



HAL
open science

A study on loss functions and decision thresholds for the segmentation of multiple sclerosis lesions on spinal cord MRI

Burhan Rashid Hussein, Cédric Meurée, Malo Gaubert, Arthur Masson, Anne Kerbrat, Benoît Combès, Francesca Galassi

► To cite this version:

Burhan Rashid Hussein, Cédric Meurée, Malo Gaubert, Arthur Masson, Anne Kerbrat, et al.. A study on loss functions and decision thresholds for the segmentation of multiple sclerosis lesions on spinal cord MRI. 2022. hal-03865212v1

HAL Id: hal-03865212

<https://hal.science/hal-03865212v1>

Preprint submitted on 22 Nov 2022 (v1), last revised 10 May 2023 (v3)

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

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

A STUDY ON LOSS FUNCTIONS AND DECISION THRESHOLDS FOR THE SEGMENTATION OF MULTIPLE SCLEROSIS LESIONS ON SPINAL CORD MRI

*Burhan Rashid Hussein*¹, *Cédric Meurée*¹, *Malo Gaubert*^{1,3}, *Arthur Masson*¹, *Anne Kerbrat*^{1,2},
Benoît Combès^{1*}, *Francesca Galassi*^{1*}

¹ Univ Rennes, Inria, CNRS, Inserm IRISA UMR 6074, Empenn ERL U 1228, Rennes, France

² Rennes University Hospital (CHU), Department of Neurology, Rennes, France

³ Rennes University Hospital (CHU), Department of Neuroradiology, Rennes, France

* Authors contributed equally

ABSTRACT

Multiple sclerosis (MS) patients often present hyper-intense T2-w lesions in the spinal cord. The severe imbalance between background and lesion classes poses a major challenge to Deep Learning segmentation approaches, requiring for ad hoc strategies. Careful selection of the loss function and adjustment of the conventional 0.5-thresholding may help mitigating this issue. Our results show the performance advantages of loss functions based on the Tversky Index and the benefits of threshold tuning over more standard settings and the state-of-the-art model for MS lesion segmentation on spinal cord MRI.

Index Terms— Multiple sclerosis, segmentation, spinal cord, loss function, decision threshold, medical imaging

1. INTRODUCTION

Multiple Sclerosis (MS) is a chronic inflammatory-demyelinating disease of the central nervous system [1]. Magnetic Resonance Imaging (MRI) is fundamental to characterizing and quantifying MS lesions. The number and volume of lesions are used for MS diagnosis, to track its progression, and to evaluate treatments [2]. Accurate identification of MS lesions in MR images is extremely difficult due to variability in lesion location, size, and shape, in addition to anatomical variability between subjects. Since manual segmentation requires expert knowledge, it is time-consuming and prone to intra- and inter-expert variability, methods to automatically segment lesions are required.

Deep Learning approaches have shown remarkable performances in medical imaging segmentation tasks. While automatic MS lesion segmentation on brain MRI is a well-studied and addressed problem [3], automatic MS lesion segmentation on spinal cord MRI has rarely been considered [2]. In this work we focus on the latter. It is important to emphasize that MS lesion segmentation in spinal cord MRI remains a more complex task than in brain MRI due to lesion contrast, image quality, a limited amount of annotated data, and image

variability arising from the diversity of MR scanners and acquisition protocols.

A fundamental challenge in MS lesions segmentation in spinal cord MRI is handling the high-class imbalance. Loss functions used in the training of deep learning models differ in their robustness to class imbalance. Careful selection of the loss function is thus crucial. To inform loss function choice, we perform a large-scale loss function comparison.

Loss functions based on the Cross-Entropy, Dice Similarity Coefficient, and Tversky Index have been proposed in the literature to mitigate the effect of imbalanced data. Ma et al. [4] conducted one of the most exhaustive evaluation of segmentation loss functions in medical imaging, comparing 20 losses over four organ segmentation tasks. None of the losses could consistently achieve the best performance on the four tasks, indicating the importance of conducting a loss function comparison to identify the optimal loss for the dataset. On the other hand, Jadon et al. [5] reported significantly improved performance metrics for the segmentation of brain tissues using Tversky and Focal Tversky losses, compared to Dice score variants. In the work from Gros et al. [2] on spinal cord MS lesions segmentation, a standard Dice Loss function was used, without reporting the investigation of alternative losses. Recent studies suggest that the decision threshold on the output probability map (p-map) may play an important role when dealing with highly imbalanced data [6]. In binary segmentation, the decision threshold is conventionally set to 0.5. Optimizing this value may improve the detection of the minority class, i.e. lesion class in our case. Compared to other approaches, adjustment of the decision threshold has the advantage of not altering the input data nor requiring re-training of the model, as it can be applied as a postprocessing step.

This work investigates state-of-the-art segmentation losses and the relevance of adjusting the decision threshold in the complex, highly imbalanced and underexplored task of MS lesion segmentation on spinal cord MRI.

2. MATERIAL AND METHODS

The methods evaluated in the context of this work are integrated into a pipeline developed in Python 3.8 and based on TensorFlow 2.9 as well as the Spinal cord toolbox (SCT) [2].

2.1. Loss functions

Model convergence at training and performance scores depend on the loss function. To handle class imbalance in image segmentation, the Dice Loss was proposed as an alternative to the widely used Binary Cross Entropy (BCE) [7], and in combination with the BCE under the name of Combo Loss [8]. The Log-Cosh Dice Loss is a recently proposed variant that tackles the non-convex nature of the Dice Loss to facilitate model convergence [5]. The Generalized Dice Loss is another variant introduced to handle highly imbalanced segmentation tasks, relying on weighting factors applied to the different classes based on their relative representation in the training data [9]. The Tversky Loss [10], an extension of the Dice Loss, incorporates weights to control the trade-off between false positive (FP) and false negative (FN) samples. The Focal Tversky Loss extends the Tversky Loss by introducing a γ parameter to make the loss more sensitive to difficult samples and well-adapted to highly imbalanced problems [11]. In the context of this work, the default parameter values proposed in the original associated publications have been selected (i.e., for the Tversky Loss, $\alpha = 0.7$, and for the Focal Tversky Loss, $\gamma = 0.75$).

2.2. MRI Data Acquisition

The dataset included 161 T2-w cervical and thoracic MRI scans from 108 subjects (35.84 ± 10.42 years old, 86 women). All patients were diagnosed with a form of MS. Acquisitions were performed in 13 different sites (four brands). Image resolution is $(0.58 \pm 0.12, 0.58 \pm 0.12, 2.81 \pm 0.20)$ mm³. Trained neurologists manually segmented lesions. Training, validation and test sets were generated by assigning a subject to a single set while balancing for lesion loads, as shown in figure 1. One test set (thr-optimization set) was used for optimising the decision threshold, the other test set (unseen set) was used at the test time only. All scans without expert-detected lesions were assigned to the unseen set.

2.3. Preprocessing

The preprocessing steps are similar to the ones proposed in [2]. Input images were reoriented and resampled to 0.5 mm isotropic images through linear interpolations. Images were cropped by selecting 48x48 2D patches in each axial plane around the SC centerline voxels [12]. 48x48x48 3D patches were then selected along the inferior-to-superior axis. An overlap of 75% between consecutive patches was introduced as a data augmentation technique. Intensity normalization




Fig. 1. Boxplots for lesion number and volume in the training (62 MRI scans), validation (13), thr-optimization (14) and unseen set (26, 46 with no MS lesions are not included here). Points are jittered to improve visualization.

was applied on stacked patches obtained from a given volumetric image to homogenize the intensity distributions [13]. Finally, each patch was normalized to zero mean and unit standard deviation.

2.4. Model Training

The models trained to segment spinal cord MS lesions were based on a 3D-Unet architecture [14]. 16, 32, and 64 filters were selected, associated with 3x3x3 convolutions, 2x2x2 max-pooling layers, and batch normalization applied after each convolution layer. We used a batch size of 8, an Adam optimizer, and an initial learning rate of 10^{-4} , except when using the Combo Loss, which suggested using an Adadelta optimization with a learning rate of 1. The learning rate was decreased by a factor of 2 after every 15 epochs if there was no decrease in validation loss. The training process was stopped after 300 epochs, the model providing the best validation loss being selected at the end of each experiment. Data augmentation operations were applied, including randomly mirroring and rotating patches by 90°.

2.5. Prediction and Postprocessing

Performing predictions on overlapping patches have been shown to generate probability maps of increased reliability [15]. As in-house experiments have confirmed this aspect, predictions from 75% overlapping patches have been selected in the context of this work. Therefore, for overlapping regions, the average probability at a given voxel was computed to generate the output p-maps.

The predicted binary map was obtained by applying a decision threshold on the corresponding p-map. In our experiments, both the conventional 0.5 and the optimal decision threshold values were assessed. The median voxel-wise F1

score calculated on the first test set was used to select the optimal threshold associated with each loss function. Threshold values were evaluated over the range of 0 to 1 with an interval of 0.05. The unseen test set was used to assess the influence of these optimal thresholds compared to the conventional value of 0.5.

2.6. Evaluation metrics

Evaluation metrics were based on those presented in Gros et al. [2] and brain MS segmentation challenges [3]. Sensitivity, precision, and voxel-wise F1-score (V-F1) were used to measure the voxel-wise agreement between the ground truth and the binarized predictions. Similar metrics were defined at the lesion level, including lesion-wise sensitivity, lesion-wise precision, and lesion-wise F1 score (L-F1). We considered a candidate lesion as correctly detected if the lesion connected voxels overlapped the ground truth connected voxels by at least 25% and did not go outside by more than 70%. Subject-level sensitivity and specificity were also calculated to assess the capability of a model to classify subjects with and without ground truth lesions correctly. The Wilcoxon signed-rank test was used to determine whether the observed metrics differences between two methods were significant.

3. EXPERIMENTAL RESULTS




Fig. 2. Boxplots at each decision threshold value for BCE, SCT, Focal Tversky and Tversky losses. One point corresponds to one subject.

The results presented in table1 indicate that loss functions based on the standard DSC lead to high voxel-wise precision and low voxel-wise sensitivity scores, the predictions being




Fig. 3. From the left: ground truth (yellow) overlaid a sagittal T2-w image (from the unseen test set); segmentation with the Tversky Loss, and with the Dice Loss. The binary output was obtained by thresholding the p-map at the optimal value. FNs and FPs voxels are white and red, respectively; TP voxels are green.

biased towards the non-lesion class, i.e. many lesion voxels are missed yielding a high number of FN voxels. The use of the Generalised Dice Loss and loss functions based on the Tversky Index results in a better trade-off between voxel-wise precision and sensitivity. The highest V-F1 are achieved with the Focal Tversky Loss (0.398 on the threshold-optimization set (0.402 on the unseen set)) and Tversky Loss (0.424 (0.398)). The parameters of these loss functions were designed to weight sensitivity higher than precision, which resulted in a consistent improvement in the performance metrics on both datasets.

The detection and count of lesions is vital in monitoring MS patients undergoing disease-modifying treatments. In this context, lesion-wise metrics have more value for clinicians than voxel-wise metrics. As both FP and FN lesions are undesirable, the L-F1 is favored, i.e. the harmonic mean of precision and sensitivity. Similarly to the voxel-wise analysis, the highest L-F1 is observed for losses based on the Tversky Index (TL: 0.523 on the thr-optimization set (0.336 on the unseen set); FTL: 0.467 on the thr-optimization set (0.367 on the unseen set)).

In this study, the decision threshold was adjusted to maximize the V-F1. Figure 3 shows an example of the segmentation outputs obtained with the DSC Loss and the Tversky Loss on an image of the unseen test set using the optimized threshold values. The results presented in table2 suggest that optimizing the decision threshold can affect the overall performance of a model with a significant boost in both the V-F1 and L-F1 metrics. Significant improvements were reported for BCE (V-F1 +20% (+32%), L-F1 +37% (+27%)), the state-of-the-art SCT model (V-F1 +14% (+16%), L-F1 +20% (+39%)), and slight improvements for DSC (V-F1 +20% (+1.5%), L-F1 +37% (+1.5%)). For all these loss functions, the performance metrics improved by lowering the decision threshold to a value lower than 0.5.

In figure 2, we report the voxel-wise F1 score boxplots for the BCE, SCT, Focal Tversky and Tversky losses. While the Tversky Loss showed minor improvements (V-F1 +2% (+2%), L-F1 +0% (+6%)), the stability of its performance scores across the range of decision thresholds suggests that

Table 1. Metrics (median (mean \pm std)) obtained on the thr-optimization and the unseen subsets (first and second row associated with each loss function respectively), using 0.5 as a decision threshold value.

| Loss | L-sensitivity | L-precision | L-F1 | V-sensitivity | V-precision | V-F1 | P-sensitivity | P-specificity |
|------------------|---|--|--|--|--|--|----------------|----------------|
| BCE | 0 (0.13 \pm 0.27) (0.124 \pm 0.338) | 0 (0.214 \pm 0.323) 0 (0.276 \pm 0.38) | 0 (0.153 \pm 0.279) 0 (0.16 \pm 0.328) | 0.113 (0.139 \pm 0.156) 0.007 (0.087 \pm 0.228) | 0.691 (0.542 \pm 0.434) 0.381 (0.472 \pm 0.417) | 0.202 (0.207 \pm 0.208) 0.014 (0.14 \pm 0.242) | 0.357 0.308 | 0.413 0.913 |
| Combo | 0.427 (0.412 \pm 0.375) 0 (0.233 \pm 0.38) | 0.417 (0.39 \pm 0.308) 0 (0.299 \pm 0.298) | 0.5 (0.346 \pm 0.247) 0 (0.246 \pm 0.295) | 0.314 (0.323 \pm 0.25) 0.124 (0.179 \pm 0.338) | 0.484 (0.48 \pm 0.328) 0.611 (0.461 \pm 0.284) | 0.413 (0.345 \pm 0.219) 0.218 (0.23 \pm 0.248) | 0.714 0.462 | 0.913 0.522 |
| Dice | 0.438 (0.455 \pm 0.34) 0.333 (0.352 \pm 0.398) | 0.5 (0.431 \pm 0.345) 0.278 (0.282 \pm 0.341) | 0.411 (0.404 \pm 0.297) 0.297 (0.286 \pm 0.318) | 0.399 (0.417 \pm 0.263) 0.279 (0.311 \pm 0.294) | 0.517 (0.482 \pm 0.288) 0.477 (0.448 \pm 0.366) | 0.443 (0.395 \pm 0.2) 0.351 (0.306 \pm 0.27) | 0.786 0.615 | 0.283 0.304 |
| Focal Tversky | 0.5 (0.468 \pm 0.379) 0.4 (0.458 \pm 0.355) | 0.375 (0.349 \pm 0.29) 0.292 (0.331 \pm 0.435) | 0.467 (0.379 \pm 0.292) 0.367 (0.341 \pm 0.362) | 0.468 (0.444 \pm 0.275) 0.373 (0.364 \pm 0.275) | 0.395 (0.411 \pm 0.303) 0.395 (0.382 \pm 0.446) | 0.398 (0.353 \pm 0.208) 0.402 (0.332 \pm 0.275) | 0.714 0.654 | 0.37 0.13 |
| Generalized Dice | 0.25 (0.327 \pm 0.327) 0.464 (0.439 \pm 0.364) | 0.171 (0.26 \pm 0.313) 0.333 (0.295 \pm 0.3) | 0.222 (0.254 \pm 0.253) 0.382 (0.326 \pm 0.296) | 0.526 (0.532 \pm 0.294) 0.523 (0.471 \pm 0.302) | 0.351 (0.325 \pm 0.228) 0.383 (0.409 \pm 0.346) | 0.415 (0.354 \pm 0.212) 0.389 (0.361 \pm 0.229) | 0.643 0.654 | 0.522 0.283 |
| Log-Cosh | 0.5 (0.535 \pm 0.366) 0.225 (0.352 \pm 0.235) | 0.472 (0.465 \pm 0.338) 0.156 (0.292 \pm 0.434) | 0.45 (0.449 \pm 0.276) 0.188 (0.292 \pm 0.273) | 0.348 (0.336 \pm 0.185) 0.242 (0.267 \pm 0.118) | 0.564 (0.517 \pm 0.308) 0.446 (0.407 \pm 0.484) | 0.392 (0.368 \pm 0.183) 0.346 (0.287 \pm 0.178) | 0.857 0.538 | 0.674 0.413 |
| SCT | 0.225 (0.265 \pm 0.304) 0 (0.278 \pm 0.392) | 0.25 (0.314 \pm 0.32) 0 (0.362 \pm 0.352) | 0.243 (0.276 \pm 0.301) 0 (0.297 \pm 0.323) | 0.17 (0.204 \pm 0.181) 0.132 (0.222 \pm 0.28) | 0.8 (0.598 \pm 0.401) 0.764 (0.525 \pm 0.353) | 0.282 (0.288 \pm 0.232) 0.23 (0.265 \pm 0.248) | 0.643 0.462 | 0.13 0.674 |
| Tversky | 0.5 (0.545 \pm 0.329) 0.367 (0.399 \pm 0.41) | 0.5 (0.499 \pm 0.33) 0.25 (0.321 \pm 0.333) | 0.523 (0.487 \pm 0.278) 0.336 (0.318 \pm 0.305) | 0.41 (0.379 \pm 0.218) 0.3 (0.317 \pm 0.303) | 0.47 (0.469 \pm 0.255) 0.52 (0.421 \pm 0.31) | 0.424 (0.399 \pm 0.206) 0.398 (0.316 \pm 0.254) | 0.857 0.615 | 0.304 0.37 |

Table 2. Metrics (median (mean \pm std)) obtained on the thr-optimization and the unseen subsets (first and second row associated with each loss function respectively), using an optimized decision threshold value based on the DSC. Significant metric differences (p-value \leq 0.05) obtained by means of the 0.5 and optimal thresholds are highlighted in bold.

| Loss | Optimal threshold | L-sensitivity | L-precision | L-F1 | V-sensitivity | V-precision | V-F1 | P-sensitivity | P-specificity |
|------------------|-------------------|--|---|---|--|--|--|----------------|----------------|
| BCE | 0.15 0.15 | 0.369 (0.408 \pm 0.366) 0.267 (0.331\pm0.355) | 0.417 (0.386 \pm 0.337) 0.25 (0.272 \pm 0.304) | 0.365 (0.354 \pm 0.289) 0.265 (0.277 \pm 0.289) | 0.438 (0.469 \pm 0.294) 0.286 (0.337\pm0.323) | 0.352 (0.394 \pm 0.274) 0.396 (0.36 \pm 0.321) | 0.404 (0.361 \pm 0.198) 0.333 (0.293\pm0.244) | 0.714 0.577 | 0.239 0.239 |
| Combo | 0.05 0.05 | 0.25 (0.37 \pm 0.401) 0.267 (0.312\pm0.354) | 0.225 (0.218 \pm 0.213) 0.131 (0.2 \pm 0.257) | 0.211 (0.26 \pm 0.257) 0.184 (0.218 \pm 0.258) | 0.475 (0.419 \pm 0.293) 0.284 (0.275\pm0.285) | 0.375 (0.35 \pm 0.272) 0.42 (0.392 \pm 0.329) | 0.419 (0.337 \pm 0.226) 0.3 (0.271\pm0.245) | 0.643 0.538 | 0.174 0.174 |
| Dice | 0.3 0.3 | 0.5 (0.491 \pm 0.358) 0.333 (0.391 \pm 0.391) | 0.5 (0.377 \pm 0.277) 0.292 (0.305 \pm 0.338) | 0.5 (0.387 \pm 0.26) 0.31 (0.308 \pm 0.318) | 0.441 (0.446 \pm 0.265) 0.324 (0.343\pm0.313) | 0.461 (0.441 \pm 0.277) 0.421 (0.411\pm0.326) | 0.459 (0.392 \pm 0.201) 0.364 (0.313 \pm 0.276) | 0.786 0.615 | 0.261 0.261 |
| Focal Tversky | 0.15 0.15 | 0.5 (0.466 \pm 0.356) 0.359 (0.41 \pm 0.382) | 0.31 (0.333 \pm 0.326) 0.171 (0.244\pm0.269) | 0.408 (0.35 \pm 0.263) 0.261 (0.281\pm0.278) | 0.546 (0.517 \pm 0.279) 0.429 (0.444\pm0.344) | 0.336 (0.352 \pm 0.292) 0.23 (0.294\pm0.263) | 0.421 (0.351 \pm 0.219) 0.306 (0.311 \pm 0.249) | 0.714 0.654 | 0.217 0.217 |
| Generalized Dice | 0.75 0.75 | 0.382 (0.379 \pm 0.332) 0.45 (0.441 \pm 0.394) | 0.196 (0.305 \pm 0.345) 0.279 (0.351 \pm 0.372) | 0.245 (0.295 \pm 0.273) 0.341 (0.353 \pm 0.333) | 0.486 (0.49 \pm 0.293) 0.438 (0.414\pm0.331) | 0.386 (0.379 \pm 0.265) 0.45 (0.464\pm0.32) | 0.443 (0.367 \pm 0.213) 0.406 (0.359 \pm 0.257) | 0.714 0.654 | 0.196 0.196 |
| Log-Cosh | 0.2 0.2 | 0.5 (0.578 \pm 0.347) 0.208 (0.326 \pm 0.37) | 0.311 (0.42 \pm 0.339) 0.121 (0.197 \pm 0.234) | 0.452 (0.419 \pm 0.25) 0.168 (0.236 \pm 0.27) | 0.409 (0.412 \pm 0.232) 0.344 (0.347\pm0.335) | 0.456 (0.449 \pm 0.301) 0.323 (0.333\pm0.323) | 0.432 (0.375 \pm 0.193) 0.368 (0.289 \pm 0.252) | 0.857 0.538 | 0.196 0.196 |
| SCT | 0.05 0.05 | 0.369 (0.459 \pm 0.399) 0.4 (0.418 \pm 0.404) | 0.438 (0.398 \pm 0.334) 0.268 (0.339 \pm 0.351) | 0.441 (0.414 \pm 0.348) 0.382 (0.339 \pm 0.316) | 0.334 (0.3 \pm 0.233) 0.296 (0.342\pm0.304) | 0.615 (0.494 \pm 0.342) 0.628 (0.494\pm0.348) | 0.42 (0.361 \pm 0.26) 0.388 (0.347\pm0.253) | 0.714 0.615 | 0.326 0.326 |
| Tversky | 0.3 0.3 | 0.55 (0.581 \pm 0.35) 0.45 (0.439 \pm 0.389) | 0.444 (0.457 \pm 0.299) 0.25 (0.329 \pm 0.315) | 0.523 (0.464 \pm 0.238) 0.4 (0.337 \pm 0.308) | 0.456 (0.419 \pm 0.235) 0.358 (0.356\pm0.318) | 0.424 (0.43 \pm 0.25) 0.454 (0.42 \pm 0.333) | 0.443 (0.399 \pm 0.207) 0.41 (0.332\pm0.242) | 0.857 0.654 | 0.261 0.261 |

it is an apt solution for our task and dataset. A similar conclusion can be drawn for the Focal Tversky Loss, for which a minor deterioration was observed across datasets while maintaining a rather stable behavior over the threshold range. A significant decrease in model performance with conventional 0.5-thresholding can be observed for BCE and SCT model. Perturbations of as little as ± 0.05 can induce a median reduction in V-F1 score of 10%.

Such observations indicate that when dealing with heterogeneous data and prone to high-class imbalance, certain schemes of the loss function and decision threshold are better suited to cope with the bias-variance trade-off and, thus, to accommodate the metrics of interest, i.e. V-F1 and L-F1. Similarly to voxel-wise and lesion-wise metrics, Tversky Loss tends to have a better trade-off when looking at the P-sensitivity and P-specificity (see table1 and table2). While Focal Tversky Loss and Generalized Dice Loss have a higher patient detection rate, both losses struggle when presented with non-lesion subjects.

To put the work in context, our results indicate that the combination of the Tversky Loss with an adjusted decision threshold performs very well compared to the latest results in T2-w MS spinal cord lesion segmentation [2].

4. CONCLUSION

This study evaluates segmentation loss functions proposed in the literature to mitigate the issue of imbalanced data, with reference to the complex task of MS lesion segmentation in spinal cord MRI. In addition, the adjustment of the decision threshold on the output p-map is explored to further improve the performance metrics of interest. Results indicate that loss functions based on the Tversky Index and with minimal adjustment of the decision threshold can yield higher median scores and less dispersed output scores than more standard settings. Future work will aim to understand how the variability due to acquisition scanners and protocol, and patient population, can affect the test metrics and can be eventually addressed.

5. COMPLIANCE WITH ETHICAL STANDARDS

All patients provided written consent and were included in one of the four clinical studies registered on clinicaltrials.gov (NCT02889965, NCT02117375, NCT04220814, NCT04918225), and compliant with French data confidentiality regulations. The study was approved by the relevant

ethics committee.

6. ACKNOWLEDGMENTS

This work was supported by a grant provided by the Ministry of Health in France, and handled by the Agence Nationale de la Recherche within the framework of the Investments for the Future program, under the reference ANR-10-COHO-002, by the Eugène Devic EDMUS Foundation against multiple sclerosis and by the ARSEP Foundation. Authors thank the Observatoire Français de la Sclérose en Plaques (OFSEP) for their fundamental input.

7. REFERENCES

- [1] S. Leguy, B. Combès, E. Bannier, and A. Kerbrat, “Prognostic value of spinal cord MRI in multiple sclerosis patients,” *Revue Neurologique*, vol. 177, no. 5, pp. 571–581, May 2021.
- [2] C. Gros et al., “Automatic segmentation of the spinal cord and intramedullary multiple sclerosis lesions with convolutional neural networks,” *NeuroImage*, vol. 184, pp. 901–915, Jan. 2019.
- [3] O. Commowick et al., “Objective evaluation of multiple sclerosis lesion segmentation using a data management and processing infrastructure,” *Scientific Reports*, vol. 8, no. 1, Sept. 2018.
- [4] J. Ma, J. Chen, M. Ng, R. Huang, Y. Li, C. Li, X. Yang, and A.L. Martel, “Loss odyssey in medical image segmentation,” *Medical Image Analysis*, vol. 71, pp. 102035, 2021.
- [5] S. Jadon, “A survey of loss functions for semantic segmentation,” in *2020 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB)*. Oct. 2020, IEEE.
- [6] N. Bice, N. Kirby, R. Li, D. Nguyen, T. Bahr, C. Kabat, P. Myers, N. Papanikolaou, and M. Fakhreddine, “A sensitivity analysis of probability maps in deep-learning-based anatomical segmentation,” *J. Appl. Clin. Med. Phys.*, vol. 22, no. 8, pp. 105–119, Aug. 2021.
- [7] F. Milletari, N. Navab, and S. Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation,” 2016.
- [8] S.A. Taghanaki, Y. Zheng, S.K. Zhou, B. Georgescu, P. Sharma, D. Xu, D. Comaniciu, and G. Hamarneh, “Combo loss: Handling input and output imbalance in multi-organ segmentation,” *Computerized Medical Imaging and Graphics*, vol. 75, pp. 24–33, 2019.
- [9] C.H. Sudre, W. Li, T. Vercauteren, S. Ourselin, and M. Jorge Cardoso, “Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations,” in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pp. 240–248. Springer International Publishing, 2017.
- [10] S.S.M. Salehi, D. Erdogmus, and A. Gholipour, “Tversky loss function for image segmentation using 3d fully convolutional deep networks,” in *Machine Learning in Medical Imaging*, pp. 379–387. Springer International Publishing, 2017.
- [11] N. Abraham and N.M. Khan, “A novel focal tversky loss function with improved attention u-net for lesion segmentation,” *CoRR*, vol. abs/1810.07842, 2018.
- [12] C. Gros, B. De Leener, S.M. Dupont, A.R. Martin, M.G. Fehlings, R. Bakshi, S. Tummala, V. Auclair, D.G. McLaren, V. Callot, J. Cohen-Adad, and M. Sdika, “Automatic spinal cord localization, robust to MRI contrasts using global curve optimization,” *Medical Image Analysis*, vol. 44, pp. 215–227, Feb. 2018.
- [13] L.G. Nyul, J.K. Udupa, and X. Zhang, “New variants of a method of MRI scale standardization,” *IEEE Transactions on Medical Imaging*, vol. 19, no. 2, pp. 143–150, 2000.
- [14] Ö. Çiçek, A. Abdulkadir, S.S. Lienkamp, T. Brox, and O. Ronneberger, “3d u-net: Learning dense volumetric segmentation from sparse annotation,” in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, pp. 424–432. Springer International Publishing, 2016.
- [15] N. Pielawski and C. Wählby, “Introducing hann windows for reducing edge-effects in patch-based image segmentation,” *PLOS ONE*, vol. 15, no. 3, pp. e0229839, Mar. 2020.