MPRAGE sequences at these same sites. Conversely, ultra-high-field (7T) data from site C displayed notably poorer performance down to 0.23 DSC and 0.29 detection F1. The error analysis clarifies the nature of this failure: of 333 false-negative lesions at site C (7T), 196 (59%) were subpial lesions and 95 (29%) were leukocortical — both types predominantly visible at 7T and therefore under-represented in the training data, which consists primarily of 3T acquisitions [40, 19]. Importantly, the dominant failure mode is *missed detection of subpial lesions*, not hallucination of artifacts: precision at site C 7T was 0.645, meaning that 64.5% of model detections were correct and the false-positive count was modest (mean 7.0 FPL/scan). It is also worth noting that in clinical practice, CL assessment is primarily detection-based: the presence of even a single confirmed CL carries diagnostic relevance under the McDonald criteria [13]. Volumetric accuracy (DSC) is therefore a less clinically relevant metric for CLs than for larger structures such as WMLs, and the model’s detection capability at 7T, while imperfect, may still provide diagnostic value. For sites A-C, leukocortical lesions were frequently missed (false negatives) by the automated segmentation approach, while numerous false positives corresponded to previously annotated WMLs. This difficulty reflects inherent ambiguity in distinguishing leukocortical from juxtacortical lesions, i.e., CL and WML, a well-known challenge in manual segmentation.

Direct comparison of our results with prior CL segmentation studies is non-trivial due to differences in detection threshold (we use a 1-voxel overlap criterion), training/test cohort composition, CL definition, and whether OOD evaluation is performed. With these caveats noted, our in-domain detection F1 of 0.53–0.73 across sites A and B is broadly consistent with results reported by La Rosa et al. [20] and Gordaliza et al. [21] under comparable conditions. Crucially, our OOD F1 of approximately 0.44–0.50 at site D represents a substantially more challenging evaluation than in prior work, which generally tested on small cohorts with similar acquisition protocols and no true distribution shift of the scale considered here. The observed inter-rater variability in our expert feedback study is also consistent with prior reports: CL annotation is known to be more variable than WML annotation [19], reflecting subtler imaging contrast and less standardised delineation criteria. Inter-rater Dice scores for WML segmentation have been reported in the range of 0.60–0.80 [41, 42], while infratentorial and spinal cord MS lesion variability is higher [43, 44, 45]; CL variability likely exceeds all of these,

given the additional challenge of cortical boundary determination.

5.3. Explainability analysis insights

The apparent clustering observed in PCA projection—where testing data from sites A and B showed shifts from training distributions—may reflect the inherent nature of learned representations or true overfitting. PCA captures only linear relationships in the high-dimensional bottleneck features, and it is expected that certain feature channels will capture scanner-specific or site-specific characteristics (acquisition noise patterns, intensity distributions, contrast variations) alongside task-relevant CL representations. This scanner-specific information naturally creates linearly separable clusters in PCA space, even when the overall learned representation remains robust for the segmentation task. Therefore, PCA alone may not be sufficient for assessing model generalization, as it cannot distinguish between harmful overfitting and expected technical variation capture. The contrasting patterns observed in UMAP projections, which preserve non-linear manifold structure, may provide more meaningful insights into the network’s actual generalization capabilities. The site- and modality-specific clustering is reduced when using UMAP projection, compared with PCA, yet it persists. This outcome underscores the critical importance of multi-center datasets and the necessity of data harmonization or domain adaptation techniques to accommodate diverse clinical contexts [46]. A centered location of the OOD data for the first three PCA components should suggest that some robust CL or T1w-modality MS brain representations were learned despite the dominant presence of some sites in the Train set.

5.4. Medical expert feedback on OOD data

The OOD evaluation should be indicative of the worst-case scenario model performance. Medical expert feedback provided a fine-grained analysis of the OOD performance, not reflected in the selected evaluation metrics. Particularly, we observed that up to 50% of FPLs were actual CLs, missed during the manual annotation. This finding has important implications for interpreting the quantitative metrics reported throughout this study: if up to half of the model’s apparent false positives are in fact true lesions overlooked by the original annotators, then the reported precision values systematically underestimate the model’s true performance, and F1-scores should be interpreted as lower bounds. This is consistent with the known difficulty of CL annotation, even for expert neurologists: subpial lesions in particular are inherently

difficult to delineate due to their thin band-like morphology and proximity to the cortical surface, and inter-rater agreement for subpial lesions is lower than for other CL subtypes [19]. Ground truth imperfection is a persistent challenge in CL research, and automated tools that can systematically flag candidate lesions may help address it by providing a structured second opinion. This shows the potential of the proposed model as a second opinion system, reinforcing medical decisions.

Some lesions are challenging for both the model and human annotators. Particularly, both raters reported FNLs to be non-MS in 40.8%-62%. The change of opinion on the ground-truth lesions and reasons for decreased confidence indicate that small lesions and lesions with unclear cortical involvement can represent a challenge for both model and human annotators. The confidence level analysis confirms that FNLs represent a significant challenge compared to TPLs for medical experts, which is not observed for FPLs vs. TPLs.

Previous discussions about the limitations of the quality measures is enhanced by this experiment results, showing that ground-truth delineation is not perfect itself. Medical experts suggest that delineation needs correction for $\approx 30\%$ FPLs and TPLs.

5.5. Limitations and future work

An evident limitation in our study arises from the exclusive reliance on T1-weighted MRI modalities (MPRAGE and MP2RAGE). The persistent inability of models—and human observers—to reliably identify subpial lesions at 3T T1-weighted imaging confirms prior research emphasizing limited subpial lesion visibility at lower magnetic field strengths [19]. Incorporating additional MRI contrasts (e.g., T2-weighted or fluid-suppressed sequences) or using ultra-high field data [20] might theoretically enhance segmentation quality. However, practical hurdles such as non-standardized acquisition protocols and limited availability across clinical sites significantly hinder this integration. This study again highlights that different CLs are visible on different MRI modalities, underscoring the importance of standardized lesion annotation criteria to facilitate consistent CL interpretation across studies and improve the clinical validity of automated segmentation tools.

Future studies should explore human-AI interaction systematically, assessing clinician trust in AI-generated segmentations and possible integration scenarios. Collecting feedback from the annotators about the missed

and false-positive lesions would be essential for understanding the DL model decisions.

From the AI-engineering perspective, there are several future directions aiming to improve in-domain performance and generalizability. We identified that varying sizes and per-subject loads, with overall small volumes of segmented instances, still represent a challenge for the proposed model, even in-domain (Figure 5). Potential remedies include size-dependent labeling for improved small lesion segmentation [47], inclusion of nDSC to the loss for improving fairness across subjects [48, 49], or using pretrained transformer-CNN architectures (e.g., UNETR or SwinUnet) for adaptive context reasoning and semantic consistency properties. To improve the performance under the domain shift, one can adopt existing pretrained models for WML models [50] or the nnU-Net framework [51], implement synthetic lesion generation strategies (e.g., CarveMix [52]), or consider test-time adaptation approaches [53].

In clinical practice, MPRAGE and MP2RAGE are typically acquired as alternatives rather than in combination, so fusing their predictions is of limited practical scope. Nonetheless, where both are available for the same subject, averaging output probability maps could offer modest gains given their partly overlapping contrast properties. A more impactful direction would be to incorporate sequences with genuinely complementary contrast, such as T2*-weighted imaging, which provides substantially improved subpial lesion visibility. The current study does not pursue multi-sequence fusion due to data constraints, but it remains a natural direction for future work.

The aforementioned approaches would not resolve fundamental ambiguities inherent to distinguishing leukocortical from juxtacortical lesions, highlighting persistent challenges due to inherent lesion variability and high prevalence. Although quantifying exact cortical involvement remains problematic, future annotation protocols might benefit from marking uncertain lesions explicitly, enabling the use of soft labels and probabilistic approaches in training automated segmentation models.

6. Conclusions

This study evaluates and compares approaches to automated cortical lesion segmentation in multiple sclerosis using the nnU-Net framework. Our results show that a standard nnU-Net architecture, with task-specific adjustments like the Blob BCE + Dice Loss, effectively segments cortical lesions across different imaging centers and MRI protocols. Cross-site performance

variability and the similarity between leukocortical and juxtacortical lesions remain challenging and necessitate clearer annotation guidelines, possibly involving soft labels.

Automated segmentation tools demonstrated a potential for clinical workflows as supportive diagnostic or research aids. Their role could include providing automatic lesion segmentation as a second opinion to highlight lesions potentially missed by expert raters. Additionally, the lesion delineation procedure can be significantly simplified with the presence of automatic tools, even if manual corrections are required.

Acknowledgment

This work was supported by the Hasler Foundation Responsible AI program (MSxplain) and the Research Commission of the Faculty of Biology and Medicine (CRFBM) of UNIL. 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). AC is supported by EUROSTAR E!113682 HORIZON2020. DSR and CT are supported by the Intramural Research Program of NINDS, NIH. This research was supported in part by the Intramural Research Program of the National Institutes of Health (NIH). The contributions of the NIH authors are considered Works of the United States Government. The findings and conclusions presented in this paper are those of the authors and do not necessarily reflect the views of the NIH or the U.S. Department of Health and Human Services. DSR received research funding from Abata and Sanofi. CVB received funding from the Fonds de Recherche Clinique (FRC) from Cliniques Universitaires Saint-Luc (CUSL). AS has the financial support of the Fédération Wallonie Bruxelles – FRIA du Fonds de la Recherche Scientifique – FNRS. SB is supported by the Funds Claire Faconnier, Ginette Kryksztein & José, and Marie Philippart-Hoffelt, managed by the King Baudouin Foundation. She has received unrelated research funding from Roche, Sanofi, and Brugmann Foundation. PM is supported by the Fondation Charcot Stichting Research Fund 2023, the Fund for Scientific Research (F.R.S, FNRS; grant #40008331), Cliniques universitaires Saint-Luc “Fonds de Recherche Clinique”, and Biogen. He has received consulting honoraria from Sanofi, Biogen and Merck.

Ethics statement

Studies involving human data were approved by the local ethics committees; informed consent was obtained from all participants before study entry.

Appendix A. Bottleneck space analysis

The relationship between projected internal model features and continuous features (DSC, patient-wise model uncertainty, lesion volume errors) is quantified in Figure A.11. Mild correlation is observed between bottleneck feature components and continuous features, with the highest overall correlation values of DSC.

The patient structural uncertainty (PSU) quantifies the disagreement between individual ensemble predictions and the average ensemble prediction, reflecting both epistemic and aleatoric uncertainty (potential model errors or ambiguous data) [54]. Figure A.12 illustrates model features with uncertainty information. Low uncertainty examples are found across all PCA clusters of the bottleneck feature space, even in clusters distant from the Site A training set (e.g., Site B MPRAGE cluster). Conversely, Site C training and test examples consistently exhibit high uncertainty. This widespread higher uncertainty in training data across different sites may indicate significant data noise, hindering the learning of consistent representations.

Appendix B. Medical expert feedback

An additional plot regarding the reasons for the decreased confidence is given in Figure B.13.


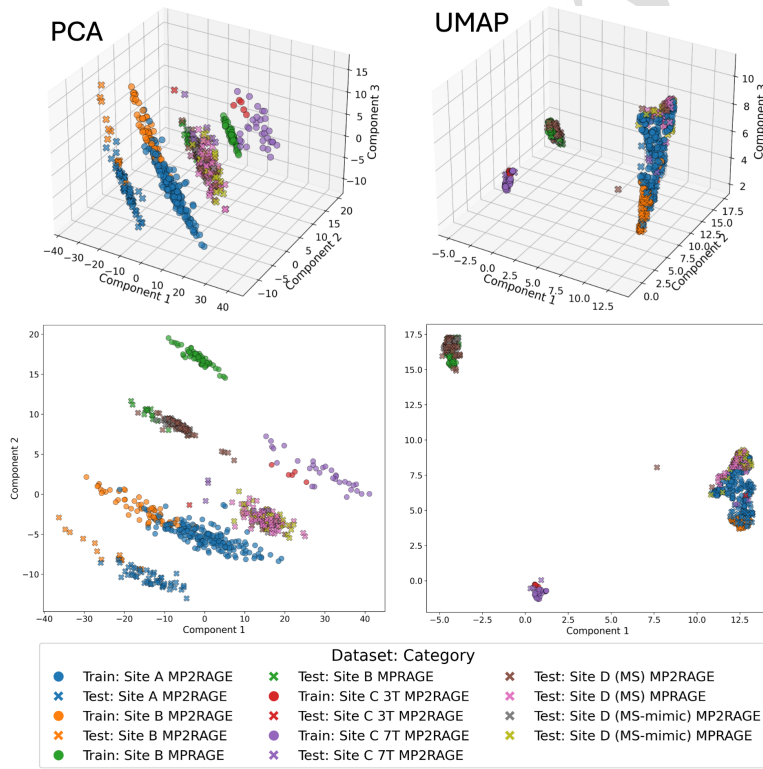


Figure 7: Visualization of the manual and automatic segmentation per site for scans with high segmentation quality ($DSC > Q3$) and median lesion count per site. Images from left to right show MRI axial slices with i) no segmentation, ii) ground truth mask overlay (blue), iii) predicted mask overlay (pink).

Figure 8: Bottleneck features visualization using PCA (left) and UMAP (right) dimensionality reduction.



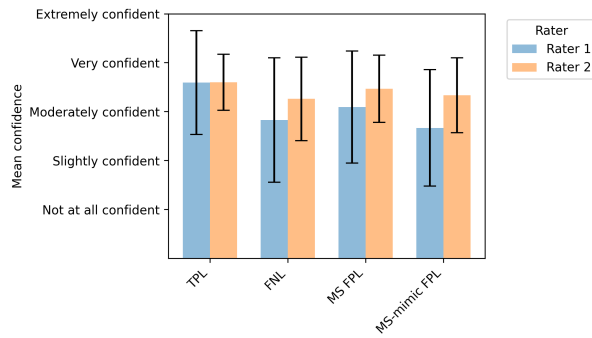


Figure 9: Mean and standard deviation of the confidence in ROI classification for each ROI type, provided by each rater.




Figure 10: Bar plot with ROI classification results for different lesion types provided by 2 raters (left and right bars respectively in each category).




Figure A.11: Pearson correlation between projected bottleneck features and continuous features.

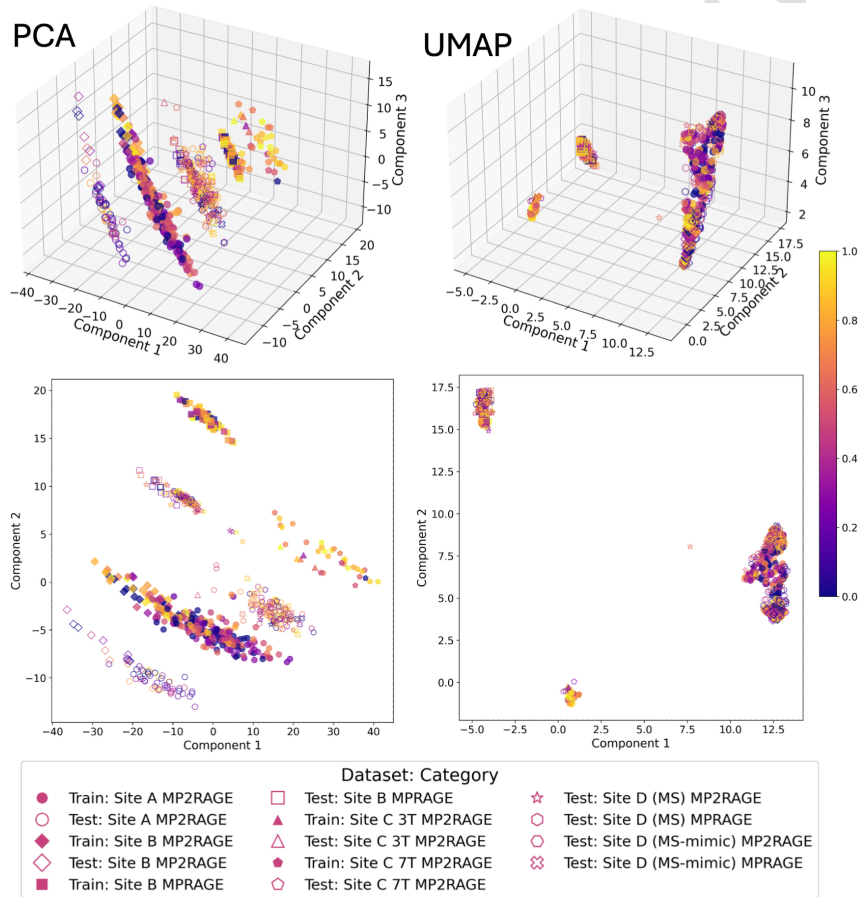


Figure A.12: Distribution of patient-wise uncertainty ranking (computed with PSU measure) with respect to the internal model representations and sites. Latent features were projected with PCA (left) and UMAP (right). Original PSU values underwent the quantile transform to accommodate a skewed uncertainty distribution.




Figure B.13: Bar plot with reasons for decreased confidence results for different lesion types provided by 2 raters (left and right).

References

- [1] D. S. Reich, C. F. Lucchinetti, P. A. Calabresi, Multiple Sclerosis, *New England Journal of Medicine* 378 (2018) 169–180. URL: <http://www.nejm.org/doi/10.1056/NEJMra1401483>. doi:10.1056/NEJMra1401483.
- [2] E. Portaccio, M. Magyari, E. K. Havrdova, A. Ruet, B. Brochet, A. Scalfari, M. Di Filippo, C. Tur, X. Montalban, M. P. Amato, Multiple sclerosis: emerging epidemiological trends and redefining the clinical course, *The Lancet Regional Health - Europe* 44 (2024) 100977. URL: <https://www.sciencedirect.com/science/article/pii/S2666776224001443>. doi:10.1016/j.lanepe.2024.100977.
- [3] C. C. Hemond, R. Bakshi, Magnetic resonance imaging in multiple sclerosis, *Cold Spring Harbor Perspectives in Medicine* 8 (2018). doi:10.1101/cshperspect.a028969, publisher: Cold Spring Harbor Laboratory Press.
- [4] M. Filippi, W. Brück, D. Chard, F. Fazekas, J. J. G. Geurts, C. Enzinger, S. Hametner, T. Kuhlmann, P. Preziosa, A. Rovira, K. Schmierer, C. Stadelmann, M. A. Rocca, Attendees of the Correlation between Pathological and MRI findings in MS workshop, Association between pathological and MRI findings in multiple sclerosis, *The Lancet. Neurology* 18 (2019) 198–210. doi:10.1016/S1474-4422(18)30451-4.
- [5] M. P. Wattjes, O. Ciccarelli, D. S. Reich, B. Banwell, N. de Stefano, C. Enzinger, F. Fazekas, M. Filippi, J. Frederiksen, C. Gasperini, Y. Hachohen, L. Kappos, D. K. B. Li, K. Mankad, X. Montalban, S. D. Newsome, J. Oh, J. Palace, M. A. Rocca, J. Sastre-Garriga, M. Tintoré, A. Traboulsee, H. Vrenken, T. Yousry, F. Barkhof, A. Rovira, M. P. Wattjes, O. Ciccarelli, N. de Stefano, C. Enzinger, F. Fazekas, M. Filippi, J. Frederiksen, C. Gasperini, Y. Hachohen, L. Kappos, K. Mankad, X. Montalban, J. Palace, M. A. Rocca, J. Sastre-Garriga, M. Tintore, H. Vrenken, T. Yousry, F. Barkhof, A. Rovira, D. K. B. Li, A. Traboulsee, S. D. Newsome, B. Banwell, J. Oh, D. S. Reich, D. S. Reich, J. Oh, 2021 MAGNIMS–CMSC–NAIMS consensus recommendations on the use of MRI in patients with multiple sclerosis, *The Lancet Neurology* 20 (2021) 653–670. URL: [https://doi.org/10.1016/S1474-4422\(21\)00095-8](https://doi.org/10.1016/S1474-4422(21)00095-8). doi:10.1016/S1474-4422(21)00095-8.

- [6] H. Kolb, O. Al-Louzi, E. S. Beck, P. Sati, M. Absinta, D. S. Reich, From pathology to MRI and back: clinically relevant biomarkers of multiple sclerosis lesions, *NeuroImage: Clinical* (2022) 103194. URL: <https://linkinghub.elsevier.com/retrieve/pii/S2213158222002595>. doi:10.1016/j.nicl.2022.103194.
- [7] E. S. Beck, J. Maranzano, N. J. Luciano, P. Parvathaneni, S. Filipini, M. Morrison, D. J. Suto, T. Wu, P. van Gelderen, J. A. de Zwart, S. Antel, D. Fetco, J. Ohayon, F. Andrada, Y. Mina, C. Thomas, S. Jacobson, J. Duyn, I. Cortese, S. Narayanan, G. Nair, P. Sati, D. S. Reich, Cortical lesion hotspots and association of subpial lesions with disability in multiple sclerosis, *Multiple Sclerosis Journal* 28 (2022) 1351–1363. URL: <https://doi.org/10.1177/13524585211069167>. doi:10.1177/13524585211069167, publisher: SAGE Publications Ltd STM.
- [8] A. Cagol, R. Cortese, M. Barakovic, S. Schaedelin, E. Ruberte, M. Absinta, F. Barkhof, M. Calabrese, M. Castellaro, O. Ciccarelli, S. Cocozza, N. De Stefano, C. Enzinger, M. Filippi, M. Jurynczyk, P. Maggi, N. Mahmoudi, S. Messina, X. Montalban, J. Palace, G. Pontillo, A.-K. Probstel, M. A. Rocca, S. Ropele, A. Rovira, M. M. Schoonheim, P. Sowa, E. Strijbis, M. P. Wattjes, M. P. Sormani, L. Kappos, C. Granziera, MAGNIMS Study Group, Diagnostic Performance of Cortical Lesions and the Central Vein Sign in Multiple Sclerosis, *JAMA Neurology* 81 (2024) 143–153. URL: <https://doi.org/10.1001/jamaneurol.2023.4737>. doi:10.1001/jamaneurol.2023.4737.
- [9] M. Absinta, P. Sati, F. Masuzzo, G. Nair, V. Sethi, H. Kolb, J. Ohayon, T. Wu, I. C. M. Cortese, D. S. Reich, Association of Chronic Active Multiple Sclerosis Lesions With Disability In Vivo, *JAMA neurology* 76 (2019) 1474–1483. doi:10.1001/jamaneurol.2019.2399.
- [10] P. Maggi, P. Sati, G. Nair, I. C. M. Cortese, S. Jacobson, B. R. Smith, A. Nath, J. Ohayon, V. v. Pesch, G. Perrotta, C. Pot, M. Théaudin, V. Martinelli, R. Scotti, T. Wu, R. D. Pasquier, P. A. Calabresi, M. Filippi, D. S. Reich, M. Absinta, Paramagnetic Rim Lesions are Specific to Multiple Sclerosis: An International Multicenter 3T MRI Study, *Annals of Neurology* 88 (2020) 1034–1042. doi:10.1002/ana.25877, publisher: John Wiley and Sons Inc.

- [11] P. Sati, J. Oh, R. T. Constable, N. Evangelou, C. R. G. Guttmann, R. G. Henry, E. C. Klawiter, C. Mainero, L. Massacesi, H. McFarland, F. Nelson, D. Ontaneda, A. Rauscher, W. D. Rooney, A. P. R. Samaraweera, R. T. Shinohara, R. A. Sobel, A. J. Solomon, C. A. Treaba, J. Wuerfel, R. Zivadinov, N. L. Sicotte, D. Pelletier, D. S. Reich, The central vein sign and its clinical evaluation for the diagnosis of multiple sclerosis: a consensus statement from the North American Imaging in Multiple Sclerosis Cooperative, *Nature Reviews Neurology* 12 (2016) 714–722. URL: <https://www.nature.com/articles/nrneurol.2016.166>. doi:10.1038/nrneurol.2016.166, publisher: Nature Publishing Group.
- [12] E. S. Beck, W. A. Mullins, J. dos Santos Silva, S. Filippini, P. Parvathani, J. Maranzano, M. Morrison, D. J. Suto, C. Donnay, H. Dieckhaus, N. J. Luciano, K. Sharma, M. I. Gaitán, J. Liu, J. A. de Zwart, P. van Gelderen, I. Cortese, S. Narayanan, J. H. Duyn, G. Nair, P. Sati, D. S. Reich, Contribution of new and chronic cortical lesions to disability accrual in multiple sclerosis, *Brain Communications* 6 (2024) fcae158. URL: <https://doi.org/10.1093/braincomms/fcae158>. doi:10.1093/braincomms/fcae158.
- [13] A. J. Thompson, B. Banwell, F. Barkhof, W. M. Carroll, T. Coetzee, G. Comi, J. Correale, F. Fazekas, M. Filippi, M. S. Freedman, K. Fujihara, S. Galetta, H. P. Hartung, L. Kappos, F. Lublin, R. A. Marrie, A. Miller, D. H. Miller, X. Montalbán, E. M. Mowry, P. S. Sørensen, M. Tintoré, A. Traboulsee, M. Trojano, B. M. J. Uitdehaag, S. Vukusic, E. Waubant, B. G. Weinshenker, S. C. Reingold, J. A. Cohen, Diagnosis of multiple sclerosis: 2017 revisions of the McDonald criteria, *Lancet Neurology* 17 (2018) 162–173. URL: [https://doi.org/10.1016/s1474-4422\(17\)30470-2](https://doi.org/10.1016/s1474-4422(17)30470-2). doi:10.1016/s1474-4422(17)30470-2, publisher: Elsevier BV.
- [14] F. L. Rosa, A. Abdulkadir, M. J. Fartaria, R. Rahmanzadeh, P. J. Lu, R. Galbusera, M. Barakovic, J. P. Thiran, C. Granziera, M. B. Cuadra, Multiple sclerosis cortical and WM lesion segmentation at 3T MRI: a deep learning method based on FLAIR and MP2RAGE, *NeuroImage: Clinical* 27 (2020). doi:10.1016/j.nicl.2020.102335, publisher: Elsevier Inc.
- [15] O. Ronneberger, P. Fischer, T. Brox, U-NET: Convolutional Networks

- for Biomedical Image Segmentation, 2015. URL: https://doi.org/10.1007/978-3-319-24574-4_28. doi:10.1007/978-3-319-24574-4_28, publication Title: Lecture notes in computer science.
- [16] O. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, O. Ronneberger, 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation, in: S. Ourselin, L. Joskowicz, M. R. Sabuncu, G. Unal, W. Wells (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016, Springer International Publishing, Cham, 2016, pp. 424–432. doi:10.1007/978-3-319-46723-8_49.
- [17] T. Kober, C. Granziera, D. Ribes, P. Browaeys, M. Schlupe, R. Meuli, R. Frackowiak, R. Gruetter, G. Krueger, MP2RAGE multiple sclerosis magnetic resonance imaging at 3 T, *Investigative Radiology* 47 (2012) 346–352. doi:10.1097/RLI.0b013e31824600e9.
- [18] M. J. Fartaria, T. Kober, C. Granziera, M. B. Cuadra, Longitudinal analysis of white matter and cortical lesions in multiple sclerosis, *NeuroImage: Clinical* 23 (2019) 101938. URL: <https://www.sciencedirect.com/science/article/pii/S2213158219302888>. doi:<https://doi.org/10.1016/j.nicl.2019.101938>.
- [19] E. Beck, P. Sati, V. Sethi, T. Kober, B. Dewey, P. Bhargava, G. Nair, I. Cortese, D. Reich, Improved Visualization of Cortical Lesions in Multiple Sclerosis Using 7T MP2RAGE, *AJNR: American Journal of Neuroradiology* 39 (2018) 459–466. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6082739/>. doi:10.3174/ajnr.A5534.
- [20] F. La Rosa, E. S. Beck, J. Maranzano, R.-A. Todea, P. van Gelderen, J. A. de Zwart, N. J. Luciano, J. H. Duyn, J.-P. Thiran, C. Granziera, D. S. Reich, P. Sati, M. Bach Cuadra, Multiple sclerosis cortical lesion detection with deep learning at ultra-high-field MRI, *NMR in Biomedicine* 35 (2022) e4730. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nbm.4730>. doi:10.1002/nbm.4730, _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nbm.4730>.
- [21] P. M. Gordaliza, J. Muller, A. Cagol, N. Molchanova, F. L. Rosa, C. Tsagkas, C. Granziera, M. B. Cuadra, Fluid and White Matter Suppression Contrasts MRI Improves Deep Learning Detection of Multiple Sclerosis Cortical Lesions, 2025. URL: <https://www.medrxiv>.

- [org/content/10.1101/2025.01.19.25320792v1](https://doi.org/10.1101/2025.01.19.25320792v1). doi:10.1101/2025.01.19.25320792, pages: 2025.01.19.25320792.
- [22] F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, K. H. Maier-Hein, nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation, *Nature Methods* 18 (2021) 203–211. URL: <https://www.nature.com/articles/s41592-020-01008-z>. doi:10.1038/s41592-020-01008-z, publisher: Nature Publishing Group.
- [23] M. G. Dwyer, N. Bergsland, A. Bartnik, D. Jakimovski, S. Noteboom, M. M. Schoonheim, M. D. Steenwijk, J. Pei, D. Clayton, R. Zivadinov, Quantifying cortical lesions in large legacy multiple sclerosis clinical trial MRI datasets using multi-contrast post-processing and deep learning, 2025. URL: <https://www.medrxiv.org/content/10.1101/2025.04.08.25325461v1>. doi:10.1101/2025.04.08.25325461, pages: 2025.04.08.25325461.
- [24] F. Isensee, M. Schell, I. Pflueger, G. Brugnara, D. Bonekamp, U. Neuberger, A. Wick, H.-P. Schlemmer, S. Heiland, W. Wick, M. Bendszus, K. Maier-Hein, P. Kickingeder, Automated brain extraction of multi-sequence MRI using artificial neural networks, *Human Brain Mapping* 40 (2019). doi:10.1002/hbm.24750.
- [25] B. Billot, D. N. Greve, O. Puonti, A. Thielscher, K. Van Leemput, B. Fischl, A. V. Dalca, J. E. Iglesias, SynthSeg: Segmentation of brain MRI scans of any contrast and resolution without retraining, *Medical Image Analysis* 86 (2023) 102789. URL: <https://www.sciencedirect.com/science/article/pii/S1361841523000506>. doi:10.1016/j.media.2023.102789.
- [26] M. A. J. Madsen, V. Wiggermann, S. Bramow, J. R. Christensen, F. Sellebjerg, H. R. Siebner, Imaging cortical multiple sclerosis lesions with ultra-high field MRI, *NeuroImage: Clinical* 32 (2021) 102847. URL: <https://www.sciencedirect.com/science/article/pii/S2213158221002916>. doi:10.1016/j.nicl.2021.102847.
- [27] F. Isensee, T. Wald, C. Ulrich, M. Baumgartner, S. Roy, K. Maier-Hein, P. F. Jäger, nnU-Net Revisited: A Call for Rigorous Validation in 3D Medical Image Segmentation, in: M. G. Linguraru, Q. Dou, A. Feragen,

- S. Giannarou, B. Glocker, K. Lekadir, J. A. Schnabel (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, Springer Nature Switzerland, Cham, 2024, pp. 488–498. doi:[10.1007/978-3-031-72114-4_47](https://doi.org/10.1007/978-3-031-72114-4_47).
- [28] D. P. Kingma, J. Ba, Adam: A Method for Stochastic Optimization, 2014. URL: <http://arxiv.org/abs/1412.6980>, arXiv: 1412.6980.
- [29] J. Ma, F. Li, B. Wang, U-Mamba: Enhancing Long-range Dependency for Biomedical Image Segmentation, 2024. URL: <http://arxiv.org/abs/2401.04722>. doi:[10.48550/arXiv.2401.04722](https://doi.org/10.48550/arXiv.2401.04722), arXiv:2401.04722 [cs, eess].
- [30] F. Kofler, S. Shit, I. Ezhov, L. Fidon, I. Horvath, R. Al-Maskari, H. Li, H. Bhatia, T. Loehr, M. Piraud, A. Erturk, J. Kirschke, J. C. Peeken, T. Vercauteren, C. Zimmer, B. Wiestler, B. Menze, blob loss: instance imbalance aware loss functions for semantic segmentation, 2023. URL: <http://arxiv.org/abs/2205.08209>. doi:[10.48550/arXiv.2205.08209](https://doi.org/10.48550/arXiv.2205.08209), arXiv:2205.08209 [cs, eess].
- [31] F. La Rosa, M. Wynen, O. Al-Louzi, E. S. Beck, T. Huelnhagen, P. Maggi, J.-P. Thiran, T. Kober, R. T. Shinohara, P. Sati, D. S. Reich, C. Granziera, M. Absinta, M. Bach Cuadra, Cortical lesions, central vein sign, and paramagnetic rim lesions in multiple sclerosis: Emerging machine learning techniques and future avenues, *NeuroImage: Clinical* 36 (2022) 103205. URL: <https://www.sciencedirect.com/science/article/pii/S2213158222002704>. doi:[10.1016/j.nicl.2022.103205](https://doi.org/10.1016/j.nicl.2022.103205).
- [32] V. Raina, N. Molchanova, M. Graziani, A. Malinin, H. Muller, M. B. Cuadra, M. Gales, Novel Structural-Scale Uncertainty Measures and Error Retention Curves: Application to Multiple Sclerosis, in: *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, 2023, pp. 1–5. doi:[10.1109/ISBI53787.2023.10230755](https://doi.org/10.1109/ISBI53787.2023.10230755).
- [33] C. M. Bishop, *Pattern recognition and machine learning*, Information science and statistics, Springer, New York, 2006.
- [34] L. McInnes, J. Healy, J. Melville, UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction, 2020. URL: <http://arxiv.org/abs/1802.03426>.

- [//arxiv.org/abs/1802.03426](https://arxiv.org/abs/1802.03426). doi:10.48550/arXiv.1802.03426, arXiv:1802.03426 [stat].
- [35] P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, G. Gerig, User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability, *NeuroImage* 31 (2006) 1116–1128. doi:10.1016/j.neuroimage.2006.01.015.
- [36] J. Jiang, M. Wang, H. Tian, L. Cheng, Y. Liu, LV-UNet: A Lightweight and Vanilla Model for Medical Image Segmentation, 2024. URL: <http://arxiv.org/abs/2408.16886>. doi:10.48550/arXiv.2408.16886, arXiv:2408.16886 [eess].
- [37] S. Javed, T. M. Khan, A. Qayyum, H. Alinejad-Rokny, A. Sowmya, I. Razzak, Advancing Medical Image Segmentation with Mini-Net: A Lightweight Solution Tailored for Efficient Segmentation of Medical Images, 2024. URL: <http://arxiv.org/abs/2405.17520>. doi:10.48550/arXiv.2405.17520, arXiv:2405.17520 [eess].
- [38] E. J. Topol, High-performance medicine: the convergence of human and artificial intelligence, *Nature Medicine* 25 (2019) 44–56. URL: <https://www.nature.com/articles/s41591-018-0300-7>. doi:10.1038/s41591-018-0300-7, publisher: Nature Publishing Group.
- [39] E. J. Topol, As artificial intelligence goes multimodal, medical applications multiply, *Science* 381 (2023) eadk6139. URL: <https://www.science.org/doi/10.1126/science.adk6139>. doi:10.1126/science.adk6139, publisher: American Association for the Advancement of Science.
- [40] J. Maranzano, M. Dadar, D. A. Rudko, D. D. Nigris, C. Elliott, J. S. Gati, S. A. Morrow, R. S. Menon, D. L. Collins, D. L. Arnold, S. Narayanan, Comparison of Multiple Sclerosis Cortical Lesion Types Detected by Multicontrast 3T and 7T MRI, *American Journal of Neuroradiology* 40 (2019) 1162–1169. URL: <https://www.ajnr.org/content/40/7/1162>. doi:10.3174/ajnr.A6099, publisher: American Journal of Neuroradiology Section: Adult Brain.

- [41] C. Egger, R. Opfer, C. Wang, T. Kepp, M. P. Sormani, L. Spies, M. Barnett, S. Schippling, MRI FLAIR lesion segmentation in multiple sclerosis: Does automated segmentation hold up with manual annotation?, *NeuroImage: Clinical* 13 (2017) 264–270. URL: <https://linkinghub.elsevier.com/retrieve/pii/S2213158216302285>. doi:10.1016/j.nicl.2016.11.020.
- [42] A. Carass, S. Roy, A. Gherman, J. C. Reinhold, A. Jesson, T. Arbel, O. Maier, H. Handels, M. Ghafoorian, B. Platel, A. Birenbaum, H. Greenspan, D. L. Pham, C. M. Crainiceanu, P. A. Calabresi, J. L. Prince, W. R. G. Roncal, R. T. Shinohara, I. Oguz, Evaluating White Matter Lesion Segmentations with Refined Sørensen-Dice Analysis, *Scientific Reports* 10 (2020) 8242. URL: <https://www.nature.com/articles/s41598-020-64803-w>. doi:10.1038/s41598-020-64803-w, bibtext:carass_evaluating_2020.
- [43] J. Krüger, A.-C. Ostwaldt, L. Spies, B. Geisler, A. Schlaefer, H. H. Kitzler, S. Schippling, R. Opfer, Infratentorial lesions in multiple sclerosis patients: intra- and inter-rater variability in comparison to a fully automated segmentation using 3D convolutional neural networks, *European Radiology* 32 (2022) 2798–2809. URL: <https://doi.org/10.1007/s00330-021-08329-3>. doi:10.1007/s00330-021-08329-3.
- [44] C. Gros, B. De Leener, A. Badji, J. Maranzano, D. Eden, S. M. Dupont, J. Talbott, R. Zhuoqiong, Y. Liu, T. Granberg, R. Ouellette, Y. Tachibana, M. Hori, K. Kamiya, L. Chougar, L. Stawiarz, J. Hillert, E. Bannier, A. Kerbrat, G. Edan, P. Labauge, V. Callot, J. Pelletier, B. Audoin, H. Rasoanandrianina, J.-C. Brisset, P. Valsasina, M. A. Rocca, M. Filippi, R. Bakshi, S. Tauhid, F. Prados, M. Yianakakos, H. Kearney, O. Ciccarelli, S. Smith, C. A. Treaba, C. Mainero, J. Lefeuvre, D. S. Reich, G. Nair, V. Auclair, D. G. McLaren, A. R. Martin, M. G. Fehlings, S. Vahdat, A. Khatibi, J. Doyon, T. Shepherd, E. Charlson, S. Narayanan, J. Cohen-Adad, Automatic segmentation of the spinal cord and intramedullary multiple sclerosis lesions with convolutional neural networks, *NeuroImage* 184 (2019) 901–915. URL: <https://www.sciencedirect.com/science/article/pii/S1053811918319578>. doi:10.1016/j.neuroimage.2018.09.081, bibtext:gros_automatic_2019.

- [45] R. Walsh, C. Meurée, A. Kerbrat, A. Masson, B. R. Hussein, M. Gaubert, F. Galassi, B. Combés, Expert Variability and Deep Learning Performance in Spinal Cord Lesion Segmentation for Multiple Sclerosis Patients, in: 2023 IEEE 36th International Symposium on Computer-Based Medical Systems (CBMS), IEEE, L'Aquila, Italy, 2023, pp. 1–8. URL: <https://inria.hal.science/hal-04090598>. doi:10.1109/CBMS58004.2023.00263, bibtex:walsh_expert_2023.
- [46] B. Glocker, R. Robinson, D. C. Castro, Q. Dou, E. Konukoglu, Machine Learning with Multi-Site Imaging Data: An Empirical Study on the Impact of Scanner Effects, 2019. URL: <http://arxiv.org/abs/1910.04597>. doi:10.48550/arXiv.1910.04597, arXiv:1910.04597 [eess].
- [47] L. Shang, Z. Lou, A. L. Alexander, V. Prabhakaran, W. A. Sethares, V. A. Nair, N. Adluru, Segmenting Small Stroke Lesions with Novel Labeling Strategies, 2024. URL: <http://arxiv.org/abs/2408.02929>. doi:10.48550/arXiv.2408.02929, arXiv:2408.02929 [cs].
- [48] V. Raina, N. Molchanova, M. Graziani, A. Malinin, H. Muller, M. B. Cuadra, M. Gales, Tackling Bias in the Dice Similarity Coefficient: Introducing NDSC for White Matter Lesion Segmentation, in: 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI), 2023, pp. 1–5. URL: <https://ieeexplore.ieee.org/abstract/document/10230755>. doi:10.1109/ISBI53787.2023.10230755, iSSN: 1945-8452.
- [49] F. Spagnolo, N. Molchanova, M. O. Pineda, L. Melie-Garcia, M. B. Cuadra, C. Granziera, V. Andrearczyk, A. Depeursinge, Exploiting XAI maps to improve MS lesion segmentation and detection in MRI, 2024. URL: <https://www.biorxiv.org/content/10.1101/2024.08.29.610090v1>. doi:10.1101/2024.08.29.610090, pages: 2024.08.29.610090 Section: New Results.
- [50] F. La Rosa, A. Abdulkadir, J.-P. Thiran, C. Granziera, M. Bach Cuadra, Software: Multiple sclerosis cortical and WM lesion segmentation at 3T MRI: A deep learning method based on FLAIR and MP2RAGE, 2020. URL: <https://zenodo.org/records/3942042>. doi:10.5281/zenodo.3942042.
- [51] T. Wald, C. Ulrich, S. Lukyanenko, A. Goncharov, A. Paderno, M. Miller, L. Maerkisch, P. F. Jäger, K. Maier-Hein, Revisiting MAE

- pre-training for 3D medical image segmentation, 2025. URL: <http://arxiv.org/abs/2410.23132>. doi:10.48550/arXiv.2410.23132, arXiv:2410.23132 [cs].
- [52] X. Zhang, C. Liu, N. Ou, X. Zeng, Z. Zhuo, Y. Duan, X. Xiong, Y. Yu, Z. Liu, Y. Liu, C. Ye, CarveMix: A simple data augmentation method for brain lesion segmentation, *NeuroImage* 271 (2023) 120041. URL: <https://www.sciencedirect.com/science/article/pii/S1053811923001878>. doi:10.1016/j.neuroimage.2023.120041.
- [53] J. M. J. Valanarasu, P. Guo, V. Vs, V. M. Patel, On-the-Fly Test-time Adaptation for Medical Image Segmentation, in: *Medical Imaging with Deep Learning*, PMLR, 2024, pp. 586–598. URL: <https://proceedings.mlr.press/v227/valanarasu24a.html>, iSSN: 2640-3498.
- [54] N. Molchanova, V. Raina, A. Malinin, F. L. Rosa, A. Depursinge, M. Gales, C. Granziera, H. Muller, M. Graziani, M. B. Cuadra, Structural-based uncertainty in deep learning across anatomical scales: Analysis in white matter lesion segmentation, *Computers in Biology and Medicine* 184 (2025) 109336. URL: <https://www.sciencedirect.com/science/article/pii/S0010482524014215>. doi:10.1016/j.combiomed.2024.109336.