



# Learning joint segmentation of tissues and brain lesions from task-specific hetero-modal domain-shifted datasets



Reuben Dorent<sup>a,\*</sup>, Thomas Booth<sup>a,b</sup>, Wenqi Li<sup>a,e</sup>, Carole H. Sudre<sup>a,c,d</sup>, Sina Kafiabadi<sup>b</sup>, Jorge Cardoso<sup>a</sup>, Sebastien Ourselin<sup>a</sup>, Tom Vercauteren<sup>a</sup>

<sup>a</sup> King's College London, School of Biomedical Engineering & Imaging Sciences, St. Thomas' Hospital, London, United Kingdom

<sup>b</sup> Department of Neuroradiology, King's College Hospital NHS Foundation Trust, London, United Kingdom

<sup>c</sup> Dementia Research Centre, UCL Institute of Neurology, UCL, London, United Kingdom

<sup>d</sup> Department of Medical Physics, UCL, London, United Kingdom

<sup>e</sup> NVIDIA, Cambridge, United Kingdom

## ARTICLE INFO

### Article history:

Received 3 October 2019

Revised 9 September 2020

Accepted 25 September 2020

Available online 9 October 2020

### Keywords:

Joint learning

Domain adaptation

Multi-Task learning

Multi-Modal

## ABSTRACT

Brain tissue segmentation from multimodal MRI is a key building block of many neuroimaging analysis pipelines. Established tissue segmentation approaches have, however, not been developed to cope with large anatomical changes resulting from pathology, such as white matter lesions or tumours, and often fail in these cases. In the meantime, with the advent of deep neural networks (DNNs), segmentation of brain lesions has matured significantly. However, few existing approaches allow for the joint segmentation of normal tissue and brain lesions. Developing a DNN for such a joint task is currently hampered by the fact that annotated datasets typically address only one specific task and rely on task-specific imaging protocols including a task-specific set of imaging modalities. In this work, we propose a novel approach to build a joint tissue and lesion segmentation model from aggregated task-specific hetero-modal domain-shifted and partially-annotated datasets. Starting from a variational formulation of the joint problem, we show how the expected risk can be decomposed and optimised empirically. We exploit an upper bound of the risk to deal with heterogeneous imaging modalities across datasets. To deal with potential domain shift, we integrated and tested three conventional techniques based on data augmentation, adversarial learning and pseudo-healthy generation. For each individual task, our joint approach reaches comparable performance to task-specific and fully-supervised models. The proposed framework is assessed on two different types of brain lesions: White matter lesions and gliomas. In the latter case, lacking a joint ground-truth for quantitative assessment purposes, we propose and use a novel clinically-relevant qualitative assessment methodology.

© 2020 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

## 1. Introduction

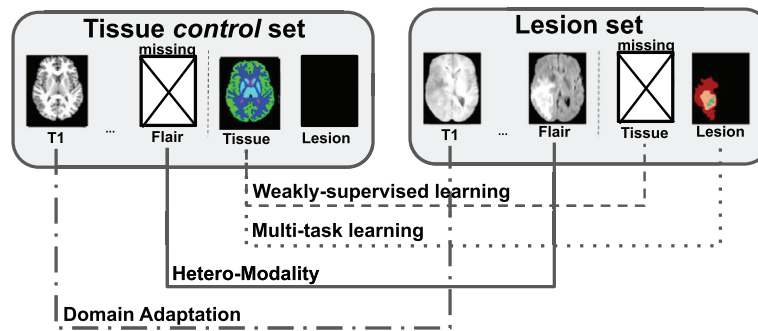
Traditional approaches for tissue segmentation used in brain segmentation / parcellation software packages such as FSL (Jenkinson et al., 2012), SPM (Ashburner and Friston, 2000) or NiftySeg (Cardoso et al., 2015) are based on subject-specific optimisation. FSL and SPM fit a Gaussian Mixture Model to the MR intensities using either a Markov Random Field (MRF) or tissue prior probability maps as regularisation. Alternatively, multi-atlas methods rely on label propagation and fusion from multiple fully-annotated images, i.e. atlases, to the target image (Cardoso

et al., 2015; Iglesias and Sabuncu, 2015). These methods typically require extensive pre-processing, e.g. skull stripping, correction of bias field and registration. They are also often time-consuming and are inherently only adapted for brains devoid of large anatomical changes induced by pathology, such as white matter lesions and brain tumours. Indeed, it has been shown that the presence of lesions can significantly distort any registration output (Sdika and Pelletier, 2009). Similarly, lesions introduce a bias in the MRF. This leads to a performance degradation in the presence of lesions for brain volume measurement (Battaglini et al., 2012) and any subsequent analysis.

While quantitative analysis is expected to play a key role in improving the diagnosis and follow-up evaluations of patients with brain lesions, current tools mostly focus on quantification of the lesions themselves and effectively discard contextual tissue

\* Corresponding author.

E-mail address: [reuben.dorent@kcl.ac.uk](mailto:reuben.dorent@kcl.ac.uk) (R. Dorent).



**Fig. 1.** Tissue and lesion segmentation, a problem at the intersection of multiple branches of Machine Learning: Multi-Task Learning (tissue + lesion segmentation), Weakly-Supervised Learning (missing annotations), Hetero-Modality (missing modalities), Domain Adaptation (different acquisition protocols).

information. Existing quantitative neuroimaging approaches allow the extraction of imaging biomarkers such as the largest diameter, volume, and count of the lesions. Such automatic segmentation of the lesions promises to speed up and improve the clinical decision-making process but more refined analysis would be feasible from tissue classification and region parcellation. In particular, brain atrophy at a global level (Popescu et al., 2013; Giorgio and De Stefano, 2013), at a cerebral level (Bermel and Bakshi, 2006), and, even more specifically, at the grey matter level (Geurts et al., 2012) have been correlated with the speed of disease progression and with physical disability (Fisniku et al., 2008). Consequently, atrophied tissue volumes in the presence of lesions are clinically relevant imaging markers (Dwyer et al., 2018). We believe that, although very few work have addressed this problem yet, a joint model for lesion and tissue segmentation is expected to bring significant clinical benefits. As representative exemplars of the technical challenges involved to build such joint models, we focus, in this work, on patients with white matter lesions or brain tumours.

Deep Neural Networks (DNNs) have become the state-of-the-art for most segmentation tasks (Simpson et al., 2019) and one would now expect these to be able to jointly perform brain tissue and pathology segmentation. However, annotated databases required to train DNNs are usually dedicated to a single task (either brain tissue segmentation or pathology delineation). In addition, the information required for brain tissue or pathology segmentation may come from different scans, leading to hetero-modal (i.e. more than one set of input imaging sequences) datasets. While  $T_1$ -weighted images provide the best grey/white matter contrast for the delineation of anatomical tissue,  $T_2$ -weighted sequences are usually more sensitive to pathological changes (Bitar et al., 2006). Choice of the used sequence or combination of sequences may also differ across pathologies.  $T_2$ -weighted FLAIR images are often used for the assessment of white matter lesions (Maillard et al., 2013) while a combination of  $T_1$  contrast-enhanced ( $T_1c$ ),  $T_2$  and FLAIR is often preferred for the characterisation of gliomas (Wen et al., 2010). Similarly, the scans may have been acquired with different magnetic resonance parameters leading to differences in resolution and contrast among databases. Consequently, the data distribution may differ between the datasets, i.e. the datasets may be domain-shifted. Given 1) the large amount of resources, time and expertise required to annotate medical images, 2) the varying imaging requirement to support each individual task, and 3) the availability of task-specific databases, it is unlikely that large databases for every joint problem, such as lesion and tissue segmentation, will become available. There is thus a need to exploit existing task-specific databases to address the joint problems. Learning a joint model from task-specific hetero-modal and domain-shifted datasets is nonetheless challenging. As shown in Fig. 1, this problem lies at the intersection of Multi-Task Learning (Zhang and

Yang, 2017), Domain Adaptation (Ben-David et al., 2010; Zhao et al., 2019) and Weakly Supervised Learning (Oquab et al., 2015; Bilen and Vedaldi, 2016; Xu et al., 2014) with singularities making individual methods from these underpinning fields insufficient to address it completely, as explained in more depth in the related work section 2.

Our approach is rooted in all these sub-domains of deep learning. The main contributions are summarised as follows:

1. We propose a joint model that performs tissue and lesion segmentation as a unique joint task and thus exploits the interdependence between the lesion and tissue segmentation tasks. Starting from a variational formulation of the joint problem, we exploit the disjoint nature of the label sets to propose a practical decomposition of the joint loss, transforming the multi-class segmentation problem into a multi-task problem.
2. We introduce feature channel averaging across modalities to adapt existing networks for our hetero-modal problem.
3. We develop a new method to minimise the expected risk under the constraint of missing modalities. Under the assumption that the network is not affected by a potential domain shift, we show that the expected risk can be further decomposed and minimised via a tractable upper bound. To our knowledge, no such optimisation method for missing modalities in deep learning has been published before.
4. Given that, in practice, the heterogeneous task-specific datasets may have been acquired with different protocols, i.e. they are domain-shifted, we integrate several existing DA techniques in our framework. These methods are based on data augmentation and adversarial training, or pseudo-healthy brain generation.
5. We demonstrate the performance of our joint approach on two clinical use cases: White matter lesions and gliomas. Our method outperforms a fully-supervised model trained on a smaller fully-annotated dataset for white matter lesions. To assess the performance of the joint model for tissue and glioma segmentation, for which no ground-truth is available, we propose a new qualitative evaluation protocol based on the ASPECTS score (Barber et al., 2000). Higher accuracy is obtained compared to time-consuming pipelines that require to mask the lesions using manual annotations.
6. Experiments show that generating pseudo-healthy annotated scans outperforms the other DA techniques, even with very few pseudo-healthy annotated scans.

This work is a substantial extension of our conference paper (Dorent et al., 2019b). Improvements include: 1) Additional mathematical proofs; 2) integration and validation of three different domain adaptation techniques to cope with domain-shifted datasets; 3) new experiments on joint brain tissue and glioma seg-

mentation; and 4) a new quantitative evaluation protocol for assessing tissue segmentation in the absence of ground-truth.

## 2. Related work

Multi-Task Learning (MTL) aims to perform several tasks simultaneously, on a single dataset, by extracting some form of common knowledge or representation and introducing a task-specific back-end. When relying on DNN for MTL, the first layers of the network are typically shared, while the last layers are trained for the different tasks (Ruder, 2017). MTL has been successfully applied to medical imaging for segmentation (Bragman et al., 2018; Moeskops et al., 2016) combined with other tasks such as detection (Saha et al., 2019) or classification (Chen et al., 2019; Le et al., 2019). The global loss function is a weighted sum of task-specific loss functions. Recently, Kendall and Gal (2017) proposed a Bayesian parameter-free method to estimate the MTL loss weights and Bragman et al. (2018) extended it to spatially adaptive task weighting and applied it to medical imaging. Although the aforementioned approaches generate different outputs from the same features, no direct interaction between the task-specific outputs is modelled in these techniques. While a joint tissue and lesion segmentation can be pursued in practice, a strong underpinning assumption is that the two outputs are conditionally independent. Consequently, these approaches do not address the problem of aggregating these outputs to generate a joint segmentation. Moreover, MTL approaches, such as (Moeskops et al., 2018; Roulet et al., 2019), do not provide any mechanism for dealing with hetero-modal datasets or changes in imaging characteristics across task-specific databases.

Domain Adaptation (DA) is a solution for dealing with domain-shifted datasets, i.e. datasets acquired with different settings. A classical strategy consists in learning a domain-invariant feature representation of the data. Csurka (2017) proposed an extensive review of these methods in deep learning. Some DA approaches have been developed to tackle a specific and identified shift. For example, data augmentation has been used for shifts caused by different MR bias fields (Sudre et al., 2017) or the presence of motion artefacts (Shaw et al., 2019); Havaei et al. (2016) and Dorent et al. (2019a) proposed network architectures for dealing with missing modalities that encodes each modality into a shared modality-agnostic latent space. Recent studies have proposed to learn a mapping between healthy and disease scans, using CycleGANs (Xia et al., 2019; Sun et al., 2018) or Variational Autoencoders (Chen and Konukoglu, 2018). Although these techniques have shown promising results, they are inherently limited to a specific type of shift. Combining causes of shift, for instance the presence/absence of lesions with different protocols of acquisition, remains an unsolved problem. In contrast, general DA approaches do not make assumptions about the nature of the shift. These methods aim to directly minimise the discrepancy between the feature distributions across the domains. Distribution dissimilarity can be assessed using correlation distances (Sun et al., 2016) or maximum mean discrepancy (Pan et al., 2011; Long et al., 2015). However, more recent techniques are mostly focused on adversarial methods, achieving promising results in medical imaging (Kamnitsas et al., 2017; Dou et al., 2018; Orbes-Arteaga et al., 2019). However, these methods are usually focus on solving a single task across domain.

Weakly-supervised Learning (WSL) deals with missing, inaccurate, or inexact annotations. Our problem is a particular case of learning with missing labels since each task-specific dataset provides a set of labels where the two sets are disjoint. Li and Hoiem (2017) proposed a method to learn a new task from a model trained on another task. This method combines DA through transfer learning and MTL. In the end, two models are created: One for the first task and one for the second one. Kim et al. (2018) ex-

tended this approach by using a knowledge distillation loss in order to create a unique multi-task model. This aims to alternatively learn one task without forgetting the other one. The WSL problem was thus decomposed into an MTL problem with aforementioned limitations for our specific use case.

This work proposes a new framework to perform a joint segmentation while dealing with task-specific, domain-shifted and hetero-modal datasets.

## 3. Tissue and lesion segmentation learning from hetero-modal and task-specific datasets: Problem definition

In order to develop a joint model, we first propose a mathematical variational formulation of the problem and introduce a network architecture to leverage existing hetero-modal and task-specific datasets for tissue and lesion segmentation.

### 3.1. Formal problem statement

Let  $x = (x^1, \dots, x^M) \in \mathcal{X} = \mathbb{R}^{N \times M}$  be a vectorized multimodal image and  $y \in \mathcal{Y} = \{0, \dots, C\}^N$  its associated tissue and lesion segmentation map.  $N$ ,  $M$  and  $C$  are respectively the number of voxels, modalities and classes. Note that images modalities are assumed to be co-registered and resampled in the same coordinate space containing  $N$  voxels. Our goal is to determine a predictive function, parametrised by the weights  $\theta \in \Theta$ ,  $h_\theta : \mathcal{X} \mapsto \mathcal{Y}$  that minimises the discrepancy between the ground truth label vector  $y$  and the prediction  $h_\theta(x)$ . Let  $\mathcal{L}$  be a loss function that estimates this discrepancy. Following the formalism used by Bottou et al. (2018), given a probability distribution  $\mathcal{D}$  over  $(\mathcal{X}, \mathcal{Y})$  and random variables under this distribution, we want to find  $\theta^*$  such that:

$$\theta^* = \arg \min_{\theta} \mathbb{E}_{(x,y) \sim \mathcal{D}} [\mathcal{L}(h_\theta(x), y)] \quad (1)$$

As is the norm in data-driven learning, we do not have access to the true joint probability  $\mathcal{D}$ . In supervised learning, the common method is to estimate the expected risk using training samples. Given a set of  $n \in \mathbb{N}$  independently drawn multimodal scans with their associated tissue and lesion segmentation map  $\{(x_k, y_k)\}_{k=1}^n$ , we want to find  $\theta^*$  that minimises the empirical risk:

$$\theta^* = \arg \min_{\theta} \sum_{k=1}^n \mathcal{L}(h_\theta(x_k), y_k) \quad (2)$$

However, in our multitask scenario, we cannot directly estimate the empirical risk since we do not have access to a fully annotated dataset for the joint task. Instead, we propose to leverage task-specific and hetero-modal datasets.

### 3.2. Task-specific and hetero-modal datasets

Let us assume that we have access to two datasets with either the tissue annotations  $y^T$  or the lesion annotations  $y^L$  (task-specificity). Let

$$S_{control} = \left\{ \left( (x_k^1, \dots, x_k^{M_T}), y_k^T \right) \right\}_{k=1}^{n_T}$$

$$S_{lesion} = \left\{ \left( (x_k^1, \dots, x_k^{M_L}), y_k^L \right) \right\}_{k=1}^{n_L}$$

denote these two training sets, where  $M_T$ ,  $M_L$ ,  $n_T$  and  $n_L$  are respectively the number of modalities in the *control* and the *lesion* sets and the size of these sets. Note that although we use the term *control* for convenience, we may expect to observe pathology with "diffuse" anatomical impact, e.g. from dementia. In addition, for the clarity of presentation, we highlight that the considered *lesions* in this work are either White Matter Hyperintensities (WMH) or gliomas.




Fig. 2. The proposed fully-convolutional network architecture: A mix a 3D U-Net (Çiçek et al., 2016a) and HeMIS (Havaei et al., 2016).

Since such datasets are typically developed in the scope of either tissue or lesion segmentation (but not both), the set of observed modalities may vary from one dataset to another (hetero-modality). Importantly, in this work, we consider that only T<sub>1</sub> scans are provided in the *control* dataset, while the *lesion* set contains either 1) the T<sub>1</sub> and the FLAIR scans for WMH segmentation, or 2) the T<sub>1</sub>, contrast-enhanced T<sub>1</sub> (T<sub>1c</sub>), T<sub>2</sub> and FLAIR scans for glioma segmentation. The full set of modalities is consequently given by the modalities in the *lesion* set, while the *control* dataset will have missing modalities. In our specific use cases, the T<sub>1</sub> modality is a shared modality across the different datasets. It will nonetheless be apparent that our method can be trivially adapted for other shared modalities.

### 3.3. On the distribution $\mathcal{D}$ in the context of heterogeneous databases

As we expect different distributions across heterogeneous databases, two probability distributions of  $(X, Y)$  over  $(\mathcal{X}, \mathcal{Y})$  can be distinguished:

- under  $\mathcal{D}_{control}$ ,  $(X, Y)$  corresponds to a multimodal scan and joint segmentation map of a patient without lesions ( $Y$  effectively being a tissue segmentation map).
- under  $\mathcal{D}_{lesion}$ ,  $(X, Y)$  corresponds to a multimodal scan and joint lesion and tissue segmentation map of a patient with lesions.

Since traditional tissue segmentation methods are not adapted in the presence of lesions, the most important and challenging distribution  $\mathcal{D}$  to address is the one for patients with lesions,  $\mathcal{D}_{lesion}$ . In the remainder of this work, we thus assume that:

$$\mathcal{D} \triangleq \mathcal{D}_{lesion}. \quad (\text{H}_1)$$

### 3.4. Hetero-modal network architecture

In order to learn from hetero-modal datasets, we need a network architecture that allows for missing modalities. Specifically, the input modalities are either a T<sub>1</sub> scan or a full set of modalities. To deal with missing modalities, arithmetic operations are employed, as originally proposed in HeMIS (Havaei et al., 2016). The network architecture is based on a U-Net (Çiçek et al., 2016b), as shown in Fig. 2. Note that, while the proposed method requires a hetero-modal network, any specific architecture can be used. The proposed network is composed of two input branches, one for the T<sub>1</sub> scan and one for the full set of modalities. Although HeMIS originally proposed to encode each modality independently, i.e. one

branch per modality, we experimentally found higher performance with these two branches. In the presence of the full set of modalities, features extracted from the T<sub>1</sub> scan and all the modalities are averaged. Consequently, the network allows for missing modalities, i.e. is hetero-modal. This hetero-modal network with weights  $\theta$  is used to capture the predictive function  $h_\theta$  that can accept either T<sub>1</sub> or the full set of modalities as input.

## 4. Optimising tissue and lesion segmentation as a joint task

Given the mathematical formulation of the problem and the hetero-modal network architecture, we propose a method to empirically optimise the joint problem of tissue and lesion segmentation.

### 4.1. Loss decomposition

Let  $C_T$ ,  $C_L$  and 0 be respectively the set of tissue classes and lesion classes and the value of the background class in the segmentation masks. Since  $C_T$  and  $C_L$  are disjoint, the segmentation map  $y$  can be decomposed into two segmentation maps  $y = y^L + y^T$  with  $y^T \in C_T \cup \{0\}$ ,  $y^L \in C_L \cup \{0\}$ .

Let us assume that the loss function  $\mathcal{L}$  can also be decomposed into a tissue loss function  $\mathcal{L}^T$  and a lesion loss function  $\mathcal{L}^L$ . This is common for multi-class segmentation loss functions in particular for those with *one-versus-all* strategies (e.g. Dice loss, Jaccard loss, Generalized Cross-Entropy). Then, the joint and multi-class segmentation problem can be formulated as a multi-task problem:

$$\mathcal{L}(h_\theta(x), y) = \mathcal{L}^T(h_\theta(x), y^T) + \mathcal{L}^L(h_\theta(x), y^L) \quad (\text{H}_2)$$

In combination with (H<sub>1</sub>), (1) can be rewritten as:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \underbrace{\mathbb{E}_{\mathcal{D}_{lesion}} [\mathcal{L}^T(h_\theta(x), y^T)]}_{\mathcal{R}^T} + \underbrace{\mathbb{E}_{\mathcal{D}_{lesion}} [\mathcal{L}^L(h_\theta(x), y^L)]}_{\mathcal{R}^L} \quad (3)$$

While the second expected risk  $\mathcal{R}^L$  can be estimated using the full set of modalities and the lesion annotations provided in the *lesion* dataset, the first expected risk  $\mathcal{R}^T$  appears to be intractable due to the missing tissue annotations in the *lesion* dataset. In the next sections, we first propose an upper bound of the expected tissue risk  $\mathcal{R}^T$  and then a means to estimate this upper bound using the *control* dataset.



## 4.2. Upper bound of the expected tissue risk $\mathcal{R}^T$

Although, thanks to its hetero-modal architecture,  $h_\theta$  may handle inputs with varying number of modalities, the current decomposition (3) assumes that all the modalities of  $x$  are available for evaluating the loss. In our scenario, the *control* set of scans with tissue annotations only contains the  $T_1$ scans. Consequently, as we do not have all the modalities with tissue annotations, and as naively evaluating a loss with missing modalities would lead to a bias, estimating  $\mathcal{R}^T$  is not straightforward.

Let us assume that the tissue loss function  $\mathcal{L}^T$  satisfies the triangle inequality:

$$\forall (a, b, c) \in \mathcal{Y}^3 : \mathcal{L}^T(a, c) \leq \mathcal{L}^T(a, b) + \mathcal{L}^T(b, c) \quad (\text{H}_3)$$

Although not all losses satisfy (H<sub>3</sub>), it is known that the binary Jaccard is a distance (Späth, 1981; Kosub, 2018) and thus satisfies the triangle inequality.

### Definition 4.1. (Binary Jaccard distance)

The binary Jaccard distance  $J_{bin}$  is defined such that:

$$\forall a, b \in \{0, 1\}^N, J_{bin}(a, b) = 1 - \frac{\sum_{i=1}^N a_i b_i}{\sum_{i=1}^N a_i + b_i - a_i b_i} \quad (4)$$

However, network outputs are typically pseudo-probabilities, and the soft version of (4) does not satisfy the triangle inequality. To satisfy (H<sub>3</sub>), we extend the binary Jaccard distance to a multi-class probabilistic formulation that coincides with the binary Jaccard for binary inputs but preserves the metric property for probabilistic inputs.

### Definition 4.2. (Probabilistic multi-class Jaccard distance)

Let  $C$  be the number of classes in  $\mathcal{C}$ ,  $N$  be the number of voxels and  $\mathcal{P} \subset [0, 1]^{C \times N}$  denote the set of probability vector maps such that for any  $p = (p_{c,i})_{c \in C, i \in [0, N]} \in \mathcal{P}$ :

$$\forall i \in [0; N], \sum_{c \in C} p_{c,i} = 1$$

The probabilistic multi-class Jaccard distance is defined for any  $(u, v) \in \mathcal{P}^2$  as:

$$\mathcal{J}(u, v) = \sum_{c \in C} \omega_c \frac{2 \sum_{i=1}^N |u_{c,i} - v_{c,i}|}{\sum_{i=1}^N |u_{c,i} + |v_{c,i}| + |u_{c,i} - v_{c,i}|} \quad (5)$$

where  $\omega_c$  are the class weights such that with  $\sum_{c \in C} \omega_c = 1$

As shown in A.1, the binary and probabilistic Jaccard distance coincide on the set of binary vectors  $\{0, 1\}^N$ . Furthermore, (5) satisfies (H<sub>3</sub>).

**Lemma 4.1.** *The probabilistic multi-class Jaccard distance is a distance and thus satisfies the triangle inequality.*

**Proof.** The proof, detailed in A.2, follows from the Steinhaus transform (Späth, 1981) applied to the metric space  $([0, 1]^N, d_1)$  where  $d_1$  is the distance induced by the  $L_1$  norm.  $\square$

Under (H<sub>3</sub>),  $\mathcal{L}^T$  satisfies the following inequality:

$$\mathcal{L}^T(h_\theta(x), y^T) \leq \mathcal{L}^T(h_\theta(x), h_\theta(x^{T_1})) + \mathcal{L}^T(h_\theta(x^{T_1}), y^T) \quad (6)$$

where  $x^{T_1}$  denotes the  $T_1$ scan associated to  $x$ . Consequently, we find an upper bound of the expected tissue risk:

$$\mathcal{R}^T(\theta) \leq \underbrace{\mathbb{E}_{\mathcal{D}_{lesion}}[\mathcal{L}^T(h_\theta(x), h_\theta(x^{T_1}))]}_{\mathcal{R}_{T_1 \rightarrow Full}^T} + \underbrace{\mathbb{E}_{\mathcal{D}_{lesion}}[\mathcal{L}^T(h_\theta(x^{T_1}), y^T)]}_{\mathcal{R}_{T_1}^T} \quad (7)$$

Minimising  $\mathcal{R}_{T_1}^T$  enforces the network to generate accurate tissue segmentation using only  $T_1$  as input. Minimising  $\mathcal{R}_{T_1 \rightarrow Full}^T$  encourages consistency between the outputs when given only  $T_1$  or the

full set of modalities as input. This latter term allows for transferring the knowledge learnt on the  $T_1$ scan to the full set of modalities.

An empirical estimation of  $\mathcal{R}_{T_1 \rightarrow Full}^T$  can be obtained by comparing the network outputs using either  $T_1$  or the full set of modalities as input. In contrast,  $\mathcal{R}_{T_1}^T$  requires comparison of inference done, under  $\mathcal{D}_{lesion}$ , from  $T_1$  inputs with ground truth tissue maps  $y^T$ . While this provides a step towards a practical evaluation of  $\mathcal{R}^T$ , we still face the challenge of not having tissue annotations  $y^T$  under  $\mathcal{D}_{lesion}$ .

## 4.3. Estimating $\mathcal{R}_{T_1}^T$ using the control dataset

To estimate  $\mathcal{R}_{T_1}^T$ , we assume that the neural network  $h_\theta$  is invariant to a potential domain shift between the  $T_1$ scans of the *control* and *lesion* datasets on the non-lesion regions. Specifically, we assume that the restriction of the feature distributions (rather than the image intensity distributions) over  $\mathcal{D}_{lesion}$  and  $\mathcal{D}_{control}$  to the non-lesion parts of the brain (i.e. the voxels  $i$  such that  $y_i \in \mathcal{C}_T$ ) are comparable, i.e.:

$$P_{\mathcal{D}_{lesion}}(h_\theta(x^{T_1})_i, y_i | y_i \in \mathcal{C}_T) = P_{\mathcal{D}_{control}}(h_\theta(x^{T_1})_i, y_i | y_i \in \mathcal{C}_T) \quad (\text{H}_4)$$

This means that the neural network  $h_\theta$  generates similar outputs on the non-lesion parts of the brain between the two datasets, leading to:

$$\mathcal{R}_{T_1}^T = \mathbb{E}_{\mathcal{D}_{lesion}}[\mathcal{L}^T(h_\theta(x^{T_1}), y^T)] = \mathbb{E}_{\mathcal{D}_{control}}[\mathcal{L}^T(h_\theta(x^{T_1}), y^T)] \quad (8)$$

Consequently, under (H<sub>4</sub>),  $\mathcal{R}_{T_1}^T$  can be estimated using the  $T_1$ scans and their tissue annotations in the *control* dataset. Section 5 presents means of ensuring that assumption (H<sub>4</sub>) is satisfied even in the presence of domain shift in the image intensity distributions.

## 4.4. Summary of the expected risk decomposition

Bringing all the terms together, given (3), (7) and (8), we seek the parameters  $\theta^*$  that optimise the tractable upper bound  $\mathcal{R}_{seg}$  of the (intractable) expected risk:

$$\theta^* = \operatorname{argmin}_\theta \left\{ \mathcal{R}_{seg} = \mathbb{E}_{\mathcal{D}_{control}}[\mathcal{L}^T(h_\theta(x^{T_1}), y^T)] + \mathbb{E}_{\mathcal{D}_{lesion}}[\mathcal{L}^L(h_\theta(x), y^L) + \mathcal{L}^T(h_\theta(x), h_\theta(x^{T_1}))] \right\} \quad (9)$$

## 5. Matching feature distributions across datasets

In this section, we explore different approaches that ensure the feature distributions extracted from the *control* and *lesion*  $T_1$ scans are comparable, i.e. we want to satisfy (H<sub>4</sub>) even in the presence of domain shift.

### 5.1. Similar acquisition protocols for the control and lesion datasets

Let's first assume that the acquisition protocols are similar for the *control* and *lesion* datasets, i.e. they are not domain-shifted. Specifically, we assume that the  $T_1$ images have been acquired with similar sequences, spacial resolution and field strength. In this case, similar to Chen and Konukoglu (2018), the restriction of the distributions  $\mathcal{D}_{lesion}$  and  $\mathcal{D}_{control}$  to the non-lesion parts of the brain can be assumed to be the same on the  $T_1$ scans, i.e.:

$$P_{\mathcal{D}_{lesion}}(x_i^{T_1}, y_i | y_i \in \mathcal{C}_T) = P_{\mathcal{D}_{control}}(x_i^{T_1}, y_i | y_i \in \mathcal{C}_T) \quad (10)$$

In the absence of domain shift, we can reasonably assume that the network produces similar outputs on the non-lesion parts of the brain for the two distributions, i.e. that (H<sub>4</sub>) is satisfied. No specific additional action thus needs to be implemented.

## 5.2. Generating pseudo-healthy scans to learn tissue segmentation from domain-shifted $T_1$ scans

Let's now consider the presence of a domain shift between the  $T_1$ control and lesion scans due to different acquisition protocols. In this section, we propose to synthesise pseudo-healthy scans from domain-shifted  $T_1$ lesion scans in order to extend the control dataset with control scans associated to the protocol of acquisition of the lesion dataset. Since the control and lesion datasets are domain-shifted, existing lesion removal approaches, based either on CycleGANs (Xia et al., 2019; Sun et al., 2018) or Variational Autoencoders (Chen and Konukoglu, 2018), are not adapted as they require training data with no domain shift beyond the presence of absence of pathology.

To tackle the presence of an acquisition-related domain shift, we propose to generate pseudo-healthy scans and their annotations using traditional image computing techniques that are inherently robust to different acquisition protocols. For example, for white matter lesions, lesion filling methods allow for transforming scans with lesions into pseudo-healthy scans (Valverde et al., 2014; Prados et al., 2016). For large and unilateral pathology, we propose to synthesise pseudo-healthy  $T_1$ scans by symmetrising the "healthy" hemisphere of the patients with lesions located in one hemisphere only. The inter-hemispheric symmetry plane is estimated via the technique described in Prima et al. (2002). Finally, the "healthy" hemisphere of those patients is mirrored in order to create a symmetric pseudo-healthy brain. Having generated pseudo-healthy images, traditional methods, designed for control scans, such as the GIF framework (Cardoso et al., 2015), can then be employed to generate the corresponding bronze standard tissue annotations.

With this set of scans  $S_{pseudo}^{T_1}$ , we have access to a pseudo-control dataset acquired with a similar protocol as in the lesion dataset and similar on the non-lesion part of the brain, and thus are in the scenario described in 5.1. Let denote  $\mathcal{D}_{pseudo}^{T_1}$  the distribution of those scans. The expected tissue risk  $\mathcal{R}_{T_1}^T$  is then equal to the expected tissue risk under  $\mathcal{D}_{pseudo}^{T_1}$ :

$$\mathcal{R}_{T_1}^T = \mathbb{E}_{\mathcal{D}_{lesion}} [\mathcal{L}^T(h_\theta(x^{T_1}), y^T)] = \mathbb{E}_{\mathcal{D}_{pseudo}^{T_1}} [\mathcal{L}^T(h_\theta(x^{T_1}), y^T)] \quad (11)$$

To take advantage of the manual annotations in the control dataset, we resort to averaging the two formulations (8,11):

$$\mathcal{R}_{T_1}^T \approx \frac{\mathbb{E}_{\mathcal{D}_{control}} [\mathcal{L}^T(h_\theta(x^{T_1}), y^T)] + \mathbb{E}_{\mathcal{D}_{pseudo}^{T_1}} [\mathcal{L}^T(h_\theta(x^{T_1}), y^T)]}{2} \quad (12)$$

Consequently, the expected tissue risk  $\mathcal{R}_{T_1}^T$  can be estimated using the control and pseudo-control  $T_1$ scans.

## 5.3. Alternative unsupervised DA techniques

In order to satisfy (H<sub>4</sub>), the feature representation of the non-lesion parts of the brain has to be invariant to the changes induced by the different protocols. A direct way to align the feature distributions restricted to the non-lesion parts would be to match the representations of pairs of scans acquired with the different settings. However, in our scenario, we do not have access to such pairs of domain-shifted scans.

In contrast, unsupervised DA allows to perform domain adaptation using unpaired and non-annotated domain-shifted scans. Unsupervised DA techniques commonly introduce an additional term ( $\mathcal{R}_{DA}$ ) that encourages the network to be invariant to the domain shift. Then, the total expected risk reads:

$$\mathcal{R}_{total} = \mathcal{R}_{seg} + \lambda \mathcal{R}_{DA} \quad (13)$$

where  $\lambda$  is a hyper-parameter that allows for balancing the segmentation risk  $\mathcal{R}_{seg}$  ((9)) with the DA regularisation  $\mathcal{R}_{DA}$ .

The definition of the DA term depends on the DA technique. In this work, two common unsupervised DA methods are considered, based either on data augmentation or adversarial learning.

### 5.3.1. Unsupervised DA via physically-inspired data augmentation

Since  $T_1$ scans play a key role for structure analysis, we expect high-resolution  $T_1$ scans for datasets developed in the scope of tissue segmentation, such as the control dataset. Conversely,  $T_1$ scans are often less critical for lesion segmentation and  $T_1$ scans may have been acquired with a lower resolution.

Let's assume that, less effort has been done to acquire high-resolution  $T_1$ lesion scans, explaining the differences in acquisition protocols. Specifically, we assume that the domain shift is caused by the presence of  $T_1$ lesion scans with artefacts (e.g. related to the MR bias field or the presence of motion artefacts) and a lower acquisition resolution. We additionally assume that differences of scanner characteristics (manufacturer, field strength) are excluded.

Then, physically-informed augmentation such as random bias field (Sudre et al., 2017) and motion artefacts (Shaw et al., 2019) and spacial smoothing can be employed to generate scans that are similar to the  $T_1$ lesion scans. Let denote  $T_\psi$  the composition of these transformations parametrised by the parameters  $\psi \sim \mathcal{D}_\psi$ . For any  $T_1$ control scan  $x_c^{T_1}$ , we can thus generate an augmented version  $T_\psi(x_c^{T_1})$ , i.e. getting access to pairs of domain-shifted  $T_1$ scans. This allows to minimise the discrepancy between the feature representations learnt by the neural network across the two domains by enforcing consistency across outputs from paired domain-shifted inputs, i.e.:

$$\mathcal{R}_{DA} = \mathbb{E}_{\mathcal{D}_\psi} \mathbb{E}_{\mathcal{D}_{control}} [\mathcal{L}^T(h_\theta(x_c^{T_1}), h_\theta \circ T_\psi(x_c^{T_1}))] \quad (14)$$

An empirical estimation of the DA regularisation term  $\mathcal{R}_{DA}$  is obtained by comparing the network outputs using the  $T_1$ control scans and their augmented versions as input.

Consequently, if the domain shift is due to different spatial resolutions and the presence of the aforementioned artefacts, the network can be trained to be invariant to the domain shift, i.e. to satisfy (H<sub>4</sub>).

### 5.3.2. Unsupervised DA via adversarial learning

Let's now assume that the domain shift cannot easily be simulated. In this case, we can use adversarial learning. Adversarial approaches for domain adaptation can be seen as a two-player game: A discriminator  $D_\phi$ , parametrised by the weights  $\phi \in \Phi$ , is trained to distinguish the source domain features from the target domain features, while the segmentation network  $h_\theta$  is simultaneously trained to confuse the domain discriminator.

The discriminator aims to predict the probability that extracted features are part of the lesion feature distributions. The discriminator accuracy can thus be seen as a measurement of the discrepancy between the lesion and control feature distributions and used as a DA regularisation term:


$$\mathcal{R}_{DA}(\phi, \theta) = \mathbb{E}_{\mathcal{D}_{lesion}} [1 - D_\phi(h_\theta(x))] + \mathbb{E}_{\mathcal{D}_{control}} [D_\phi(h_\theta(x))] \quad (15)$$

This DA term can be estimated by using features extracted from  $T_1$ control and lesion scans as input of the discriminator.

The following proposition shows that the discriminator accuracy is a principled measurement of the feature distribution discrepancy:

**Proposition 1.** *Let assume that  $\mathcal{L}$  satisfies the triangle inequality and is bounded. Let us also assume that the family of domain discriminators  $\mathcal{H}_\Phi = \{D_\phi\}_{\phi \in \Phi}$  is rich enough. Then there is a constant  $K$  such that:*

$$\mathbb{E}_{\mathcal{D}_{lesion}} [\mathcal{L}(h_\theta(x), y)] \leq \mathcal{R}_{seg} + K \sup_{\phi} \mathcal{R}_{DA}(\phi, \theta) + \epsilon(\Theta) \quad (16)$$



**Fig. 3.** The training procedure using samples from the *control* and *lesion* datasets. The different elements of the decomposed loss upper bound are computed and minimised at each training iteration. The same network is used for all the different hetero-modal inputs. Note that domain adaptation is not represented.

where  $\epsilon(\Theta)$  is independent of the network parameters  $\theta$  and corresponds to the accuracy of the best (and unknown) segmenter in the family of functions parametrised in  $\Theta$ .

**Proof.** The proof uses (3), (7), is based on Ben-David et al. (2010) and Long et al. (2018) and detailed in Appendix B.  $\square$

(16) shows that the intractable expected loss is bounded by a weighted sum of the tractable segmentation risk (9) and the accuracy of the best discriminator, up to a constant w.r.t the network parameters.

Moreover, the alternative optimisation strategy can be seen as a way to estimate the best discriminator while minimising the upper bound defined in (16).

Note that (16) stands for features extracted at any level of  $h_\theta$ . In this work and similarly to Orbes-Arteaga et al. (2019), the contracting path features from the U-Net are used as input of the discriminator.

## 6. Implementation of the joint model optimisation

Given the formulation of the joint model and our proposed computationally tractable decomposition, we present in this section the implementation of our framework.

### 6.1. Stochastic optimisation of the joint model

We use a stochastic gradient descent approach to minimise the expected risk decomposition (9) and to enforce the network to be invariant to a potential domain shift between the datasets. The total loss function reads:

$$\mathcal{L}_{total} = \mathcal{L}_{seg} + \lambda \mathcal{L}_{DA} \quad (17)$$

where  $\lambda$  is a hyper-parameter that allows for balancing the segmentation loss  $\mathcal{L}_{seg}$  (associated to  $\mathcal{R}_{seg}$ ) with the domain adaptation loss  $\mathcal{L}_{DA}$  (associated to  $\mathcal{R}_{DA}$ ). Fig. 3 shows the training procedure without DA. The weights of the segmentation loss are given by the decomposition of the problem. The domain adaptation parameter  $\lambda$  is a hyper-parameter that is experimentally chosen.

At each training iteration, we draw pairs of samples  $(x_l, y_l^l)$  and  $(x_c^T, y_c^T)$  from  $\mathcal{S}_{lesion}$  and  $\mathcal{S}_{control}$  and compute in each mini-batch the following loss functions and associated gradient. Note that there is no natural pairing between  $(x_l, y_l^l)$  and  $(x_c^T, y_c^T)$ . Our paired sampling procedure thus exploits random pairing.

As presented in Section 5, different scenarios are considered.

**Similar acquisition protocols** If the datasets are not domain-shifted, no DA is required ( $\lambda = 0$ ), and the segmentation loss is:

$$\begin{aligned} \mathcal{L}_{seg} = & \mathcal{L}^L(h_\theta(x_l), y_l^l) + \mathcal{L}^T(h_\theta(x_l), h_\theta(x_c^T)) \\ & + \mathcal{L}^T(h_\theta(x_c^T), y_c^T) \end{aligned} \quad (18)$$

We experimentally found that the inter-modality tissue loss  $\mathcal{R}_{T_1 \rightarrow Full}^T$  has to be skipped for few epochs (50 in our experiments).

**Pseudo-healthy generation** Given a pseudo-healthy annotated set of scans  $\mathcal{S}_{pseudo}^{T_1} = \{x_{pseudo}^{T_1}, y_{pseudo}^T\}$ , no DA is employed ( $\lambda = 0$ ), and the segmentation loss, defined by (12), is:

$$\begin{aligned} \mathcal{L}_{seg} = & \mathcal{L}^L(h_\theta(x_l), y_l^l) + \mathcal{L}^T(h_\theta(x_l), h_\theta(x_c^T)) \\ & + \frac{1}{2} [\mathcal{L}^T(h_\theta(x_{pseudo}^{T_1}), y_{pseudo}^T) + \mathcal{L}^T(h_\theta(x_c^T), y_c^T)] \end{aligned} \quad (19)$$

**DA via augmentation** If we assume that the differences of protocols can be simulated (random bias field, motion artefacts and spatial smoothing), the domain invariance (14) is learnt by minimising the inter-domain feature discrepancy defined as:

$$\mathcal{L}_{DA} = \mathcal{L}^T(h_\theta(x_c^T), h_\theta(T_\phi(x_c^T)))$$

where  $T_\phi$  corresponds to a composition of these transformations. The segmentation loss  $\mathcal{L}_{seg}$  is the same as in (18).

**DA via adversarial learning** If adversarial learning is employed, a discriminator  $D$  is trained to discriminate scans from the two domains by maximising the domain classification accuracy. For computational stability, the  $L_1$  distance defined in (15) has been replaced by the cross-entropy. Conversely, the segmenter  $h_\theta$  is trained to minimise this domain classification accuracy, i.e.:

$$\mathcal{L}_{DA} = \log(D_\psi(h_\theta(x_c^T))) + \log(1 - D_\psi(h_\theta(x_l^T)))$$

As in Kamnitsas et al. (2017); Orbes-Arteaga et al. (2019), the DA loss is skipped for few epochs (20 in our experiments) in order to initialise the discriminator. The segmentation loss  $\mathcal{L}_{seg}$  is the same as in (18).

### 6.2. Implementation details

We implemented our network in PyTorch, using TorchIO (Pérez-García et al., 2020). Codes are available at <http://github.com/ReubenDo/jSTABL>.

Convolutional layers are initialised such as proposed in He et al. (2015). The scaling and shifting parameters in the batch normalisation layers were initialised to 1 and 0 respectively. As suggested by Ulyanov et al. (2016), we used instance normalisation. We used the same discriminator as in Orbes-Arteaga et al. (2019).

**Table 1**  
Summary of data characteristics for white matter lesion segmentation.

	Control data		Lesion data			Fully anno. data
	OASIS1 (Marcus et al., 2007)	ADNI2 (Jack Jr. et al., 2008)	WMH-Utrecht (Kuijf et al., 2019)	WMH-Singapore	WMH-Amsterdam	MRBrainS18 (Kuijf and Bennink, 2018)
Sequences	3D MP-RAGE T <sub>1</sub> ×	3D MP-RAGE T <sub>1</sub> ×	3D MP-RAGE T <sub>1</sub> 2D FLAIR	3D MP-RAGE T <sub>1</sub> 2D FLAIR	3D MP-RAGE T <sub>1</sub> 3D FLAIR	MP-RAGE 3D T <sub>1</sub> 2D FLAIR
MRI scanner	Siemens Vision	Various	Philips Achieva	Siemens TrioTim	GE Signa HDxt	Philips Achieva
Field Strength	1.5T	3T	3T	3T	3T	3T
Voxel size (mm <sup>3</sup> )	1.00 × 1.00 × 1.00	1.20 × 1.05 × 1.05	0.96 × 0.95 × 3.00	1.00 × 1.00 × 3.00	1.20 × 0.98 × 3.00	0.96 × 0.96 × 3.00
Annotations	143 structures	143 structures	WMH	WMH	WMH	6 Tissues+WMH+CSF
# scans available	35	25	20	20	20	30 (7 available)
Training data in:	Tissue + jSTABL	Tissue + jSTABL	Lesion + jSTABL	Lesion + jSTABL	Lesion + jSTABL	Fully

We performed a 3-fold cross validation. For each fold, we randomly split the data into 70% for training, 10% for validation and 20% for testing. We used a batch of 2 *lesion* scans, and 2 *control* scans. Note that, for the DA approach based on data augmentation, the batch of 2 *control* scans consists in a pair of non-augmented/augmented *control* scans. As a data augmentation, a rotation with a random angle in  $[-10^\circ, 10^\circ]$  and a random Gaussian noise are employed. The network was trained using Adam optimiser (Kingma and Ba, 2015) the learning rates  $l_R$ ,  $\beta_1$ ,  $\beta_2$  were initially respectively set up to  $5.10^{-4}$ , 0.9 and 0.999.  $l_R$  was progressively reduced by a factor of 2 every 10,000 iterations. We employed the training strategy used for the nnU-Net (Isensee et al., 2019): The learning rate is reduced by a factor 2 after 15 epochs without reduction of the exponential moving average of the loss on the validation split.

We used the probabilistic version of the multi-class Jaccard distance (5) as the segmentation loss function. In order to give the same weight to the lesion segmentation and the tissue segmentation, we choose  $\omega$  such that

$$\sum_{c \in C_{\text{tissue}}} \omega_c = \sum_{c' \in C_{\text{lesion}}} \omega_{c'} = \frac{1}{2}.$$

## 7. Experiments and results

### 7.1. Joint white matter lesion and tissue segmentation

#### 7.1.1. Task and datasets

In this first set of experiments, we focus on the segmentation of white matter lesions and six tissue classes (white matter, grey matter, basal ganglia, ventricles, cerebellum, brainstem), as well as the background. As detailed in Table 1, we used 2 *control* datasets and 2 *lesion* datasets:

- **Lesion data  $S_{\text{lesion}}$ :** The White Matter Hyperintensities (WMH) training database (Kuijf et al., 2019) consists of 60 sets of brain MR images (T<sub>1</sub> and FLAIR,  $M = 2$ ) with manual annotations of WMHs. The data comes from three different institutes. Note that images modalities are be co-registered and resampled in the FLAIR coordinate space.
- **Tissue data  $S_{\text{control}}$ :** Consists of 35 T<sub>1</sub>scans ( $M' = 1$ ) from the OASIS project (Marcus et al., 2007) with annotations of 143 structures of the brain provided by Neuromorphometrics, Inc. (<http://Neuromorphometrics.com/>) under academic subscription. From the 143 structures, we deduct the 6 tissue classes. In order to have balanced training datasets between the two datasets, to include data acquired at the same field strength (3T) as the *lesion* data, and similar to Li et al. (2017), we added 25 T<sub>1</sub>control scans from the Alzheimer's Disease Neuroimaging Initiative 2 (ADNI-2) database (Jack Jr. et al. (2008), [adni.loni.usc.edu](http://adni.loni.usc.edu)) with bronze standard parcellation of the brain structures computed with the accurate but time-consuming algorithm of Cardoso et al. (2015).

- **Fully annotated data  $S_{\text{fully}}$ :** MRBrainS18 (<http://mrbrains18.isi.uu.nl/>) is composed of 30 sets of brain MR images with tissue and lesion manual annotations. 7 scans are publicly available for training and validation. Although the cerebrospinal fluid (CSF) has been annotated in MRBrainS18, it was considered as background to have the same set of tissue classes as in  $S_{\text{control}}$  where the CSF was not labelled. Note that image modalities are be co-registered and resampled in the FLAIR coordinate space.

#### 7.1.2. Similar acquisition protocol for the T<sub>1</sub>scans

Despite the differences in scanners, all T<sub>1</sub>acquisitions across the datasets followed very similar protocols (MP-RAGE) (see Table 1). Therefore, they were considered as following a similar distribution and the data was only pre-processed as follows.

- **Skull stripping:** All the scans were skull-stripped using ROBEX (Iglesias et al., 2011).
- **Resampling:** All the scans in  $S_{\text{control}}$  are resampled into the transversal direction with slices of 3 mm thickness to obtain a similar spacing  $1 \times 1 \times 3 \text{ mm}^3$  in the datasets.
- **Intensity normalisation:** We used a zero-mean unit-variance normalisation in order to match the intensity distributions.

#### 7.1.3. Description of the compared models

We considered three different models in our experiments.

- **Pipeline model (Pipeline):** This model corresponds to the combination of two task-specific models:  
A *Tissue* segmentation model that only performs tissue segmentation and is trained on the T<sub>1</sub>scans from the dataset with tissue annotations  $S_{\text{control}}$ .  
A *Lesion* model that only performs lesion segmentation and is trained using the T<sub>1</sub> and FLAIR scans from the dataset with lesion annotations  $S_{\text{lesion}}$ .  
The two models are combined such that the predicted lesion mask has the priority over the predicted tissue mask. Consequently, the background of the *Lesion* output is replaced by the *Tissue* output.
- **Fully-supervised model (Fully-Sup):** This joint model performs tissue and lesion segmentation and is trained using the T<sub>1</sub> and FLAIR scans from the small fully-annotated dataset  $S_{\text{fully}}$ .
- **Proposed joint model (jSTABL):** Our proposed model for joint Segmentation of Tissues and Brain Lesions is trained using both the T<sub>1</sub>scans from  $S_{\text{control}}$  with the tissue annotations and the T<sub>1</sub> and FLAIR scans from  $S_{\text{lesion}}$  with the lesion annotations.

Each model used the architecture presented in Fig. 2. Consequently, the *Pipeline* model has twice as many parameters as the other models.

In this set of experiments, the skull-stripped images are first cropped to remove the blank spaces and then padded to size of (144,192,48).



**Table 2**

Evaluation of our framework (jSTABL) on patients with White Matter Lesion in comparison with baseline methods. We report means and standard deviations for Dice scores. Means only are reported in the online leader-board, leading to missing standard deviations.

Classes	OASIS1+ADNI2			WMH			MRBrainS18			
	Tissue	Fully-Sup	jSTABL	Pipeline	Fully-Sup	jSTABL	SPM	Pipeline	Fully-Sup	jSTABL
Grey matter	88.3 (3.4)	81.6 (2.5)	88.3 (3.4)	88.3 (2.1)	85.4 (2.7)	<b>88.8 (2.1)</b>	76.5	82.3	83.7	82.2
White mater	92.8 (2.3)	83.3 (2.6)	92.3 (2.7)	92.1 (1.8)	85.4 (2.6)	<b>92.4 (1.5)</b>	75.7	85.0	85.7	85.6
Brainstem	93.5 (1.0)	71.7 (2.7)	93.0 (0.9)	93.6 (1.0)	77.1 (2.4)	<b>94.2 (0.9)</b>	76.5	72.8	85.0	73.3
Basal ganglia	89.5 (3.0)	69.6 (4.1)	88.4 (2.7)	<b>86.3 (4.3)</b>	74.2 (2.0)	85.1 (3.4)	74.7	77.4	79.7	78.0
Ventricles	90.3 (4.3)	70.5 (18.0)	90.6 (3.8)	94.7 (2.3)	92.1 (4.2)	<b>95.7 (1.4)</b>	80.9	91.8	92.2	92.9
Cerebellum	95.0 (1.2)	92.0 (1.4)	94.9 (1.1)	95.7 (1.0)	93.8 (2.0)	<b>96.0 (0.9)</b>	89.4	89.2	93.2	90.4
White matter Lesion				77.4 (9.6)	60.1 (19.1)	<b>77.6 (9.2)</b>	40.8	58.4	56.2	59.4

**Table 3**

Evaluation of our framework (jSTABL) on patients with White Matter Lesion in comparison with baseline methods. We report means and standard deviations for 95th-percentile Hausdorff distances. Means only are reported in the online leader-board, leading to missing standard deviations.

Classes	OASIS1+ADNI2			WMH			MRBrainS18			
	Tissue	Fully-Sup	jSTABL	Pipeline	Fully-Sup	jSTABL	SPM	Pipeline	Fully-Sup	jSTABL
Grey matter	1.3 (0.4)	2.0 (0.4)	1.4 (0.5)	1.2 (0.2)	1.3 (0.3)	<b>1.1 (0.2)</b>	2.9	1.9	1.9	2.1
White mater	1.2 (0.5)	3.0 (0.1)	1.2 (0.5)	<b>1.1 (0.2)</b>	2.0 (0.3)	<b>1.1 (0.1)</b>	4.9	3.2	2.9	3.2
Brainstem	1.4 (0.4)	10.5 (2.3)	1.7 (0.4)	1.7 (0.4)	8.8 (2.2)	<b>1.3 (0.4)</b>	25.3	11.6	6.65	11.6
Basal ganglia	1.8 (0.5)	4.9 (0.8)	2.1 (0.3)	<b>2.0 (0.6)</b>	3.6 (0.4)	2.5 (0.3)	7.1	4.3	4.3	4.0
Ventricles	1.9 (2.5)	20.4 (12.2)	1.8 (2.5)	1.4 (1.2)	5.1 (10.5)	<b>1.0 (0.1)</b>	5.8	3.2	3.0	2.9
Cerebellum	2.3 (0.6)	3.3 (0.4)	2.4 (0.6)	1.8 (0.7)	3.2 (1.7)	<b>1.8 (0.6)</b>	4.3	5.1	3.7	4.8
White matter Lesion				4.6 (3.8)	11.0 (9.4)	<b>4.2 (3.6)</b>	25.3	10.2	13.3	7.2

#### 7.1.4. Method for assessing the models

The performance of the three models was evaluated on the three datasets using Dice Score and 95% Hausdorff distance. On the *control* data (OASIS1+ADNI2) and WMH, scores were computed on the testing splits, while on MRBrainS18, models were submitted to the challenge MRBrainS18.

For the *control* data (OASIS1+ADNI2) and MRBrainS18, the full set of annotations allows a direct assessment of the tissue and the lesion segmentation performance. For WMH, only the *lesion* annotations are provided. In order to assess both the tissue and lesion segmentation on WMH, the lesions are filled as normal-appearing white matter on  $T_1$  images using the method described in Prados et al. (2016) and implemented in NiftySeg (Cardoso et al., 2015). Then, GIF framework (Cardoso et al., 2015) was performed on the modified  $T_1$  scans to obtain bronze standard tissue annotations. The tissue mask and lesion annotations were then merged by completing the non-lesion parts with the tissue mask. Finally, the model outputs are compared to the merged tissue and lesion masks. In the end, for each model and each dataset, we can assess the performance of tissue and lesion segmentation.

Given that participants to the MRBrainS18 challenge do not have access to the held-out evaluation data set and that the Jaccard score is not provided by the challenge organisers, only the Dice Similarity Coefficient (DSC) and 95th-percentile Hausdorff distance are reported for each class.

#### 7.1.5. Results

The main results are shown in Table 2 for the Dice Similarity Coefficient and in Table 3 for the 95th-percentile Hausdorff distance.

Firstly, our proposed method (jSTABL) achieves comparable performance to the single-task models on the *control* data (*Tissue*) and on WMH (priority of *Lesion* in *Pipeline*). This suggests that learning from hetero-modal datasets via our method does not degrade the performance on the tasks characterising the task-specific datasets.

Secondly, jSTABL slightly outperforms *Pipeline* on segmenting the tissues in WMH for the two sets of metrics. This shows that the tissue knowledge learnt from  $T_1$  scans has been well generalised to multi-modal scans. Although we could have expected that the presence of lesions would create perturbations for the *Tissue* model, this latter model in fact ignores the lesions and mostly

classifies them as white matter. Given that the white matter lesions are usually surrounded by white matter, the *Pipeline* predictions are consequently not too degraded. However, some artefacts around the lesions in the *Pipeline* outputs can be observed, in particular in the ventricles for patients with large lesions surrounding them. Fig. 4(a) shows an example for which parts of the ventricles are classified as background. In contrast, we did not observe such artefacts with jSTABL predictions.

Thirdly, jSTABL outperforms the fully-supervised model (*Fully-Sup*) on the *control* data (OASIS1+ADNI2) and WMH, while reaching comparable performance on MRBrainS18. This demonstrates the two main advantages of our method. First, without using any fully-annotated data, our model performs as well as a fully-supervised model that could be considered as an upper bound for our method, especially when the testing and training splits are from the same dataset (MRBrainS18). Secondly, our method takes advantage of large task-specific datasets: Unlike the fully-supervised model (*Fully-Sup*), jSTABL generalises well on unseen data (MRBrainS18). While the fully-supervised model (*Fully-Sup*) fails to segment scans from OASIS1, ADNI2 and WMH, the jSTABL model obtains relatively good performance on all the datasets we use for tissue and lesion segmentation. In particular, jSTABL outperforms SPM on 6 of the 7 classes. In fact, the only class that is significantly underperformed compared to the fully-supervised model (*Fully-Sup*) is the brain stem. This is due to observed differences in the annotation protocol across the *control* and MRBrainS datasets. Fig. 5 shows these differences and the consequences on the prediction.

## 7.2. Glioma and tissue segmentation

### 7.2.1. Task and datasets

Additionally, we assess our framework on another main types of brain lesions: Gliomas. Our goal is to segment the 6 tissue classes and three tumour classes (whole tumour, core tumour, enhancing tumour). In this case, domain adaptation was required and its evaluation is the focus of this section. We used two sets of data in these experiments.




Fig. 4. Examples of output for the different models: Pipeline; Fully-sup; jSTABL; and the combination of the manual annotation and GIF. (a) Tissue model used in Pipeline can be perturbed by the presence of lesions (arrow). (b) Example for which the fully-supervised model largely fails to segment the tissue and the lesions.

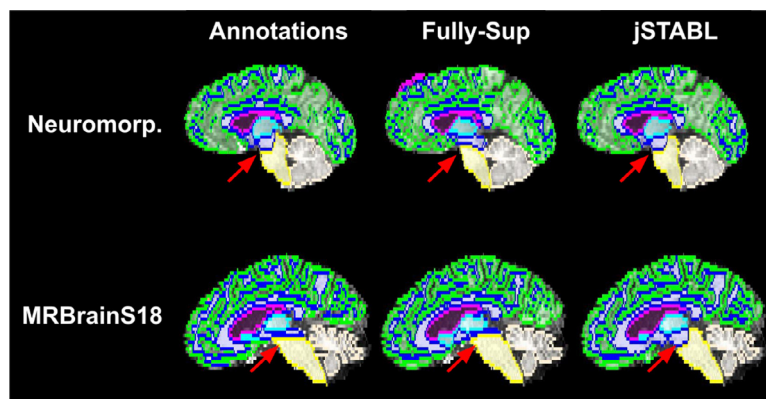


Fig. 5. Comparison of the brainstem annotations by Neuromorphometrics and in MRBrainS18 and between the outputs of the Fully-Sup and jSTABL models. Arrows show the annotations protocol differences.

- **Tissue data  $S_{control}$** : again we used OASIS1 data and the same 25  $T_1$  control scans from ADNI2 with tissue annotations as presented in section 7.1.1.
- **Lesion data  $S_{lesion}$** : We evaluate our method on the training set of BraTS18 (Menze et al., 2015; Bakas et al., 0000) which contains the scans of 285 patients, 210 with high grade glioma and 75 with low grade glioma. 129 patients have a tumour located in one hemisphere only. Four scans ( $T_1$ ,  $T_1c$ ,  $T_2$  and FLAIR) have been acquired for each patient and pre-processed by the organisers: Co-registration, skull-stripping and re-sampling to an isotropic 1mm resolution. Manual annotations include three tumour labels: 1) Necrotic core and non-enhancing tumour; 2) oedema; and 3) enhancing core.

The acquisition protocols of the  $T_1$  scans in the two datasets are inconsistent. Specifically, MP-RAGE was used for the tissue data  $S_{control}$ , while we observed other protocols such as fast spin echo (SE) for  $S_{lesion}$ . Note that the detailed acquisition settings for  $S_{lesion}$  are not publicly available.

### 7.2.2. Description of the compared models

In order to evaluate our framework with and without the domain adaptation (DA) component, different models are considered, as presented in 5.

- **Pipeline model (Pipeline)**: This model corresponds to the combination of two task-specific models:

A Tissue segmentation model that only performs tissue segmentation and is trained on the  $T_1$  scans from the dataset with tissue annotations  $S_{control}$ .


A Lesion model that only performs lesion segmentation and is trained using the  $T_1$ ,  $T_1c$ ,  $T_2$  and FLAIR scans from the dataset with lesion annotations  $S_{lesion}$ .

- **Proposed joint model without DA (jSTABL)**: Our joint Segmentation Tissue And Brain Lesion model is trained using our training procedure without domain adaptation, tissue segmentation is learned from the  $T_1$  scans in  $S_{control}$ .
- **jSTABL + data augmentation (jSTABL+Augm)**: Corresponds to our jSTABL model with DA based on data augmentation.
- **jSTABL + adversarial DA (jSTABL+Adv)**: Corresponds to our jSTABL model with DA based on adversarial learning.
- **jSTABL + 5 synthetic control scans (jSTABL+5)**: Corresponds to our jSTABL model with only 5 additional pseudo-healthy scans.
- **jSTABL + 90 synthetic control scans (jSTABL+90)**: Corresponds to our jSTABL model with 90 additional pseudo-healthy scans.

Note that the pseudo-healthy scans used for training were generated from the training lesion scans to avoid introducing bias at testing stage.

In this set of experiments, the skull-stripped images are first cropped to remove the blank spaces and then random patches of size (112,112,112) are fed to the network.





**Fig. 6.** Examples of multi-modal outputs shown on T<sub>1</sub> scans from BraTS18 for the different models: *Pipeline*, *jSTABL*, *jSTABL+Augm*, *jSTABL+Adv*, *jSTABL+5*, *jSTABL+90* and *Les. An. + GIF* pipeline. Scans with different resolutions, contrasts and grades (High Grade for (a) and (b), Low Grade for (c)) are presented.

### 7.2.3. Method for assessing the models

While the evaluation of the tumour segmentation on BraTS18 and the tissue segmentation on the *control* data (OASIS1+ADNI2) is straightforward using the manual annotations, the tissue segmentation performance cannot be assessed on BraTS18 due to the missing tissue annotations. For this reason we propose two methods to assess quantitatively and qualitatively the tissue segmentation on BraTS18.

**Quantitative assessment using the symmetrised data** Firstly, we propose to use the 129 patients from BraTS18 with a tumour located in one side to generate 129 pseudo-healthy symmetrised data with the bronze standard tissue annotations from GIF as ground truth. Examples are shown in Fig. C.9. By computing the Dice Score Coefficient between the predictions on the symmetrised BraTS18 data and the bronze standard ground truth, we quantitatively evaluate our model on the pseudo-healthy hemisphere of BraTS18 samples.

**Qualitative assessment on anatomical landmarks** Secondly, in order to assess the accuracy of the models on the tissues surrounding the tumour, we propose a new qualitative protocol. This protocol is based on the Alberta Stroke Program Early CT Score (ASPECTS) which was originally proposed to assess early ischaemic cerebral changes on CT or MRI scans (Barber et al., 2000). The stroke scores are obtained by assessing the integrity of 10 anatomical landmarks as shown on 8. Scores and associated template are commonly used in clinical practice.

The landmarks were chosen because they are easily identifiable, reliable amongst readers and capture a large cerebral coverage. The landmarks include or delineate our tissue classes of interest: Grey matter; white matter; basal ganglia; and ventricles. Instead of evaluating loss of clarity of landmarks due to ischemia we evaluated loss of clarity of landmarks due to incorrect tissue predictions. Unlike ASPECTS, which excludes infratentorial structures which are difficult to evaluate on CT, we added the brainstem and the cerebellum as two additional landmarks for a total of 12 anatomical landmarks. We named our assessment method Anatomy ASPECTS+. For each landmark, 3 scores are possible: 0 = anatomy inaccurate; 0.5 = anatomy mostly accurate; and 1 = anatomy highly accurate. Anatomical landmarks that were infiltrated with substantial tumour were excluded.

For our experiments, we randomly drew 20 patients from the testing sets of BraTS18 and two senior neuro-radiologists independently evaluated the quality of the predictions using Anatomy ASPECTS+ for 4 methods: 1) *Les. An. + GIF* pipeline that uses the tumour annotations and tissue segmentation obtained by GIF while the tumour is masked; 2) *Pipeline*; 3) *jSTABL*; 4) *jSTABL+10*. The neuro-radiologists assessed blindly the 4 methods, in a randomised order, for the 20 patients.


### 7.2.4. Results

Table 4 shows the DSC for the 6 tissue classes and the three tumour classes on BraTS18.

**Table 4**

Evaluation of our framework (jSTABL) on patients with gliomas in comparison to baseline methods. We report means and standard deviations for Dice scores. Metrics were computed on the BraTS 2018 validation dataset.

Models	w/o DA		w/ DA		w/ pseudo-healthy gen.	
	Pipeline	jSTABL	jSTABL + Adv	jSTABL + Augm	jSTABL + 5	jSTABL + 90
Grey Matter	76.1 (8.4)	79.1 (5.2)	81.1 (4.6)	82.8 (4.7)	88.3 (3.9)	88.8 (4.0)
White Matter	85.4 (5.7)	87.0 (4.4)	88.1 (4.4)	90.3 (2.8)	93.1 (2.5)	93.3 (2.6)
Brainstem	81.5 (17.6)	92.4 (2.5)	92.6 (1.9)	92.4 (2.0)	94.9 (1.4)	95.5 (1.5)
Basal Ganglia	72.7 (20.1)	73.1 (7.3)	77.7 (7.1)	84.7 (5.1)	89.7 (3.5)	90.5 (3.2)
Ventricles	75.0 (26.9)	91.8 (6.4)	92.5 (4.5)	93.4 (4.4)	94.7 (4.2)	95.1 (3.8)
Cerebellum	86.5 (11.3)	93.2 (4.6)	93.7 (3.5)	94.0 (2.4)	94.7 (4.5)	95.1 (2.9)
Whole Tumour	87.9 (8.7)	88.1 (6.7)	87.7 (8.1)	88.3 (9.0)	88.2 (8.1)	88.1 (9.4)
Core Tumour	78.6 (20.5)	79.1 (19.6)	79.5 (18.9)	80.4 (18.6)	80.5 (18.1)	80.9 (17.9)
Enhancing Tumour	69.9 (29.1)	70.0 (28.6)	70.0 (28.9)	71.4 (27.7)	70.3 (28.8)	71.0 (28.3)



**Fig. 7. Comparison of our method (jSTABL) with the GIF framework using the proposed Anatomy ASPECTS+ qualitative assessment methodology.** BR = Brainstem, C = Caudate, CE = Cerebellum, I = Insula, IC = Internal Capsule, L = Lentiform Nucleus, M1 = Frontal operculum, M2 = Anterior temporal lobe, M3 = Posterior temporal lobe, M4 = Anterior MCA, M5 = Lateral MCA, M6 = Posterior MCA.

Firstly, *jSTABL* model outperforms *Pipeline* on tissue segmentation. We observed that the presence of a large tumour creates major perturbations for the *Tissue* model. For example, we found samples for which the tumour and the surrounding tissues were partially classified as cerebellum, even though the tumour was far from the cerebellum, as shown in Fig. 6. In contrast, such artefacts were not observed for *jSTABL* model, demonstrating again advantages of our method compared to a simpler *Pipeline* approach.

Secondly, while obtaining relatively good performance on most of the tissue classes, *jSTABL* model fails to segment correctly grey matter and basal ganglia. This highlights the needs for domain adaptation.

Thirdly, learning from pseudo-healthy annotated scans (*jSTABL+5* and *jSTABL+90*) outperforms the other unsupervised DA strategies based either on data augmentation (*jSTABL+Augm*) and adversarial learning (*jSTABL+Adv*). This demonstrates the benefits of using a supervised approach for our problem. Moreover, only 5 pseudo-healthy annotated scans are required to obtain an accuracy similar to the one on the *control* data (see Table 2), i.e. to bridge the domain gap. Fig. 6 shows that learning from pseudo-healthy annotated scans allows the network to be robust to variations in resolution, contrast or glioma grade, even with few samples used for domain adaptation.

Finally, Anatomy ASPECTS+ is employed to provide a quantitative assessment of the segmentation of the tissues surrounding the tumour. Four models are compared by two neuro-radiologists: *Pipeline*, *jSTABL*, *jSTABL+5* and the time-consuming *Les. Ann. + GIF* pipeline that requires manual annotations of the lesions. Results are presented in Fig. 7. First, *jSTABL* is more often “mostly accurate” than the *Pipeline* (mean score - 75% vs 66%). Again, this highlights the strength of our joint model compared to a pipeline ap-

proach. Secondly, *jSTABL+5* is more often “highly accurate” than the *Les. Ann. + GIF* pipeline (mean score - 46% vs 24%). This shows that our fast and fully automatic method can be considered as a new state-of-the-art for performing joint tissue and lesion segmentation.

## 8. Discussion

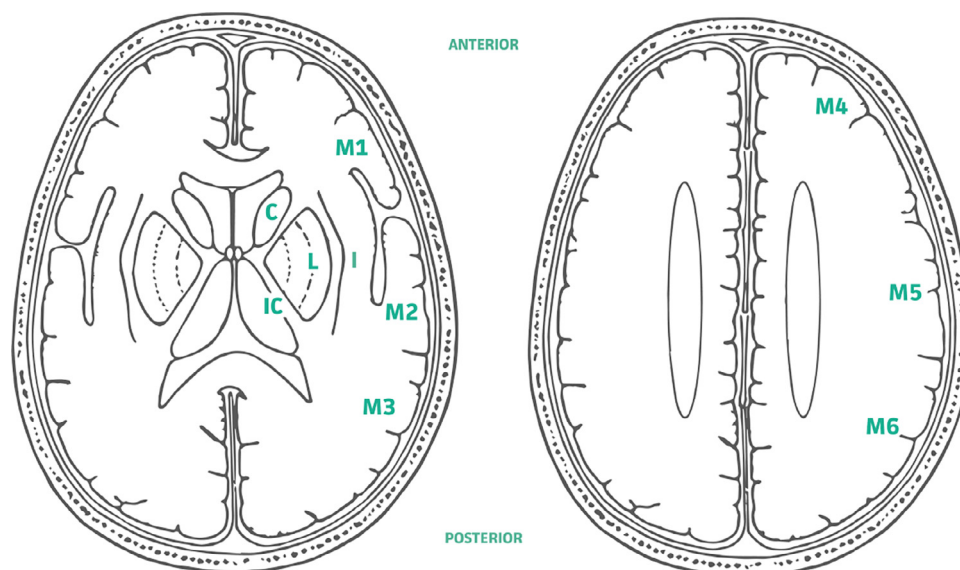
In this section, we discuss some of the limitations of the different methods.

Firstly, a common modality across the task-specific dataset is required to transfer the knowledge learned between the task-specific sets of modality. Without this common modality, the upper bound is not tractable anymore and our method cannot be applied.

Secondly, our approach relies on a simple hetero-modal architecture that aims to encode modalities in a common shared feature space. Yet, averaging the feature maps doesn’t enforce the network to learn a shared feature representation. To tackle this problem, a hetero-modal variational auto-encoder architecture has been recently introduced Dorent et al. (2019a). Based on a principled formulation of the problem, the induced loss function is the cross-entropy, which does not satisfy the triangle inequality. Consequently, without further research, this approach cannot be directly integrated in the formulation of our problem.

Thirdly, we found that the presence of lesions does not always perturb a network trained on *control* data, especially for small lesions. Consequently, our method didn’t always show large improvements compared to a simpler *Pipeline* approach on WMH. However, we observed that *Pipeline* was perturbed by larger pathology and thus is less robust.





**Fig. 8.** ASPECTS anatomical landmarks (Barber et al., 2000) used in our qualitative assessment methodology in the absence of joint ground truth. C = Caudate, I = Insula, IC = Internal Capsule, L = Lentiform Nucleus, M1 = Frontal operculum, M2 = Anterior temporal lobe, M3 = Posterior temporal lobe, M4 = Anterior MCA, M5 = Lateral MCA, M6 = Posterior MCA Illustration courtesy of P.A. Barber.

## 9. Conclusion

This work addresses the challenge of learning a joint brain tissue and lesion segmentation with disjoint heterogeneous annotations. Our novel approach is mathematically grounded, conceptually simple, and relies on reasonable assumptions.

The main contribution of this work is to overcome the challenge of the lack of fully-annotated data for joint problems. We demonstrate that a model trained on databases providing either the tissue or the lesion annotations and with different modalities can achieve similar performance to a model trained on a fully-annotated joint dataset. Our work also shows that the knowledge learnt from one modality can be preserved when more modalities are used as input. Finally, domain adaptation for image segmentation can be performed with a small set of data related to the target distribution.

In the future, we will evaluate our approach on new datasets with other lesions. Furthermore, we would like to extend our method to include the full parcellation of the brain (143 structures). Finally, we plan to integrate uncertainty measures in our framework as a future work. As one of the first work to methodologically address the problem of joint learning from hetero-modal and domain-shifted datasets, we believe that our approach will help DNN make further impact in clinical scenarios.

## Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This work was supported by the Wellcome Trust [203148/Z/16/Z] and the Engineering and Physical Sciences Research Council (EPSRC) [NS/A000049/1]. TV is supported by a Medtronic / Royal Academy of Engineering Research Chair [RC-SRF1819\734]. CS is supported by an the Alzheimer's Society Junior Fellowship [AS-JF-17-011]. We thank Prof. Alexander Hammers, Samuel Joutard and Lucas Fidon for their useful comments. Some

of the data used in preparation of this article was obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database ([adni.loni.usc.edu](http://adni.loni.usc.edu)). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at [http://adni.loni.usc.edu/wp-content/uploads/how\\_to\\_apply/ADNI\\_Acknowledgement\\_List.pdf](http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf).

## Appendix A. Probabilistic multi-class Jaccard loss function

First we recall the definitions of the binary Jaccard distance and the proposed extension to probabilistic inputs. The binary Jaccard distance  $J_{bin}$  is defined such that:

$$\forall a, b \in \{0, 1\}^N, J_{bin}(a, b) = 1 - \frac{\sum_{i=1}^N a_i b_i}{\sum_{i=1}^N a_i + b_i - a_i b_i} \quad (\text{A.1})$$

### Definition A.1. (Probabilistic multi-class Jaccard distance)

Let  $C$  be the number of classes in  $\mathcal{C}$ ,  $N$  be the number of voxels and  $\mathcal{P} \subset [0, 1]^{C \times N}$  denote the set of probability vector map such that for any  $p = (p_{c,i})_{c \in \mathcal{C}, i \in [0, N]} \in \mathcal{P}$ :

$$\forall i \in [0; N], \sum_{c \in \mathcal{C}} p_{c,i} = 1$$

The probabilistic multi-class Jaccard distance is defined for any  $(u, v) \in \mathcal{P}^2$  as:

$$\mathcal{J}(u, v) = \sum_{c \in \mathcal{C}} \omega_c \frac{2 \sum_{i=1}^N |u_{c,i} - v_{c,i}|}{\underbrace{\sum_{i=1}^N |u_{c,i}| + |v_{c,i}| + |u_{c,i} - v_{c,i}|}_{\mathcal{J}_c}} \quad (\text{A.2})$$

where  $\omega_c$  are class-specific weights summing up to one.

### A1. Relation between the probabilistic jaccard loss and the binary jaccard distance

The binary case corresponds to a two-class problem, i.e.  $C = \{0, 1\}$ . Let  $a, b \in \{0, 1\}^N$  be two binary vectors of size  $n$  and let  $u$  and  $v$  denote respectively the categorical encodings of

$a$  and  $b$ , i.e.  $a_i = 1 \iff (u_{1,i} = 1 \text{ and } u_{0,i} = 0)$  and  $b_i = 1 \iff (v_{1,i} = 1 \text{ and } v_{0,i} = 0)$ . The binary Jaccard distance can be rewritten as:

$$\mathcal{J}_{bin}(a, b) = 1 - \frac{\sum_{i=1}^N a_i b_i}{\sum_{i=1}^N a_i + b_i - a_i b_i} = \frac{\sum_{i=1}^N a_i + b_i - 2a_i b_i}{\sum_{i=1}^N a_i + b_i - a_i b_i}$$

Given that for all  $i \in \{1, \dots, N\}$ ,  $a_i = 0$  or  $a_i = 1$ , we observe that  $a_i^2 = a_i$ . Using the same property for  $b$ , we get:

$$\mathcal{J}_{bin}(a, b) = \frac{\sum_{i=1}^N (a_i - b_i)^2}{\sum_{i=1}^N a_i^2 + b_i^2 - a_i b_i} = \frac{2 \sum_{i=1}^N (a_i - b_i)^2}{\sum_{i=1}^N a_i^2 + b_i^2 + (a_i - b_i)^2}$$

Finally, given that for all  $i \in \{1, \dots, N\}$ ,  $(a_i - b_i)^2 = 0$  or  $(a_i - b_i)^2 = 1$ , we have that  $(a_i - b_i)^2 = |a_i - b_i|$  and conclude that:

$$\mathcal{J}_{bin}(a, b) = \frac{2 \sum_{i=1}^N |a_i - b_i|}{\sum_{i=1}^N a_i^2 + b_i^2 + |a_i - b_i|} = \mathcal{J}_1(u, v)$$

## A2. Proof of lemma 4.1

The proof that the probabilistic Jaccard is a distance is based on the Steinhaus transform (Späth, 1981). Given a metric space  $(E, d)$  with a distance  $d$  and given a fixed point  $\alpha \in E$ , we can define a new distance  $d_{new}$  as:

$$d_{new}(x, y) = \frac{d(x, y)}{d(x, \alpha) + d(y, \alpha) + d(x, y)}$$

Consequently, the probabilistic Jaccard loss distance  $\mathcal{J}_c$  defined in (A.2):

$$\forall u, v \in [0, 1]^{C \times N}, \mathcal{J}_c(u, v) = \frac{2 \|u_c - v_c\|_1}{\|u_c\|_1 + \|v_c\|_1 + \|u_c - v_c\|_1} \quad (\text{A.3})$$

can be seen as a Steinhaus transform of the metric space  $([0, 1]^N, \|\cdot\|_1)$  with  $\alpha = 0$  and thus is a distance. Given that the weighted sum of distances is a distance, we finally conclude that the probabilistic multi-class Jaccard defined as:

$$\mathcal{J} = \sum_{c \in C} \mathcal{J}_c \quad (\text{A.4})$$

is a distance.

## Appendix B. Proof of Proposition 1

First, Eqs. (3), (9) and (7) are combined:

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_{lesion}} [\mathcal{L}(h_\theta(x), y)] &\leq \mathcal{R}_{seg} + \epsilon_{lesion}(\theta) - \epsilon_{control}(\theta) \\ &\leq \mathcal{R}_{seg} + |\epsilon_{lesion}(\theta) - \epsilon_{control}(\theta)| \end{aligned} \quad (\text{B.1})$$

where  $\epsilon_{lesion}(\theta)$  and  $\epsilon_{control}(\theta)$  denote the expected tissue loss on the *lesion* and *control* domains, defined as:

$$\begin{aligned} \epsilon_{lesion}(\theta) &= \mathbb{E}_{\mathcal{D}_{lesion}} [\mathcal{L}^T(h_\theta(x^{T_1}), y^T)] \\ \epsilon_{control}(\theta) &= \mathbb{E}_{\mathcal{D}_{control}} [\mathcal{L}^T(h_\theta(x^{T_1}), y^T)] \end{aligned} \quad (\text{B.2})$$

Let  $\theta^* = \arg \min_{\theta \in \Theta} \epsilon_{lesion}(\theta) + \epsilon_{control}(\theta)$  be the parameters of the ideal (and unknown) segmenter that minimises the two expected tissue losses. Then let denote  $\epsilon_{lesion}(\theta, \theta^*)$  and  $\epsilon_{control}(\theta, \theta^*)$  the performance gap between the segmenter parametrised by  $\theta$  and this ideal segmenter:

$$\begin{aligned} \epsilon_{lesion}(\theta, \theta^*) &= \mathbb{E}_{\mathcal{D}_{lesion}} [\mathcal{L}^T(h_\theta(x^{T_1}), h_{\theta^*}(x^{T_1}))] \\ \epsilon_{control}(\theta, \theta^*) &= \mathbb{E}_{\mathcal{D}_{control}} [\mathcal{L}^T(h_\theta(x^{T_1}), h_{\theta^*}(x^{T_1}))] \end{aligned} \quad (\text{B.3})$$

Using the fact that the loss function satisfies the triangle inequality:

$$\begin{aligned} \epsilon_{lesion}(\theta) &\leq \epsilon_{lesion}(\theta, \theta^*) + \epsilon_{lesion}(\theta^*) \\ \epsilon_{control}(\theta) &\leq \epsilon_{control}(\theta, \theta^*) + \epsilon_{control}(\theta^*) \end{aligned} \quad (\text{B.4})$$

Then, the performance gap between the two domains  $|\epsilon_{lesion}(\theta) - \epsilon_{control}(\theta)|$  can be bounded as follows:

$$|\epsilon_{lesion}(\theta) - \epsilon_{control}(\theta)| \leq |\epsilon_{lesion}(\theta, \theta^*) - \epsilon_{control}(\theta, \theta^*)| + \epsilon(\Theta) \quad (\text{B.5})$$

where  $\epsilon(\Theta)$  is the tissue expected loss of the ideal segmenter:

$$\epsilon(\Theta) = \epsilon_{lesion}(\theta^*) + \epsilon_{control}(\theta^*)$$

(B.5) can be found in Ben-David et al. (2010) in which the loss function is assumed to be the  $L_1$  distance.

Given that  $\epsilon(\Theta)$  is a constant w.r.t the network parameters  $\theta$ , the goal of domain adaptation is to reduce the distribution discrepancy  $d_{DA}(\theta)$  defined as:

$$d_{DA}(\theta) = |\epsilon_{lesion}(\theta, \theta^*) - \epsilon_{control}(\theta, \theta^*)| \quad (\text{B.6})$$

Similarly to Long et al. (2018), we demonstrate that this discrepancy  $d_{DA}(\theta)$  can be estimated using the discriminator accuracy.

Let denote  $\mathcal{D}_l^\theta = (x^{T_1}, f_\theta(x^{T_1}))_{x^{T_1} \sim \mathcal{D}_{lesion}}$  and  $\mathcal{D}_c^\theta = (x^{T_1}, f_\theta(x^{T_1}))_{x^{T_1} \sim \mathcal{D}_{control}}$  the proxies of the distributions  $\mathcal{D}_{lesion}$  and  $\mathcal{D}_{control}$ . Then, the two performance gaps with the ideal segmenter can be re-written as:

$$\begin{aligned} \epsilon_{lesion}(\theta, \theta^*) &= \mathbb{E}_{(x, f) \sim \mathcal{D}_l^\theta} [\mathcal{L}^T(h_{\theta^*}(x^{T_1}), f)] \\ \epsilon_{control}(\theta, \theta^*) &= \mathbb{E}_{(x, f) \sim \mathcal{D}_c^\theta} [\mathcal{L}^T(h_{\theta^*}(x^{T_1}), f)] \end{aligned} \quad (\text{B.7})$$

Let also define a difference hypothesis space  $\Delta$ :

$$\Delta \triangleq \{\delta_{\theta'} : (x^{T_1}, f) \mapsto \mathcal{L}^T(h_{\theta'}(x^{T_1}), f), \theta' \in \Theta\}$$

Moreover, we define the  $\Delta$ -distance between the two distributions  $\mathcal{D}_l^\theta$  and  $\mathcal{D}_c^\theta$  as:

$$d_\Delta(\mathcal{D}_l^\theta, \mathcal{D}_c^\theta) \triangleq \sup_{\delta_{\theta'} \in \Delta} \left| \mathbb{E}_{\mathcal{D}_l^\theta} [\delta_{\theta'}(x^{T_1}, f)] - \mathbb{E}_{\mathcal{D}_c^\theta} [\delta_{\theta'}(x^{T_1}, f)] \right| \quad (\text{B.8})$$

Finally, by combining (B.6), (B.7) and (B.8), we obtain the following upper bound for the distribution discrepancy  $d_{DA}(\theta)$ :

$$\begin{aligned} d_{DA}(\theta) &= \sup_{\delta_{\theta'} \in \Delta} \left| \mathbb{E}_{\mathcal{D}_l^\theta} [\delta_{\theta'}(x^{T_1}, f)] - \mathbb{E}_{\mathcal{D}_c^\theta} [\delta_{\theta'}(x^{T_1}, f)] \right| \\ &= \sup_{\theta' \in \Theta} \left| \mathbb{E}_{\mathcal{D}_l^\theta} [\mathcal{L}^T(h_{\theta'}(x^{T_1}), f)] - \mathbb{E}_{\mathcal{D}_c^\theta} [\mathcal{L}^T(h_{\theta'}(x^{T_1}), f)] \right| \\ &\geq \underbrace{\left| \mathbb{E}_{\mathcal{D}_l^\theta} [\mathcal{L}^T(h_{\theta^*}(x^{T_1}), f)] - \mathbb{E}_{\mathcal{D}_c^\theta} [\mathcal{L}^T(h_{\theta^*}(x^{T_1}), f)] \right|}_{=d_{DA}(\theta)} \end{aligned} \quad (\text{B.9})$$

Finally, let  $K > 0$  be the upper bound of the loss function  $\mathcal{L}^T$  and  $\mathcal{H}_\Phi^K$  denote the family of the discriminators multiplied by  $K$ :


$$\mathcal{H}_\Phi^K \triangleq \{K D_\phi, \phi \in \Phi\}$$

$\mathcal{H}_\Phi^K$  and the difference hypothesis space  $\Delta$  are two continuous and  $K$ -bounded function classes. Similarly to Long et al. (2018), let's assume that the family of the discriminator  $\mathcal{H}_\Phi^K$  is rich enough to contain the difference hypothesis space  $\Delta$ . Given that a multilayer perceptrons that can fit any functions, this assumption is not unrealistic. Then, we show that the discriminator accuracy of the best discriminator is an upper bound of the  $\Delta$ -distance:

$$\begin{aligned} d_\Delta(\mathcal{D}_l^\theta, \mathcal{D}_c^\theta) &\leq K \sup_{\phi} \left| \mathbb{E}_{\mathcal{D}_l^\theta} [D(x^{T_1}, f)] - \mathbb{E}_{\mathcal{D}_c^\theta} [D(x^{T_1}, f)] \right| \\ &\leq K \sup_{\phi} \left| \mathbb{E}_{\mathcal{D}_l^\theta} [D(x^{T_1}, f)] + \mathbb{E}_{\mathcal{D}_c^\theta} [1 - D(x^{T_1}, f)] \right| \quad (\text{B.10}) \\ &= K \sup_{\phi} \mathcal{R}_{DA}(\phi, \theta) \end{aligned}$$

Finally, by combining (B.1), (B.5), (B.9) (B.10), we obtain the following tractable upper bound for the expected segmentation loss:

$$\mathbb{E}_{\mathcal{D}_{lesion}} [\mathcal{L}(h_\theta(x), y)] \leq \mathcal{R}_{seg} + K \sup_{\phi} \mathcal{R}_{DA}(\phi, \theta) + \epsilon(\Theta) \quad (\text{B.11})$$



**Fig. C.9.** In order to synthesise a pseudo-healthy set of scans, we symmetrized the “healthy” hemisphere of brains from BraTS. GIF framework is then used to generate tissue ground truth.  $T_1$  scans are shown with the tissue segmentation.

### Appendix C. Visualisation symmetrised brain scans

Fig. C.9 shows some examples of pseudo-healthy scans, with their tissue annotations, synthesised as described in 5.2.

### References

- Ashburner, J., Friston, K.J., 2000. Voxel-based morphometry – the methods. *Neuroimage* 11 (6), 805–821.
- Bakas, S., Reyes, M., et al., Identifying the Best Machine Learning Algorithms for Brain Tumor Segmentation, Progression Assessment, and Overall Survival Prediction in the BRATS Challenge. arXiv:1811.02629.
- Barber, P.A., Demchuk, A.M., Zhang, J., Buchan, A.M., 2000. Validity and reliability of a quantitative computed tomography score in predicting outcome of hyperacute stroke before thrombolytic therapy. *The Lancet* 355 (9216), 1670–1674.
- Battaglini, M., Jenkinson, M., De Stefano, N., 2012. Evaluating and reducing the impact of white matter lesions on brain volume measurements. *Hum. Brain Mapp.* 33 (9), 2062–2071.
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W., 2010. A theory of learning from different domains. *Mach. Learn.* 79 (1), 151–175.
- Bermel, R.A., Bakshi, R., 2006. The measurement and clinical relevance of brain atrophy in multiple sclerosis. *The Lancet Neurology* 5 (2), 158–170.
- Bilen, H., Vedaldi, A., 2016. Weakly supervised deep detection networks. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2846–2854.
- Bitar, R., Leung, G., Perng, R., Tadros, S., Moody, A.R., Sarrazin, J., McGregor, C., Christakis, M., Symons, S., Nelson, A., Roberts, T.P., 2006. MR Pulse sequences: what every radiologist wants to know but is afraid to ask. *Radiographics* 26 (2), 513–537.
- Bottou, L., Curtis, F.E., Nocedal, J., 2018. Optimization methods for large-scale machine learning. *SIAM Rev.* 60 (2), 223–311.
- Bragman, F.J.S., Tanno, R., Eaton-Rosen, Z., Li, W., Hawkes, D.J., Ourselin, S., Alexander, D.C., McClelland, J.R., Cardoso, M.J., 2018. Uncertainty in multitask learning: Joint representations for probabilistic MR-only radiotherapy planning. In: *Proceedings of the 21st International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'18)*, pp. 3–11.
- Cardoso, M.J., Modat, M., Wolz, R., Melbourne, A., Cash, D., Rueckert, D., Ourselin, S., 2015. Geodesic information flows: spatially-variant graphs and their application to segmentation and fusion. *IEEE Trans. Med. Imaging* 34 (9), 1976–1988.
- Chen, S., Bortsova, G., García-Uceda Juárez, A., van Tulder, G., de Bruijne, M., 2019. Multi-task attention-based semi-supervised learning for medical image segmentation. In: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.-T., Khan, A. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*. Springer International Publishing, Cham, pp. 457–465.
- Chen, X., Konukoglu, E., 2018. Unsupervised detection of lesions in brain MRI using constrained adversarial auto-encoders. *CoRR* abs/1806.04972.
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O., 2016. 3d u-net: Learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*. Springer International Publishing, Cham, pp. 424–432.
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O., 2016. 3d u-net: Learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*. Springer International Publishing, Cham, pp. 424–432.
- Csurka, G., 2017. A Comprehensive Survey on Domain Adaptation for Visual Applications. In: *Domain Adaptation in Computer Vision Applications*. Springer International Publishing, pp. 1–35.
- Dorent, R., Joutard, S., Modat, M., Ourselin, S., Vercauteren, T., 2019. Hetero-modal variational encoder-decoder for joint modality completion and segmentation. In: *Proceedings of the 22nd International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'19)*, p. Inpress.
- Dorent, R., Li, W., Ekanayake, J., Ourselin, S., Vercauteren, T., 2019. Learning joint lesion and tissue segmentation from task-specific hetero-modal datasets. In: *Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning*. PMLR, London, United Kingdom, pp. 164–174.
- Dou, Q., Ouyang, C., Chen, C., Chen, H., Heng, P.-A., 2018. Unsupervised cross-modal-ity domain adaptation of convnets for biomedical image segmentations with adversarial loss. In: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, pp. 691–697.
- Dwyer, M.G., Bergsland, N., Ramasamy, D.P., Jakimovski, D., Weinstock-Guttman, B., Zivadinov, R., 2018. Atrophied brain lesion volume: a new imaging biomarker in multiple sclerosis. *Journal of Neuroimaging* 28 (5), 490–495.
- Fisniku, L.K., Chard, D.T., Jackson, J.S., Anderson, V.M., Altmann, D.R., Miszkil, K.A., Thompson, A.J., Miller, D.H., 2008. Gray matter atrophy is related to long-term disability in multiple sclerosis. *Ann. Neurol.* 64 (3), 247–254.
- Geurts, J.J., Calabrese, M., Fisher, E., Rudick, R.A., 2012. Measurement and clinical effect of grey matter pathology in multiple sclerosis. *The Lancet Neurology* 11 (12), 1082–1092.
- Giorgio, A., De Stefano, N., 2013. Clinical use of brain volumetry. *J. Magn. Reson. Imaging* 37 (1), 1–14.
- Havaei, M., Guizard, N., Chapados, N., Bengio, Y., 2016. HeMIS: Hetero-modal image segmentation. In: *Proceedings of the 19th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'16)*, pp. 469–477.
- He, K., Zhang, X., Ren, S., Sun, J., 2015. Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In: *Proceedings of the 25th International Conference on Computer Vision (ICCV'15)*, pp. 1026–1034.
- Iglesias, J.E., Liu, C., Thompson, P.M., Tu, Z., 2011. Robust brain extraction across datasets and comparison with publicly available methods. *IEEE Trans. Med. Imaging* 30 (9), 1617–1634.
- Iglesias, J.E., Sabuncu, M.R., 2015. Multi-atlas segmentation of biomedical images: a survey. *Med. Image Anal.* 24 (1), 205–219.
- Isensee, F., Petersen, J., Kohl, S.A.A., Jäger, P.F., Maier-Hein, K.H., 2019. Nnu-net: breaking the spell on successful medical image segmentation. *CoRR* abs/1904.08128.
- Jack Jr, C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britton, P.J., L. Whitwell, J., Ward, C., Dale, A.M., Felmlee, J.P., Gunter, J.L., Hill, D.L., Killiany, R., Schuff, N., Fox-Bosetti, S., Lin, C., Studholme, C., DeCarli, C.S., Krueger, G., Ward, H.A., Metzger, G.J., Scott, K.T., Mallozzi, R., Blezek, D., Levy, J., Debbins, J.P., Fleisher, A.S., Albert, M., Green, R., Bartzokis, G.,



- Glover, G., Mugler, J., Weiner, M.W., 2008. The alzheimer's disease neuroimaging initiative (adni): mri methods. *J. Magn. Reson. Imaging* 27 (4), 685–691.
- Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M., 2012. FSL. *Neuroimage* 62 (2), 782–790.
- Kamnitsas, K., Baumgartner, C., Ledig, C., Newcombe, V., Simpson, J., Kane, A., Menon, D., Nori, A., Criminisi, A., Rueckert, D., et al., 2017. Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. In: Proceedings of the 20th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'17), pp. 597–609.
- Kendall, A., Gal, Y., 2017. What uncertainties do we need in Bayesian deep learning for computer vision? In: Proceedings of Advances in Neural Information Processing Systems 30 (NIPS 2017), pp. 5574–5584.
- Kim, D.-J., Choi, J., Oh, T.-H., Yoon, Y., Kweon, I.S., 2018. Disjoint multi-task learning between heterogeneous human-centric tasks. In: Proceedings of the 2018 IEEE Winter Conference on Applications of Computer Vision (WACV) (2018), pp. 1699–1708.
- Kingma, D.P., Ba, J., 2015. Adam: A method for stochastic optimization. In: 3rd International Conference on Learning Representations, ICLR 2015, pp. 1–15.
- Kosub, S., 2018. A note on the triangle inequality for the jaccard distance. *Pattern Recognit. Lett.*
- Kuijff, H.J., Bennink, E., 2018. MRBrainS18. <http://mrbrains18.isi.uu.nl/>.
- Kuijff, H.J., Biesbroek, J.M., de Bresser, J., Heinen, R., Andermatt, S., Bento, M., Berseth, M., Belyaev, M., Cardoso, M.J., Casamitjana, A., Collins, D.L., Dadar, M., Georgiou, A., Ghafoorian, M., Jin, D., Khademi, A., Knight, J., Li, H., Lladó, X., Luna, M., Mahmood, Q., McKinley, R., Mehrtash, A., Ourselin, S., Park, B., Park, H., Park, S.H., Pezold, S., Puybareau, E., Rittner, L., Sudre, C.H., Valverde, S., Vilaplana, V., Wiest, R., Xu, Y., Xu, Z., Zeng, G., Zhang, J., Zheng, G., Chen, C., van der Flier, W., Barkhof, F., Viergever, M.A., Biessels, G.J., 2019. Standardized assessment of automatic segmentation of white matter hyperintensities; results of the WMH segmentation challenge. *IEEE Trans. Med. Imaging*.
- Le, T.-L.-T., Thome, N., Bernard, S., Bismuth, V., Patoureaux, F., 2019. Multitask classification and segmentation for cancer diagnosis in mammography. In: Cardoso, M.J., Feragen, A., Glocker, B., Konukoglu, E., Oguz, I., Unal, G., Vercauteren, T. (Eds.), Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning – Extended abstract track, pp. 1–4.
- Li, W., Wang, G., Fidon, L., Ourselin, S., Cardoso, M.J., Vercauteren, T., 2017. On the compactness, efficiency, and representation of 3D convolutional networks: Brain parcellation as a pretext task. In: Proceedings of Information Processing in Medical Imaging (IPMI'17), pp. 348–360.
- Li, Z., Hoiem, D., 2017. Learning without forgetting. *IEEE Trans. Pattern Anal. Mach. Intell.* 40 (12), 2935–2947.
- Long, M., Cao, Y., Wang, J., Jordan, M., 2015. Learning transferable features with deep adaptation networks. In: Bach, F., Blei, D. (Eds.), Proceedings of the 32nd International Conference on Machine Learning. PMLR, Lille, France, pp. 97–105.
- Long, M., Cao, Y., Wang, J., Jordan, M.I., 2018. Conditional Adversarial Domain Adaptation. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (Eds.), Advances in Neural Information Processing Systems 31. Curran Associates, Inc., pp. 1640–1650.
- Maillard, P., Carmichael, O., Harvey, D., Fletcher, E., Reed, B., Mungas, D., DeCarli, C., 2013. FLAIR And diffusion MRI signals are independent predictors of white matter hyperintensities. *American Journal of Neuroradiology* 34 (1), 54–61.
- Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L., 2007. Open access series of imaging studies (OASIS): cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *J. Cogn. Neurosci.* 19 (9), 1498–1507.
- Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., Weber, M., Arbel, T., Avants, B.B., Ayache, N., Buendia, P., Collins, D.L., Cordier, N., Corso, J.J., Criminisi, A., Das, T., Delingette, H., Demiralp, Ç., Durst, C.R., Dojat, M., Doyle, S., Festa, J., Forbes, F., Geremia, E., Glocker, B., Golland, P., Guo, X., Hamamci, A., Iftekharuddin, K.M., Jena, R., John, N.M., Konukoglu, E., Lashkari, D., Mariz, J.A., Meier, R., Pereira, S., Precup, D., Price, S.J., Raviv, T.R., Reza, S.M.S., Ryan, M., Sarikaya, D., Schwartz, L., Shin, H., Shotton, J., Silva, C.A., Sousa, N., Subbanna, N.K., Szekely, G., Taylor, T.J., Thomas, O.M., Tustison, N.J., Unal, G., Vasseur, F., Wintermark, M., Ye, D.H., Zhao, L., Zhao, B., Zikic, D., Prastawa, M., Reyes, M., Van Leemput, K., 2015. The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Trans. Med. Imaging* 34 (10), 1993–2024.
- Moeskops, P., de Bresser, J., Kuijff, H.J., Mendrik, A.M., Biessels, G.J., Pluim, J.P., Išgum, I., 2018. Evaluation of a deep learning approach for the segmentation of brain tissues and white matter hyperintensities of presumed vascular origin in MRI. *NeuroImage: Clinical* 17, 251–262.
- Moeskops, P., Wolterink, J.M., van der Velden, B.H.M., Gilhuijs, K.G.A., Leiner, T., Viergever, M.A., Išgum, I., 2016. Deep learning for multi-task medical image segmentation in multiple modalities. In: Ourselin, S., Joscovic, L., Sabuncu, M.R., Unal, G., Wells, W. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016. Springer International Publishing, Cham, pp. 478–486.
- Oquab, M., Bottou, L., Laptev, I., Sivic, J., 2015. Is object localization for free? – Weakly-supervised learning with convolutional neural networks. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 685–694.
- Orbes-Arteaga, M., Varsavsky, T., Sudre, C.H., Eaton-Rosen, Z., Haddow, L.J., Sørensen, L., Nielsen, M., Pai, A., Ourselin, S., Modat, M., Nachev, P., Cardoso, M.J., 2019. Multi-domain adaptation in brain mri through paired consistency and adversarial learning. In: DART/MIL3ID/MICCAI, pp. 54–62.
- Pan, S.J., Tsang, I.W., Kwok, J.T., Yang, Q., 2011. Domain adaptation via transfer component analysis. *IEEE Trans. Neural Networks* 22 (2), 199–210.
- Pérez-García, F., Sparks, R., Ourselin, S., 2020. Torchio: a python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning arXiv:2003.04696.
- Popescu, V., Agosta, F., Hulst, H.E., Sluimer, I.C., Knol, D.L., Sormani, M.P., Enzinger, C., Ropele, S., Alonso, J., Sastre-Garriga, J., Rovira, A., Montalban, X., Bodini, B., Ciccarelli, O., Khaleeli, Z., Chard, D.T., Matthews, L., Palace, J., Giorgio, A., De Stefano, N., Eisele, P., Gass, A., Polman, C.H., Uitdehaag, B.M.J., Messina, M.J., Comi, G., Filippi, M., Barkhof, F., Vrenken, H., 2013. Brain atrophy and lesion load predict long term disability in multiple sclerosis. *Journal of Neurology, Neurosurgery & Psychiatry* 84 (10), 1082–1091. doi:10.1136/jnnp-2012-304094.
- Prados, F., Cardoso, M.J., Kanber, B., Ciccarelli, O., Kapoor, R., Wheeler-Kingshott, C.A.G., Ourselin, S., 2016. A multi-time-point modality-agnostic patch-based method for lesion filling in multiple sclerosis. *Neuroimage* 139, 376–384.
- Prima, S., Ourselin, S., Ayache, N., 2002. Computation of the mid-sagittal plane in 3-D brain images. *IEEE Trans. Med. Imaging* 21 (2), 122–138.
- Roulet, N., Slezak, D.F., Ferrante, E., 2019. Joint learning of brain lesion and anatomy segmentation from heterogeneous datasets. In: Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning. PMLR, pp. 401–413.
- Ruder, S., 2017. An overview of multi-task learning in deep neural networks arXiv:1706.05098.
- Saha, O., Sathish, R., Sheet, D., 2019. Learning with multitask adversaries using weakly labelled data for semantic segmentation in retinal images. In: Cardoso, M.J., Feragen, A., Glocker, B., Konukoglu, E., Oguz, I., Unal, G., Vercauteren, T. (Eds.), Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning. PMLR, London, United Kingdom, pp. 414–426.
- Sdika, M., Pelletier, D., 2009. Nonrigid registration of multiple sclerosis brain images using lesion inpainting for morphometry or lesion mapping. *Hum. Brain Mapp.* 30 (4), 1060–1067.
- Shaw, R., Sudre, C., Ourselin, S., Cardoso, M.J., 2019. Mri k-space motion artefact augmentation: Model robustness and task-specific uncertainty. In: Cardoso, M.J., Feragen, A., Glocker, B., Konukoglu, E., Oguz, I., Unal, G., Vercauteren, T. (Eds.), Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning. PMLR, London, United Kingdom, pp. 427–436.
- Simpson, A.L., Antonelli, M., Bakas, S., Bilello, M., Farahani, K., van Ginneken, B., Kopp-Schneider, A., Landman, B.A., Litjens, G. J.S., Menze, B.H., Ronneberger, O., Summers, R.M., Bilic, P., Christ, P.F., Do, R.K.G., Gollub, M., Golia-Pernicka, J., Heckers, S., Jarnagin, W.R., McHugo, M., Napel, S., Vorontsov, E., Maier-Hein, L., Cardoso, M.J., 2019. A large annotated medical image dataset for the development and evaluation of segmentation algorithms. arXiv:1902.09063.
- Späth, H., 1981. The minimum location problem for the jaccard metric. *Operations-Research-Spektrum* 3 (2), 91–94.
- Sudre, C.H., Cardoso, M.J., Ourselin, S., 2017. Longitudinal segmentation of age-related white matter hyperintensities. *Med. Image Anal.* 38, 50–64.
- Sun, B., Feng, J., Saenko, K., 2016. Return of frustratingly easy domain adaptation. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence. AAAI Press, pp. 2058–2065.
- Sun, L., Wang, J., Ding, X., Huang, Y., Paisley, J.W., 2018. An adversarial learning approach to medical image synthesis for lesion removal arXiv:1810.10850.
- Ulyanov, D., Vedaldi, A., Lempitsky, V.S., 2016. Instance normalization: The missing ingredient for fast stylization. arXiv:1607.08022.
- Valverde, S., Oliver, A., Lladó, X., 2014. A white matter lesion-filling approach to improve brain tissue volume measurements. *NeuroImage: Clinical* 6, 86–92.
- Wen, P.Y., Macdonald, D.R., Reardon, D.A., Cloughesy, T.F., Sorensen, A.G., Galanis, E., DeGroot, J., Wick, W., Gilbert, M.R., Lassman, A.B., Tsien, C., Mikkelsen, T., Wong, E.T., Chamberlain, M.C., Stupp, R., Lamborn, K.R., Vogelbaum, M.A., van den Bent, M.J., Chang, S.M., 2010. Updated response assessment criteria for high-grade gliomas: response assessment in neuro-oncology working group. *Journal of Clinical Oncology* 28 (11), 1963–1972.
- Xia, T., Chartsias, A., Tsaftaris, S.A., 2019. Adversarial pseudo healthy synthesis needs pathology factorization. In: Cardoso, M.J., Feragen, A., Glocker, B., Konukoglu, E., Oguz, I., Unal, G., Vercauteren, T. (Eds.), Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning. PMLR, London, United Kingdom, pp. 512–526.
- Xu, Y., Zhu, J.-Y., Chang, E.I.-C., Lai, M., Tu, Z., 2014. Weakly supervised histopathology cancer image segmentation and classification. *Med. Image Anal.* 18 (3), 591–604.
- Zhang, Y., Yang, Q., 2017. A survey on multi-task learning. arXiv:1707.08114.
- Zhao, H., des Combes, R.T., Zhang, K., Gordon, G.J., 2019. On learning invariant representations for domain adaptation. In: Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9–15 June 2019, Long Beach, California, USA, pp. 7523–7532.