Realistic Image Normalization for Multi-Domain Segmentation

Pierre-Luc Delisle[a,*], Benoit Anctil-Robitaille[a], Christian Desrosiers[a], Herve Lombaert[a]

[a]*Department of Computer and Software Engineering, ETS Montreal, Canada*

Abstract

Image normalization is a building block in medical image analysis. Conventional approaches are customarily utilized on a per-dataset basis. This strategy, however, prevents the current normalization algorithms from fully exploiting the complex joint information available across multiple datasets. Consequently, ignoring such joint information has a direct impact on the performance of segmentation algorithms. This paper proposes to revisit the conventional image normalization approach by instead learning a common normalizing function across multiple datasets. Jointly normalizing multiple datasets is shown to yield consistent normalized images as well as an improved image segmentation. To do so, a fully automated adversarial and task-driven normalization approach is employed as it facilitates the training of realistic and interpretable images while keeping performance on-par with the state-of-the-art. The adversarial training of our network aims at finding the optimal transfer function to improve both the segmentation accuracy and the generation of realistic images. We evaluated the performance of our normalizer on both infant and adult brains images from the iSEG, MRBrainS and ABIDE datasets. Results reveal the potential of our normalization approach for segmentation, with Dice improvements of up to 57.5% over our baseline. Our method can also enhance data availability by increasing the number of samples available when learning from multiple imaging domains.

Keywords: 3D MRI, Brain segmentation, Data harmonization, Generative adversarial networks, Intensity normalization

1. Introduction

Powered by their capacity to learn hierarchical feature representations from data, deep learning algorithms have achieved unprecedented performance in a broad range of medical imaging applications. Notably, deep convolutional neural networks (CNNs) have helped improve the segmentation of various anatomical structures in medical images, for instance brain regions in 3D MRI (Dolz et al., 2019; Kamnitsas et al., 2017). However, supervised learning algorithms typically require a large amount of labeled data for training. Obtaining such quantities is often difficult, since the manual labeling of images is a complex and time-consuming process performed by highly-trained clinical experts.

A possible approach to both alleviate the lack of training data and increase the generalization performance of the learning algorithm is to use data acquired from multiple sites. However, medical images from separate datasets can be acquired with distinct scanner models or parameters, and therefore may present drastic differences in their images intensities. Acquisition standards could be normalized but reconstructed images still exhibit important differences across sites (Kochunov et al., 2014). Another potential source of variability in intensities can also arise from differences in patient demographics between datasets. This problem is illustrated in Fig. 2, where image intensity histograms are shown for two public datasets, MR-BrainS, on adult brains, and iSEG, on 6-8 months old infants.
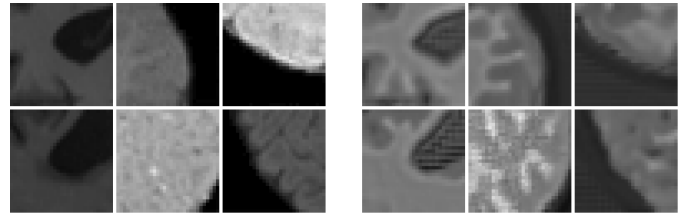

Figure 1: Mixed iSEG and MRBrainS inputs (**left**) and images generated with two pipelined FCNs without constraint on realism using only Dice loss (**right**). Images generated with only Dice loss preserve the structure required for segmentation but lack realism.

These histograms indicate differences in the distribution overlap of tissue classes across datasets, which directly impair any subsequent segmentation processing.

Standard deep learning models are sensitive to the data distribution on which they are trained. This leads to sub-optimal performance when evaluating on different sets of medical images (Mårtensson et al., 2020). A common strategy to address this problem is to normalize images in a pre-processing step, for instance, so that their intensities fall in the same range or have the same global mean (Onofrey et al., 2019). However, this naive approach is generally insufficient for tasks such as segmentation since it does not consider the intensity distribution of individual regions in the image. Hence, a normalization of images yielding the same global average intensity may still result in different class-specific distributions, directly impacting any downstream analysis.

Recent work has investigated the potential of deep neural networks for data-driven image normalization. Instead of feed-

---
[*]Corresponding author: pierre-luc.delisle.1@etsmtl.net

ing images directly to a segmentation network, Drozdzal et al. (2018) employed a second CNN as a pre-processing network to normalize images prior to their segmentation. Their strategy of learning image-specific normalization has led to a better segmentation performance with images of different characteristics. However, as shown in Fig. 1, since there is no constraint on the realism of images produced by the normalization network, these images typically lack interpretability across datasets. While their work has focused on a single data source, other studies have considered the problem of harmonizing data across multiple sites. For instance, DeepHarmony (Dewey et al., 2019) uses a fully-convolutional CNN architecture to translate images from one acquisition protocol to another. Despite improved results, important limitations remain: 1) it requires having images of the same subjects for different protocols; 2) it cannot be easily extended to more than two sites since it relies on learning a protocol-to-protocol mapping; 3) it still needs several pre-processing steps to mitigate image inhomogeneity and perform gain correction. A normalizing approach without such paired images of the same subjects across protocols has been explored (Modanwal et al., 2020) using a cycle-consistent generative adversarial network (CycleGAN). However, such approach is not tailored to a specific task such as segmentation, and can consequently lead to sub-optimal results in the downstream analyses. Nonetheless, GANs proved to be powerful at generating medical images (Nie et al., 2018).

1.1. Contributions

We address the limitations of existing image normalization approaches with a novel adversarial learning method that generates normalized images that are both interpretable by clinicians and optimized for a downstream segmentation task. Our method leverages information from multiple datasets by learning a joint normalizing transformation accounting for large image variability. This is achieved with a deep learning architecture comprised of two fully-convolutional 3D CNNs (Long et al., 2015), the first one acting as a normalized image generator and the second one used as segmentation network. During training, our model also includes a 3D CNN-based discriminator (He et al., 2016) which serves as a domain classifier. In standard adversarial learning approaches for image generation, the discriminator tries to classify images as real or fake (Goodfellow et al., 2014). On the other hand, typical approaches for domain adaptation instead use the discriminator for predicting if an image is from a source or a target domain (Kamnitsas et al., 2017; van Opbroek et al., 2015; Cheplygina et al., 2019). In our proposed normalization method, the discriminator distinguishes images between all input domains (i.e., acquisition site and/or protocol) *as well* as an additional "generated" class. Hence, the produced images are both realistic and domain-invariant.

Our contributions can be summarized as follows:

- A first learned normalization method for medical images producing images that are both optimized for segmentation and interpretable by clinicians. Compared to recent approaches, such as those based on CycleGANs, our
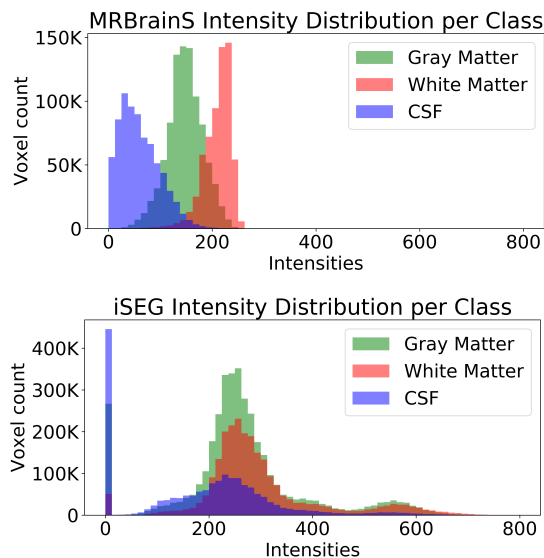

Figure 2: Intensity histograms of different brain tissue classes for adult brains in the MRBrainS dataset and 6-8 month infants in the iSEG dataset. We can see the important overlap in intensities, especially for iSEG, which is the cause of misleading classifiers. One can also notice the intensity range being largely different between both datasets.

method can accommodate an arbitrary number of data source domains without additional complexity.

- A novel adversarial learning model for 3D image processing that jointly optimizes three convolutional neural networks. Unlike standard adversarial techniques which have separate discriminators for domain classification and differentiating generated images from real ones, our model combines these two tasks in a single network using an additional domain class. As theoretical contribution, we show that optimizing this model corresponds to minimizing the KL divergence between the generated image probability distribution of each domain and the mean distribution of real images.

- One of the most comprehensive experimental analyses of learned normalization models for medical images, which includes three very different brain MRI datasets and also evaluates our method on multi-modal data and images degraded by bias field. This analysis demonstrates the advantages of our approach compared to the recent state-of-the-art.

After an overview of related work, the methodology describes our task-driven normalization method. The experimental results follow with an evaluation on: 1) the cross-data performance of a standard supervised baseline; 2) the normalization and segmentation performance of our method on data from two or more sites; 3) the capacity to normalize multi-modal images; and 4) the ability to correct for intensity inhomogeneity in images with strong bias field. The conclusion finally summarizes the main contributions and results of our work, and discusses possible extensions.

1.2. Related Work

Image Normalization. A plethora of pre-processing techniques exists to normalize medical images prior to any image analysis. One common approach, known as standardization (Birenbaum and Greenspan, 2017; Kamnitsas et al., 2017; Casamitjana et al., 2017; Chen et al., 2018), consists of normalizing each pixel intensity value in an input image by subtracting from it the image average intensity and dividing it by the its standard deviation. However, this simple strategy does not take into account the global statistics of the dataset. Other pre-processing approaches, such as histogram equalization (Onofrey et al., 2019) and bias field correction (Birenbaum and Greenspan, 2017; Baid et al., 2018; Feng et al., 2019), are also commonly used to mitigate the problem of intensity inhomogeneity in images. For instance, Onofrey et al. (2019) evaluate the benefit of using different normalization techniques to multi-site prostate MRI before applying deep learning-based segmentation. Recently, a few studies have explored the potential of learning methods for dynamic data augmentation and normalization (Drozdzal et al., 2018; Ciga et al., 2019; Hesse et al., 2020) as well as image denoising (Oguz et al., 2020). Drozdzal et al. (2018) use two consecutive fully-convolutional CNNs, a pre-processor network followed by a segmentation network trained with a Dice metric, to normalize an input image prior to segmentation. Despite showing a better performance compared to the segmentation of unnormalized images, this prior work has several limitations. First, since there is no realism constraint on the intermediate images produced by the pre-processing network, these images lack interpretability across multiple datasets. Moreover, because images are encoded for a specific network, they cannot be used with other segmentation models without retraining. Last, the feasibility of this approach was not demonstrated for multi-site settings where the multiple datasets are used jointly to learn the normalization. A domain adaptation strategy that adds a domain classifier at the end of each layer of a classification network to extract domain-invariant features is explored in (Ciga et al., 2019). While it supports multiple domains, such as datasets acquired with different imaging protocols, this strategy is not tailored to a specific task like segmentation. Thus, it yields suboptimal results compared to a task-driven normalization approach. To this effect, we show in our experiments that learning a segmentation network jointly with the normalization network actually improves the contrast between different regions of interest (ROIs) in the normalized images.

Data Harmonization. The harmonization of data across multiple sites has also sparked interests in research, specially to increase the sample size of statistical studies (Logue et al., 2018). Shinohara et al. (2017) show that employing scanners from the same vendor and carefully harmonizing the protocols for the acquisition of multicenter 3D MRI brain data still results in systematic image differences. This impacts the accuracy of volumetric analyses, notably introducing bias in measured white and gray matter volumes. Moreover, multiplying the number of sites can introduce nonlinear age-related differences in ROIs within the brain. Pomponio et al. (2020) use data harmonization to remove site-related demographics effects in the cross-

sectional LIFESPAN dataset. However, data harmonization is only done after the segmentation of ROIs, adding an extra step to the processing pipeline. Data harmonization approaches based on deep learning have also been explored in recent work. Dewey et al. (2019) proposed an FCN architecture based on 3D U-Net for contrast harmonization between two different protocols, demonstrating a more consistent volume quantification across these protocols. However, training this architecture requires paired images of the same subjects with different acquisition protocols, which is challenging to obtain in practice and does not scale to multi-site data. Modanwal et al. (2020) use a cycle-consistent generative adversarial network (CycleGAN) to generate harmonized structural breast images between two different types of scanner. The proposed method leverages unpaired data with two generator-discriminator pairs to bypass common limitation of image translation algorithms that require paired data. This approach has also been explored for image denoising (Oguz et al., 2020). While the results show visually-realistic harmonized images, it does not consider the specific image analysis task performed after pre-processing. Moreover, this harmonization approach is limited to only two domains, and extending it to additional ones requires substantial modifications. In contrast, our method accommodates an arbitrary number of data sites without added complexity.

2. Material and Methods

2.1. The Proposed Architecture

Our adversarial image normalization architecture, illustrated in Fig. 3, consists of three main components: 1) a fully-convolutional Generator (G) which acts as an image pre-processor to produce normalized images; 2) a Segmenter (S) based on the same fully-convolutional network that outputs a segmentation map from normalized images; 3) a Discriminator (D) which tries to classify the domain of normalized images. In the following, we describe each of these components in greater details.

Generator. Although any other FCN architecture could be used, our Generator network G is based on a modified 3D U-Net (Çiçek et al., 2016) where the change resides primarily in the expanding path. The original 3D U-Net uses transposed 3D convolutions and concatenates feature maps to recover spatial resolution. Because of GPU memory constraints, we employ a simpler upscaling operator to upsample the resolution at each level of the U-Net decoder coupled with a feature concatenation. This enables the reduction of the total number of parameters in our model. We kept the same encoding path as in (Çiçek et al., 2016), which consists of alternating convolutions and max-pooling layers. The network also uses shortcut connections from encoding to decoding path between layers of equal resolution to help recover the high-resolution features. Each convolution is followed by a batch normalization and a ReLU activation function. We also kept the same number of feature maps in each convolution. During training and testing, the model takes a 3D patch $\mathbf{x} \in \mathbb{R}^{|\Omega|}$, where Ω is the set of patch voxels, and transforms it into a cross-domain normalized image $\widehat{\mathbf{x}} = G(\mathbf{x})$.
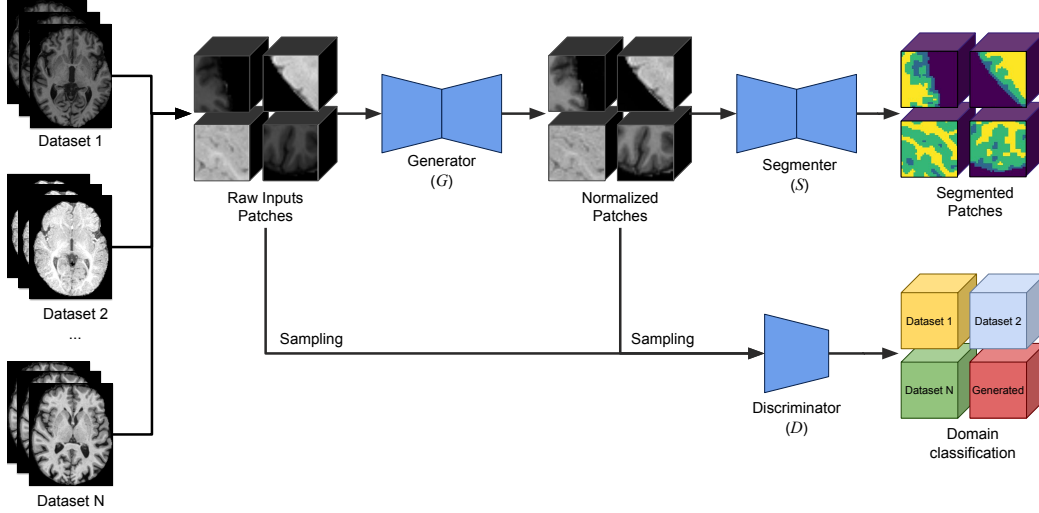
3

Figure 3: Proposed architecture. A first FCN *Generator* network (G) takes a non-normalized patch and generates a normalized patch. The normalized patch is input to a second FCN *Segmenter* network (S) for proper segmentation. *Discriminator* (D) network apply the constraint of realism on the normalized output. The algorithm learns the optimal normalizing function based on the observed differences between input datasets.

Segmenter. The segmentation network S uses the same 3D U-Net architecture as our Generator. This network receives the normalized output of the image generator and performs voxel-wise classification using a final 1^3 convolution before the softmax layer. The output of this network is the segmentation map of the input patch $S(\widehat{\mathbf{x}})$. While any other segmentation loss can be employed to train S, we chose the widely used Dice loss (Milletari et al., 2016; Carass et al., 2020) defined as

$$\mathcal{L}_{\text{seg}}(\mathbf{s},\mathbf{y}) \;=\; 1 - \frac{\epsilon + 2\sum_c \omega_c \sum_{v\in\Omega} s_{v,c} \cdot y_{v,c}}{\epsilon + \sum_c \omega_c \sum_{v\in\Omega} (s_{v,c} + y_{v,c})}, \qquad (1)$$

where $s_{v,c} \in [0,1]$ is the softmax output of S for voxel v and class c, $y_{v,c}$ is the corresponding ground-truth label (i.e., $y_{v,c}$ is 1 if the ground-truth class of voxel v is c, else it is 0), and ϵ is a small constant to avoid zero-division. The weight ω_c controls the influence of class c in the loss. It is typically tuned to alleviate the problem of class imbalance by giving a higher weight to smaller-region classes in images. Throughout all experiments, we empirically fixed these weights to $\omega_{\text{BG}} = 0.22$, $\omega_{\text{WM}} = 0.28$, $\omega_{\text{GM}} = 0.20$ and $\omega_{\text{CSF}} = 0.30$.

Discriminator. For the domain classifier D, we chose the DC-GAN (Radford et al., 2016) discriminator's architecture with 5 layers and adapted its implementation to 3D volumes. The first 4 layers consist of a 3D convolutional operation, a LeakyReLU activation, and a dropout operation. A final linear layer follows to ensure classification. Note that we also tried the more recent ResNet model, which employs residual connections to improve gradient flow during training. However, we found that this model tends to overfit and leads to a worse performance than DCGAN's discriminator. Our final classification network receives as inputs both raw dataset patches and normalized patches from the Generator. Each patch has an image domain label $z \in \{1, \dots, K+1\}$ which determines from which dataset the patch comes from (labels 1 to K) or if the patch is generated (label

$K+1$). The role of this discriminator is to ensure that images produced by G are both realistic and domain-invariant. Although other classification losses could be considered, we used the negative log likelihood (cross entropy) loss, i.e.,

$$\mathcal{L}_{\text{dis}}(D(\mathbf{x}), z) \;=\; -\log D_z(\mathbf{x}), \qquad (2)$$

where $D_z(\mathbf{x})$ is the softmax probability for class z. We note that, since $\sum_{z=1}^{K+1} D_z(\mathbf{x}) = 1$, the loss for the $(K+1)$-th class corresponding to generated examples can be written in terms of domain classes as

$$\mathcal{L}_{\text{dis}}(D(\mathbf{x}), K+1) \;=\; -\log\Big(1 - \sum_{z=1}^{K} D_z(\mathbf{x})\Big).$$

2.1.1. Adversarial Training

The three networks of our model are trained together in an adversarial manner by optimizing the following loss function:

$$\min_{G,S} \max_{D} \mathcal{L}(G,S,D) \;=\; \mathbb{E}_{\mathbf{x},\mathbf{y}}\big[\mathcal{L}_{\text{seg}}(S(G(\mathbf{x})),\mathbf{y})\big]$$
$$- \lambda \mathbb{E}_{\mathbf{x},z}\big[\mathcal{L}_{\text{dis}}(D(\mathbf{x}),z) + \mathcal{L}_{\text{dis}}(D(G(\mathbf{x})),K+1)\big] \tag{3}$$

where $\mathcal{L}_{\text{seg}}$ and $\mathcal{L}_{\text{dis}}$ are the segmentation and discriminator losses, respectively defined in Eq. (1) and Eq. (2). Hyper-parameter λ controls the trade-off between having a good segmentation accuracy (first loss term) and having normalized images which are domain-invariant (last loss term). By using $\lambda = 0$, the model becomes similar to (Drozdzal et al., 2018), where the generator is not constrained to produce realistic images. In contrast, for a large λ, our model becomes similar to an adversarial domain classifier (Ciga et al., 2019), where generated images are normalized across different domains but not optimal for segmentation, with added realism constraints.

4

Algorithm 1: Training of our adversarial image normalization method, with the interacting updates of the Discriminator, the Segmenter and the Generator.

Input: Training set $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i, z_i)\}_{i=1}^{|\mathcal{D}|}$
Input: Batch size m, number epochs, iterations and steps
   $(n_{epochs}, n_{iter}, n_{steps})$, and learning rates η_G, η_S, η_D;
Output: Network parameters $\theta_G, \theta_S, \theta_D$;

Randomly initialize network parameters $\theta_G, \theta_S, \theta_D$;

for $epoch = 1, \ldots, n_{epochs}$ **do**
  for $iteration = 1, \ldots, n_{iter}$ **do**
    for $step = 1, \ldots, n_{steps}$ **do**
      Sample batch of m examples from all domains
        $\{(\mathbf{x}_i, z_i)\}_{i=1}^m$;
      Update the Discriminator D:

$$\theta_D \leftarrow \theta_D - \frac{\eta_D}{m} \sum_{i=1}^m \Big(\nabla_D \mathcal{L}_{\text{dis}}(D(\mathbf{x}_i), z_i) \\ + \nabla_D \mathcal{L}_{\text{dis}}(D(G(\mathbf{x}_i)), K+1) \Big);$$

    Sample batch m examples from all domains
      $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^m$;
    Update the Segmenter S and Generator G:

$$\theta_S \leftarrow \theta_S - \frac{\eta_S}{m} \sum_{i=1}^m \nabla_S \mathcal{L}_{\text{seg}}(S(G(\mathbf{x}_i)), \mathbf{y}_i);$$

$$\theta_G \leftarrow \theta_G - \frac{\eta_G}{m} \sum_{i=1}^m \Big(\nabla_G \mathcal{L}_{\text{seg}}(S(G(\mathbf{x}_i)), \mathbf{y}_i) \\ - \lambda \nabla_G \mathcal{L}_{\text{dis}}(D(G(\mathbf{x}_i)), K+1) \Big);$$

  Adjust learning rates η_G, η_S, η_D;
return $\theta_G, \theta_S, \theta_D$;

The training procedure of our method is detailed in Algorithm 1. As in standard adversarial learning approaches, we train our model by updating the generator and discriminator in two separate steps. The discriminator is updated $n_{steps} = 3$ times to maintain near optimal solution of domain classification while G is updated less frequently. The segmentation network is then updated at the same frequency as G. We adopt a mini-batch stochastic gradient descent (SGD) technique and use Adam (Kingma and Ba, 2015) optimizer to update parameters at each step, where the gradient is estimated using a batch of m training examples, and the update step size is controlled by learning rate η_G, η_S, η_D.

Although it is possible to enforce the realism and domain invariance of normalized images via two separate discriminators, employing a single discriminator provides several important advantages. First, it avoids the problem of instability which results from training discriminators with competing losses. Instead of treating image domain and realism as unrelated properties, our model predicts them jointly within a single network. Moreover, our single discriminator model has fewer hyper-parameters to tune and is less expensive in terms of computation and mem-

ory. In addition to its higher simplicity, our adversarial model without the segmentation loss can also be shown, under mild assumptions, to lead to the desirable solution where normalized images are generated from the mean distribution of real images. This is done in the following theorem.

Theorem 1. *Let $p_r(\mathbf{x}|z)$ and $p_g(\mathbf{x}|z)$ be the probabilities that $\mathbf{x}$ is a real or a generated image, respectively, from source dataset z. The minimax optimization problem of Eq. (3) without the segmentation term corresponds to minimizing the divergence between $p_g(\mathbf{x}|z)$ for each z and the mean distribution of real images $\overline{p}_r(\mathbf{x}) = \frac{1}{k} \sum_{z=1}^K p_r(\mathbf{x}|z)$.*

PROOF. See Appendix Appendix A.

2.2. Data and performance metrics

To evaluate the performance of our method, we selected three databases with important differences in their intensity profile and subject demographics.

iSEG. The first dataset, iSEG (Wang et al., 2019), comprises 10 T1 and T2 MRI data of 6-8 month old infants acquired with a 3 Tesla scanner. The ground truth is the segmentation mask of the three main brain tissues, white matter (WM), gray matter (GM) and cerebrospinal fluids (CSF), which are critical for detecting abnormalities in brain development. Images are sampled into an isotropic $1.0\,\text{mm}^3$ resolution. This dataset is particularly challenging for segmentation since subjects are in an isointense phase where the white matter and gray matter voxel intensities greatly overlap, thus leading to a lower tissue contrast. Its images are also noisier because of the shorter scanning time used to avoid motion artifacts.

MRBrainS. The MRBrainS13 (Mendrik et al., 2015) dataset contains 5 healthy adult subjects with T1 and T2 FLAIR modalities. Images were acquired from a 3 Tesla scanner following a voxel size of $0.958\,\text{mm} \times 0.958\,\text{mm} \times 3.0\,\text{mm}$. This dataset has the same ground-truth classes as iSEG.

ABIDE. The Autism Brain Imaging Data Exchange (ABIDE I) (Di Martino et al. et al., 2014) was also used to further validate our method on an independent multi-site dataset. It comprises 1,112 images of normal and autism spectrum disorder (ASD) subjects. Images were acquired across 17 international sites, thus providing a high variance in intensity distribution. The anatomical scan parameters for each site are available on the ABIDE website[1]. A different source of variability in this dataset comes from the broad age span of its subjects, ranging from 7 to 64 years. The detailed demographics of subjects are shown in Table 1. Note that the data of 9 subjects were excluded from our study due to poor image image quality (e.g., important motion artifacts). Since this dataset lacks the manual segmentation masks, we instead considered the segmentation maps produced by the FreeSurfer *recon-all*[2] pipeline as ground-truth.

---

[1] http://fcon_1000.projects.nitrc.org/indi/abide/
[2] https://surfer.nmr.mgh.harvard.edu/fswiki/recon-all

Table 1: Demographics of subjects in the ABIDE dataset (Di Martino et al. et al., 2014).

Group	n	Male	Female	Age (mean ± stdev)
Control	539	474	65	17.01 ± 8.36
ASD	573	474	99	17.08 ± 7.72

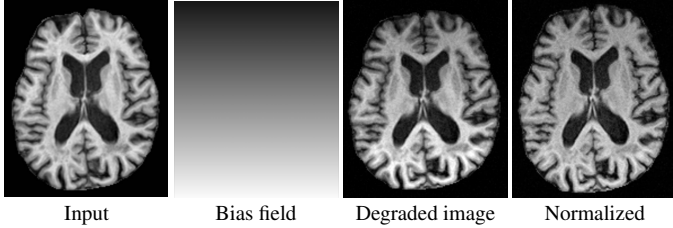| Input | Bias field | Degraded image | Normalized |

Figure 4: Normalization of an image degraded with bias field of strength $\alpha = 0.5$. **Left**: Input test image. **Middle-left**: Applied bias field. **Middle-right**: Resulting degraded test input. **Right**: Normalized image produced by the generator. One can notice the more uniform distribution of intensities in the normalized image.

For iSEG, 8 subjects were randomly selected for training, while another was kept for validation and the remaining one was kept for testing. For MRBrainS, 3 images were randomly selected for training, 1 image was kept for validation and the last one for test purposes. For ABIDE, which has more images available for learning, all 1,103 usable images were randomly split in three sets: 60% for training, 20% for validation, and 20% for testing.

We evaluate the segmentation performance of tested methods using the mean Dice Similarity Coefficient (DSC), which measures the degree of overlap between the predicted segmentation map $\mathbf{S}$ and the ground-truth $\mathbf{G}$:

$$\text{DSC}(\mathbf{S}, \mathbf{G}) = \frac{2|\mathbf{S} \cap \mathbf{G}|}{|\mathbf{S} \cup \mathbf{G}|}. \tag{4}$$

We also consider the Mean Hausdorff Distance (MHD) to measure the segmentation quality in terms of its boundary:

$$\text{MHD}(\mathbf{S}, \mathbf{G}) = \frac{1}{2}(\text{dist}(\mathbf{S}, \mathbf{G}) + \text{dist}(\mathbf{G}, \mathbf{S})). \tag{5}$$

Here, $\text{dist}(\cdot, \cdot)$ is the maximum Euclidean distance between a point in the predicted segmentation map and its nearest point in the ground-truth (or vice-versa).

2.3. Pre-processing

For all three datasets, images were cropped to brain dimensions and then padded with zero-intensity voxels to have a fixed size corresponding to the largest brain. Skull stripping was performed using the segmentation map of MRBrainS. Images were resampled to be $1.0\,\text{mm}^3$ isotropic in order to match the iSEG resolution. Due to limited GPU memory, volumes were processed (normalization and segmentation) in separate, overlapping 3D patches of size 32^3 voxels. To train the model, 40,000 patches are randomly sampled among training images of each

dataset, all centered on a voxel of a foreground class. Likewise, 12,000 patches are sampled from validation images in each dataset to evaluate the model at each epoch. The same subjects and patches were selected across all experiments through a common random seed. For testing, we process full volumes in 32^3 sub-patches extracted with a stride of 8^3 voxels. Whole-image segmentation is obtained by averaging the class probabilities of each voxel across patches containing this voxel.

To test the robustness of our model to intensity inhomogeneity, we apply a data augmentation strategy where input patches have a 50% probability of being transformed with a multiplicative bias field. Let $I(x,y,z)$ be the intensity of the input image at voxel (x,y,z), the transformed image can be defined as in (Song et al., 2017):

$$I'(x,y,z) = I(x,y,z)B(x,y,z) + \eta(x,y,z) \tag{6}$$

where $B(x,y,z)$ is the bias field at (x,y,z) and $\eta(x,y,z)$ is additive noise at the same voxel. In our experiments, we implemented a simple noiseless model where intensities are scaled linearly along the y-axis:

$$B(x,y,z) = \frac{y}{H}\alpha + (1-\alpha), \quad \alpha \in [0,1]. \tag{7}$$

In the above expression, α is the coefficient defining the slope of the bias field and H is the height of the image. While $\alpha = 0$ corresponds to having no bias field, $\alpha = 1$ scales intensities from -100% when y = 0 to $+100\%$ when y = H. An example of a transformed test image is shown in Fig. 4. Note that the bias field slope for each data augmentation is randomly selected based on a uniform distribution.

2.4. Implementation Details

Training was performed on an NVIDIA Tesla V100 32 GB GPU, with a total of 120 epochs for the dual dataset configurations (iSEG and MRBrainS) and 70 epochs when training with three datasets (iSEG, MRBrainS and ABIDE). All experiments were run to find the best ratio between the segmentation loss and discriminator loss, controlled by the hyper-parameter λ. For the main results, we selected a value of $\lambda = 1.5$ which gave a good segmentation accuracy on the validation examples, while giving plausible generated images. We used a weight decay of 0.001 on all experiments. For the generator and segmentation network, we initialized the learning rate to 0.001 and updated it with a multi-step strategy, decreasing it by a factor of 10 at epochs 50 and 75. Since the discriminator takes longer to train, a different strategy was adopted. Starting with a learning rate of 0.0001, we instead used a reduce-on-plateau strategy which applies a decay of 0.1 when the validation loss does not decrease for 7 consecutive epochs. We trained each model with an Adam optimizer. The model was implemented using the PyTorch[3] deep learning framework.

---

[3] http://pytorch.org

Table 2: Dice scores when using different training and testing sets, without any normalization. The non-adaptation of the acquisition domains seriously impacts the segmentation accuracy.

Training	Testing	Dice			
		CSF	GM	WM	Mean
iSEG	iSEG	0.920	0.857	0.828	0.868
MRBrainS	MRBrainS	0.861	0.789	0.839	0.830
iSEG	MRBrainS	0.401	0.354	0.519	0.425
MRBrainS	iSEG	0.293	0.082	0.563	0.313

3. Experiments and Results

A series of experiments is conducted to assess the benefits of the proposed adversarial normalization method. We start by showing the poor generalization performance of a baseline segmentation model without normalization, when trained and tested on different datasets. We then evaluate our method in a dual-site setting with the iSEG and MRBrainS datasets, and compare its segmentation and normalization performance against that of two competing approaches. Afterwards, we demonstrate our method in a multi-site setting by adding the large-scale ABIDE dataset and test it on multi-modal data. Last, we show its ability to correct intensity inhomogeneity and evaluate the impact of hyper-parameter λ on performance.

3.1. Cross-dataset Baseline Performance

We first establish a baseline measuring the cross-dataset performance of our segmentation network without any normalization. For this experiment, we use the same 3D U-Net architecture as in our adversarial method and train it on unnormalized T1 images from different datasets using a Dice loss. Four different scenarios are tested: 1) training and testing on iSEG data only; 2) training and testing on MRBrainS data only; 3) training on iSEG and testing on MRBrainS; 4) training on MRBrainS and testing on iSEG. The last two scenarios evaluate the segmentation model's ability to generalize across datasets with different characteristics, including subject demographics and acquisition protocol.

Results of this experiment are shown in Table 2. As can be seen, the performance decreases considerably when testing on a different dataset from the one used for training. Thus, the mean Dice score of the model trained on iSEG drops by 51.0% when tested on MRBrainS. Likewise, we observe a 62.3% drop in mean Dice when testing on iSEG the model trained on MRBrainS. These results demonstrate the high sensitivity of deep learning segmentation models to the training data, thus validating the need for a data-driven normalization method.

3.2. Dual-site Evaluation

Next, we compare our adversarial normalization method against a commonly-used standardization technique (Birenbaum and Greenspan, 2017) and the learned normalization approach of Drozdzal et al. (2018), on images from two different datasets: iSEG and MRBrainS. As mentioned earlier, these datasets have distinct characteristics, iSEG containing T1 MRIs of infants in
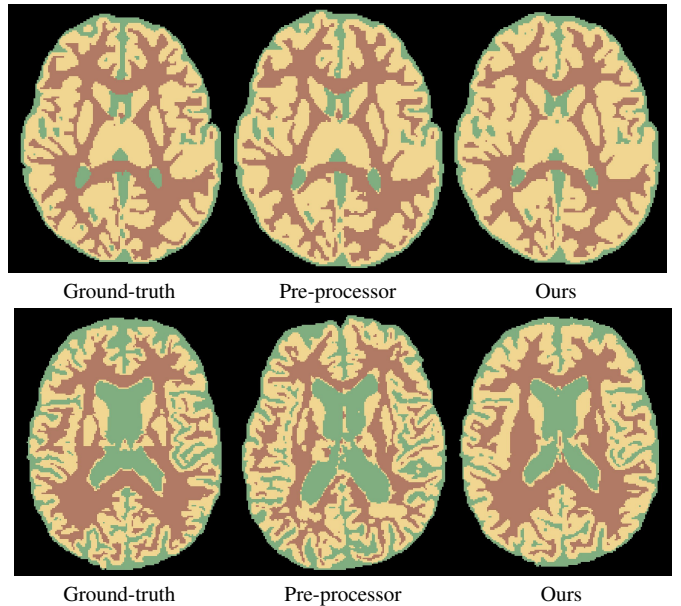

Ground-truth   Pre-processor   Ours

Ground-truth   Pre-processor   Ours

Figure 5: Visualization of predicted segmentation masks for two test images. **Top row**: iSEG image. **Bottom row**: MRBrainS image.

the isointense stage and MRBrainS the T1 MRIs of adult subjects. Therefore, standard per-image normalization may not be effective. The standardization technique tested in this experiment normalizes the intensity of each voxel in a given volume by subtracting from it the volume mean and dividing it by the standard deviation. The learned normalization approach, which we call *Pre-processor* in the results, contains the same pipeline as our method (i.e., generator and segmentation networks), but without its discriminator. This baseline is used to asses the contribution of the adversarial learning in obtaining realistic normalized images. As this segmentation-optimized approach does not impose any constraint on realism, it should be an "upper-bound" on the segmentation accuracy achievable by our method.

3.2.1. Segmentation Performance

The Dice scores obtained by the Standardization technique, the Pre-processor without realism constraints and our adversarial normalization method are reported in Table 3 for the iSEG and MRBrainS datasets (dual-site setting). As expected, both our method and the learned pre-processor yield a large gain in performance over the fixed standardization technique, with a mean Dice score (DSC in %) and mean Hausdorff distance (MHD in millimeters) of 86.7% / 0.444 mm and 87.0% / 0.482 mm respectively, compared to 84.1% / 0.675 mm for a conventional Standardization. Surprisingly, our method achieves a performance on par with the learned Pre-processor, obtaining a slightly lower mean Dice score but slightly better mean Hausdorff distance. This indicates that imposing realism constraints, while also considering the downstream segmentation task, does not impact the segmentation accuracy. We also note that the performance of our joint normalization method, when trained with both iSEG and MRBrainS datasets, is greater than that of the segmentation network when trained and tested independently

Table 3: Comparative Dice scores of different model architectures and data. The proposed method yields an important performance improvement over training and testing on single-domain or on standardized inputs.

| Setting	Method	CSF		GM		WM		Mean	
		DSC	MHD	DSC	MHD	DSC	MHD	DSC	MHD
Dual-site (iSEG + MRBrainS)	Standardization	0.897	0.792	0.836	0.498	0.790	0.734	0.841	0.675
	Pre-processor	**0.919**	**0.227**	**0.860**	0.517	0.831	0.702	**0.870**	0.482
	Ours	0.912	0.245	0.853	**0.492**	**0.836**	**0.595**	0.867	**0.444**
Multi-site (iSEG + MRBrainS + ABIDE)	Standardization	0.860	0.264	0.881	0.684	0.856	0.812	0.866	0.587
	Pre-processor	**0.922**	**0.251**	**0.895**	**0.392**	0.870	**0.530**	**0.896**	**0.390**
	Ours	0.913	0.293	0.887	0.422	**0.870**	0.598	0.890	0.438
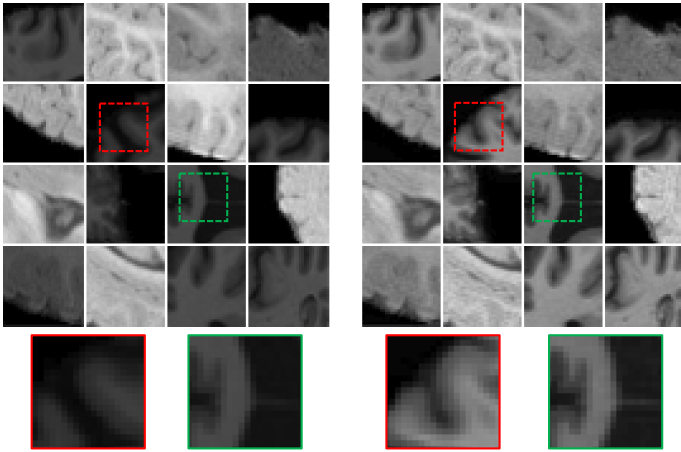
Figure 6: Mixed iSEG and MRBrainS inputs (**left**) and the generated images with adversarial normalization (**right**). Notice the improved homogeneity of intensities in the normalized images and enhanced tissue contrast (see zoomed regions).

on these datasets (Table 2). These results demonstrate the benefit of exploiting jointly-normalized datasets to improve the overall performance. Examples of predicted segmentation for iSEG and MRBrainS test images are shown in Fig. 5. We see that our method yields a segmentation mask close to the ground-truth for the two datasets and for all brain tissue classes.
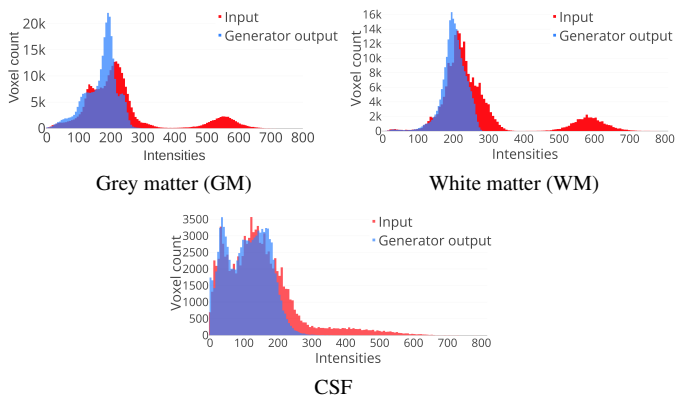

Figure 7: Histograms of tests images with constraint of realism. One can notice the intensities have been modified and follow a more Gaussian-like curve with overlapping intensities for each classes.

Table 4: Jensen-Shannon Divergence (JSD) of input and normalized images from the generator. A lower value corresponds to more similar distributions. Dataset normalization performance is on par with the method in Drozdzal et al. (2018) while still offering visually realistic images.

Datasets	Input data	Pre-processor	Ours
iSEG + MRBrainS	2.0839	0.2793	0.2788
iSEG + MRBrainS + ABIDE	12.2212	0.4185	0.4180

3.2.2. Normalization Performance

The advantage of our method in generating realistic images is illustrated in Fig. 6, where randomly-selected patches from input images and the corresponding output of our generator are shown. Normalized patches present a more uniform intensity distribution, while also preserving the realism and details of the original patches. Our method also exhibits a better contrast enhancement in the generated images, as seen in the zoomed regions of the figure. This is enabled by our task-driven approach which also minimizes the segmentation loss, thereby increasing the contrast along region boundaries.

The normalization effect of our method can be further appreciated from Fig. 7, which shows the histogram of intensities for input images and the output of our generator (combined iSEG and MRBrainS images). While input images have a broader spread in distributions with distinct modes, the normalized images have a narrower distribution more centered around a single mode. This helps reducing the intra-class variance and, therefore, increases the segmentation accuracy.

3.3. Multi-site Evaluation

In the next experiment, we evaluate our adversarial normalization method in a multi-site scenario involving a third dataset, ABIDE. This large-scale dataset contains images obtained from 17 sites with different acquisition protocols. The segmentation performance of the Standardization technique, the learned Pre-processor approach of Drozdzal et al. (2018) and our method, when trained with all three datasets, is reported in Table 3 (multi-site setting). Once again, both the Pre-processor and our method achieve a considerable improvement compared to employing a fixed standardization technique. Specifically, our method yields a 2.40% improvement in mean Dice over this technique. Its performance is also similar to the segmentation-optimized Pre-processor, which does not preserve realism.

Table 5: Dice score of our normalization method when using only T1 images or T1 + T2 images as input. The higher performance for T1 + T2 inputs demonstrates the capability of our method to process multi-modal data.

Modality	CSF		GM		WM		Mean	
	DSC	MHD	DSC	MHD	DSC	MHD	DSC	MHD
T1 only	**0.912**	0.245	0.853	0.492	0.836	0.595	0.867	0.444
T1 + T2	0.910	**0.165**	**0.889**	**0.431**	**0.862**	**0.493**	**0.887**	**0.363**

Table 6: Pearson correlation coefficient (ρ) between voxel intensity and y-axis position in an axial slice of an MRBrainS and iSEG test subject before and after normalization. A higher correlation implies a stronger bias field degradation. Note that the proposed method reduces correlation in all cases.

α	MRBrainS		iSEG	
	Degraded	Normalized	Degraded	Normalized
Original	−0.0199	–	0.0271	–
0.3	0.0982	0.0534	0.2560	0.1192
0.5	0.1963	0.1422	0.4264	0.2430
0.7	0.3103	0.2393	0.5914	0.4599
0.9	0.4368	0.3649	0.7316	0.5928

Table 2 gives the Jensen-Shannon Divergence (JSD) between the intensity distributions of images from different datasets. Lower values correspond to more similar distributions between input images and the generator output. We see an important decrease in JSD when using normalized images. This illustrates the normalization effect of both our method and the learned Pre-processor. Comparing the two approaches, our method leads to slightly smaller JSD values, suggesting a more uniform distribution of intensities across different datasets. Note that this metric does not evaluate realism of images, which is the main advantage of our method.

3.4. Multi-modal Testing

Recent approaches that use multiple image modalities can increase the accuracy of learned tasks (Dolz et al., 2019). For instance, T1 images typically provide a higher contrast between gray and white matter tissues, while T2 images offer a better contrast between brain tissue and CSF. Combining both T1 and T2 is, therefore, expected to improve brain segmentation. We assess the improvement in segmentation accuracy of our normalization method when used with multi-modal images.

Table 5 compares the Dice scores obtained by our method when using T1 images only as input or both pre-aligned T1 and T2 images as inputs. The accuracy is measured as higher when employing the two MRI sequences, with a mean Dice score improvement of 0.020 and an average Hausdorff Distance improvement of 0.081 mm.

3.5. Robustness to Image Degradation

Our task-driven normalization method also demonstrates the ability to enhance images with non-homogeneous intensity. To evaluate this, we trained the method with 50% of input images augmented with a random bias field as described in Section 2.3. Since the discriminator must discern the generator output for these degraded images from the real, non-degraded images, it
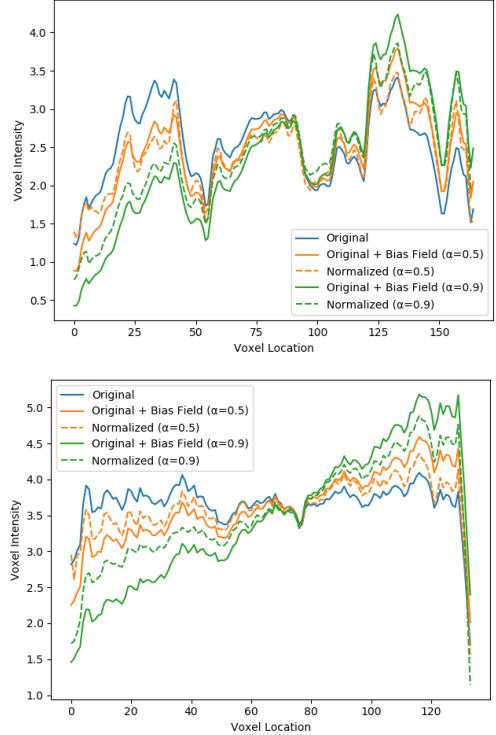

Figure 8: Mean voxel intensity in an axial slice of an MRBrainS (left) and iSEG (right) test subject across the brain for bias field strength of $\alpha = 0.5$ and $\alpha = 0.9$. By bringing intensity values closer to those of the original image, our normalization method helps correct the bias field in degraded images.

encourages the generator to remove the bias field while also preserving realism.

We measure our method's ability to correct the bias field by computing the Pearson correlation between the intensity and position of voxels along the field's direction (i.e., y-axis). Since we used a linear bias field, a higher correlation corresponds to a stronger degradation of the image. Table 6 gives the mean correlation for test images of the MRBrainS and iSEG datasets, the same images degraded with a bias field of increasing strength $\alpha \in \{0.3, 0.5, 0.7, 0.9\}$, and the output of our generator for these images. The mean intensity as function of the y-axis position, for $\alpha = 0.5$ and $\alpha = 0.9$, is shown in Fig. 8. We observe that, for both datasets, intensity is noticeably less correlated to the y-axis position in normalized images than in degraded images, illustrating our method's ability to correct inhomogeneity. The benefit of our method is particularly important for stronger bias fields (i.e., $\alpha = 0.7$ and $\alpha = 0.7$). An example of a test image with bias field of strength $\alpha = 0.5$ and corresponding normalized output of the generator is provided in Figure 4. One can see that intensities in the normalized image are more uniform.

3.6. Impact of Hyper-parameter Lambda

As defined in Eq. (3), hyper-parameter λ has a direct impact on the level of realism in generated images. A lower value emphasizes segmentation accuracy, while a higher value prioritizes the generation of more realistic, domain-invariant images. In this last experiment, we analyse the impact of this crucial hyper-parameter on image segmentation and normalization.
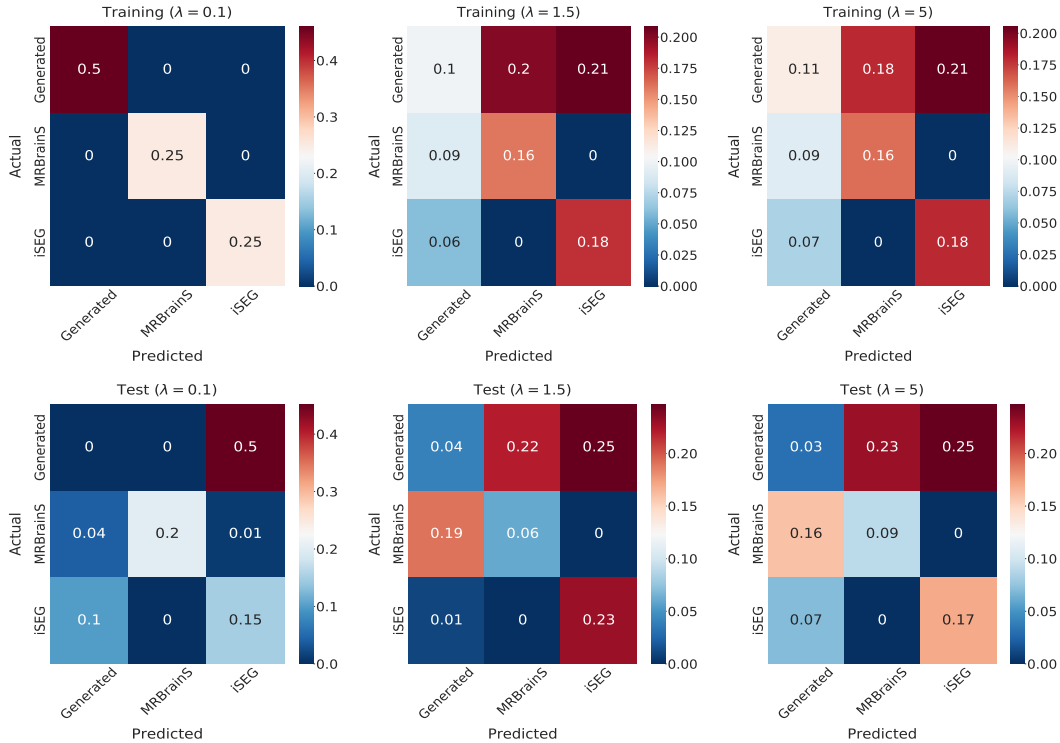
Figure 9: Normalized confusion matrix of the domain classifier (discriminator) for different values of hyper-parameter λ. **Top row**: training examples. **Bottom row**: test examples. For $\lambda = 1$ generated training images are easily differentiated from real ones. In contrast, when using $\lambda \geq 1.5$, the generator confuses generated images for real ones, with a uniform distribution between the MRBrainS and iSEG classes.

Table 7: Mean Dice score for different values of hyper-parameter λ. A lower λ emphasizes segmentation performance while a higher value gives more importance to the realism of the generated image.

| Lambda (λ) | Mean DSC | Discr. accuracy (%) | |
		Train	Test
0.1	0.875	98.95	34.49
1.5	0.866	44.10	32.58
5.0	0.851	44.07	28.58

Table 7 gives the mean Dice scores and discriminator accuracy for our training and test samples with $\lambda \in \{0.1, 1.5, 5.0\}$. As expected, the segmentation accuracy decreases with higher λ values, since the model focuses more on generating realistic images and less on obtaining a precise segmentation. Hence, we observe a 2.4% drop in mean Dice scores when increasing from $\lambda = 0.1$ to $\lambda = 5.0$. The realism of the generated images can be estimated with the discriminator accuracy. A higher value indicates that the generated images can be more easily differentiated from real ones. For a small λ value of 0.1, the discriminator can classify almost perfectly all our training examples, indicating that images produced by the generator are indeed very different from real images. As λ is increased, generated images become more similar to real ones, resulting in a lower discriminator accuracy both in training and testing samples.

The behavior of our discriminator for the λ values is further analyzed in Fig. 9 which shows the normalized confusion matrix[4] for training and testing samples. The normalized training images are correctly identified as generated with $\lambda = 0.1$, but uniformly predicted as MRBrainS or iSEG with $\lambda = 1.5$ or $\lambda = 5.0$. This corresponds to the scenario expected from Theorem 2, where generated images from different datasets follow the same distribution and the discriminator predicts domain classes uniformly. We also note that the discriminator overfits the training data with $\lambda = 0.1$, resulting in a poor classification of the generated samples during testing.

4. Conclusions

This paper presents a realistic task-and-data-driven normalization method that improves the segmentation of images by exploiting multiple datasets simultaneously. Our method leverages an adversarial learning strategy that involves three networks: a generator which normalizes input images while preserving their realism, a task network that predicts an accurate segmentation from normalized images, and a discriminator which classifies the domain of these images. Unlike traditional adversarial approaches for image synthesis or domain adaptation, where a discriminator distinguishes real images from fake ones or predicts the domain of images, our discriminator is trained in a (K+1)-class classification problem with K domain classes corresponding to the originating dataset (or site) of real images,

---

[4]The matrix is normalized by dividing each value by the total number of samples.

and an additional class corresponding to the generated normalized images. By maximizing the discriminator loss and simultaneously minimizing the segmentation loss during training, the generator consequently learns to produce images that are both harmonized and realistic across all datasets, while still optimizing for segmentation. Compared to the recent data harmonization techniques (Modanwal et al., 2020), our method has less hyper-parameters to tune and can more easily adapt to the addition of new datasets or image modalities.

The advantage of our method has been demonstrated in a comprehensive set of experiments involving three largely different brain MRI datasets: iSEG, MRBrainS and ABIDE. In an experiment with iSEG and MRBrainS images, we first established that a standard segmentation network performs poorly when trained and tested across different datasets (Table 2). In contrast, our adversarial normalization method achieves a consistently higher accuracy using images with dissimilar intensity distributions (Table 3 and Fig. 5). Moreover, our normalization network also provides clinically interpretable images when compared to state-of-the-art approaches (Drozdzal et al., 2018) (Fig. 6).

Our experiments also showed our normalization strategy to yield good performance on data from more than two sites or with multiple image modalities. While trained on all three datasets, our method achieved a higher mean Dice score of 0.890 compared to 0.867 when employing only iSEG and MR-BrainS (Table 3). In addition, it considerably reduced the variability of intensities in normalized images from the different datasets with a JSD of 0.418 compared to 12.221 in the case of unnormalized images (Table 4). Furthermore, when employing an additional image modality as input (T2-weighted MRI), our method also obtained a better segmentation accuracy, with a mean Dice improvement of 0.02 with respect to using only T1-weighted images (Table 5).

We further evaluated our method by analyzing its robustness to image degradation and its sensitivity to hyper-parameters. In addition to harmonizing images from different sites, our adversarial normalization method could also remove intensity inhomogeneity without requiring additional processing (Table 6). This could help achieve a much faster and more reliable analysis of large-scale brain MRI datasets compared to traditional processing pipelines (Table 7 and Fig. 9). Training our model with different values of hyper-parameter λ enabled us to study the trade-off between segmentation accuracy and normalized image realism. Best results were found with $\lambda = 1.5$.

Our task-driven adversarial normalization approach unlocks the training of deep learning models with data from multiple sites by improving both realism and accuracy of normalized images across datasets. A potential technical limitation of the proposed method is the need to process 3D images in smaller subregions due to the current limitation in GPU memory. Although we obtained spatially-smooth generated images and segmentation maps, the overall result may be sub-optimal since the global context and intensity distribution of images is not fully considered when working with local patches. In future work, we plan to tackle this problem by incorporating global image statistics in the loss function, and by exploring a 2.5D approach (Xue

et al., 2020) where slices from different view planes are processed simultaneously. Moreover, while the discriminator architecture employed in our model generally led to good results, it can sometimes suffer from vanishing gradient or mode collapse. Our future work will also investigate adversarial models such as the Wasserstein GAN (Arjovsky et al., 2017) which are less prone to these problems. Finally, we aim to demonstrate our method in a broader set of applications, including the segmentation of brain lesions where preserving fine regions during normalization is critical.

The source code of this article is publicly available on Github[5].

Appendix A. Proof of Theorem 1

Theorem 2. *Let $p_r(\mathbf{x}|z)$ and $p_g(\mathbf{x}|z)$ be the probabilities that $\mathbf{x}$ is a real or a generated image, respectively, from source dataset z. The minimax optimization problem of Eq. (3) without the segmentation term corresponds to minimizing the divergence between $p_g(\mathbf{x}|z)$ for each z and the mean distribution of real images $\overline{p}_r(\mathbf{x}) = \frac{1}{k}\sum_{z=1}^{K} p_r(\mathbf{x}|z)$.*

PROOF. If we ignore the segmentation loss $\mathcal{L}_{\mathrm{seg}}$, the optimization problem is given by

$$\min_G \max_D \mathcal{L}(G,D) = \mathbb{E}_{\mathbf{x},z}\Big[\log D_z(\mathbf{x}) + \log D_{K+1}(G(\mathbf{x}))\Big]$$

$$= \mathbb{E}_{\mathbf{x},z}\Big[\log D_z(\mathbf{x}) + \log\Big(1 - \sum_{z'=1}^{K} D_{z'}(G(\mathbf{x}))\Big)\Big] \tag{A.1}$$

Suppose generator G is fixed, the optimal discriminator D^* is found by minimizing

$$\mathcal{L}_G(D) = -\sum_{z=1}^{K} p(z) \int_{\mathbf{x}} \Big[p_r(\mathbf{x}|z)\log D_z(\mathbf{x}) + p_g(\mathbf{x}|z)\log\Big(1 - \sum_{z'=1}^{K} D_{z'}(\mathbf{x})\Big)\Big] d\mathbf{x}. \tag{A.2}$$

We obtain the optimum value for each $\mathbf{x}$ by deriving this function with respect to $D_z(\mathbf{x})$

$$\frac{\partial \mathcal{L}_D}{\partial D_z(\mathbf{x})} = -\frac{p(z)\,p_r(\mathbf{x}|z)}{D_z(\mathbf{x})} + \frac{p(z)\,p_g(\mathbf{x}|z)}{1 - \sum_{z'=1}^{K} D_{z'}(\mathbf{x})}. \tag{A.3}$$

Setting this to zero yields

$$\frac{D_z^*(\mathbf{x})}{1 - \sum_{z'=1}^{K} D_{z'}^*(\mathbf{x})} = \frac{p_r(\mathbf{x}|z)}{p_g(\mathbf{x}|z)}. \tag{A.4}$$

---
[5] https://github.com/pldelisle/deepNormalize

Summing both sides of the equation over z, and using $D_{K+1}(\mathbf{x}) = 1 - \sum_z D_z(\mathbf{x})$, we then get

$$\frac{\sum_{z=1}^{K} D_z^*(\mathbf{x})}{1 - \sum_{z'=1}^{K} D_{z'}^*(\mathbf{x})} = \sum_{z=1}^{K} \frac{p_r(\mathbf{x}|z)}{p_g(\mathbf{x}|z)} = \frac{1 - D_{K+1}^*(\mathbf{x})}{D_{K+1}^*(\mathbf{x})} \qquad (A.5)$$

and therefore

$$D_z^*(\mathbf{x}) = \frac{\frac{p_r(\mathbf{x}|z)}{p_g(\mathbf{x}|z)}}{1 + \sum_{z'=1}^{K} \frac{p_r(\mathbf{x}|z')}{p_g(\mathbf{x}|z')}}, \quad z = 1, \dots, K; \qquad (A.6)$$

$$D_{K+1}^*(\mathbf{x}) = \frac{1}{1 + \sum_{z'=1}^{K} \frac{p_r(\mathbf{x}|z')}{p_g(\mathbf{x}|z')}} \qquad (A.7)$$

Next, we use this result to find the optimal generator. Toward this goal, we plug (A.7) into the loss of Eq. (A.1) and minimize

$$\begin{aligned}
\mathcal{L}_{D^*}(G) &= \sum_{z=1}^{K} p(z) \int_{\mathbf{x}} p_g(\mathbf{x}|z) \log D_{K+1}^*(\mathbf{x}) \, d\mathbf{x} \\
&= \sum_{z=1}^{K} p(z) \int_{\mathbf{x}} p_g(\mathbf{x}|z) \log\left(\frac{1}{1 + \sum_{z'=1}^{K} \frac{p_r(\mathbf{x}|z')}{p_g(\mathbf{x}|z')}} \right) d\mathbf{x} \\
&= \sum_{z=1}^{K} p(z) \int_{\mathbf{x}} p_g(\mathbf{x}|z) \log\left(\frac{p_g(\mathbf{x}|z)}{p_g(\mathbf{x}|z) + \sum_{z'=1}^{K} \frac{p_g(\mathbf{x}|z)}{p_g(\mathbf{x}|z')} p_r(\mathbf{x}|z')} \right) d\mathbf{x}
\end{aligned}$$
$$(A.8)$$

Let $q(\mathbf{x}|z)$ the probability distribution defined as

$$q(\mathbf{x}|z) = \frac{1}{\mathcal{Z}(z)} \Big[p_g(\mathbf{x}|z) + \sum_{z'=1}^{K} \frac{p_g(\mathbf{x}|z)}{p_g(\mathbf{x}|z')} p_r(\mathbf{x}|z') \Big], \qquad (A.9)$$

where $\mathcal{Z}(z)$ is a normalization constant. Suppose that the generator produces images similar to the input, i.e. $p_g(\mathbf{x}|z') \approx p_r(\mathbf{x}|z')$, this constant can be estimated as

$$\mathcal{Z}(z) = \int_{\mathbf{x}} p_g(\mathbf{x}|z) d\mathbf{x} + \sum_{z'=1}^{K} \int_x \frac{p_g(\mathbf{x}|z)}{p_g(\mathbf{x}|z')} p_r(\mathbf{x}|z') d\mathbf{x} \approx K+1.$$
$$(A.10)$$

Hence, we can rewrite the generator loss in (A.8) as

$$\begin{aligned}
\mathcal{L}_{D^*}(G) &= \sum_{z=1}^{K} p(z) \int_{\mathbf{x}} p_g(\mathbf{x}|z) \log\left(\frac{p_g(\mathbf{x}|z)}{(K+1)q(\mathbf{x}|z)} \right) d\mathbf{x} \\
&= \sum_{z=1}^{K} p(z) D_{\mathrm{KL}}(p_g(\cdot|z) \,\|\, q(\cdot|z)) - \log(K+1),
\end{aligned}$$
$$(A.11)$$

where $D_{\mathrm{KL}}(p \,\|\, q) \geq 0$ is the KL divergence. Assuming that $p(z)$ is uniform, the optimal generator G^* is therefore such that $p_g(\cdot|z) = q(\cdot|z)$, for $z = 1, \dots, K$. Considering Eq. (A.9), this can only be achieved if $p_g(\cdot|z) = p_g(\cdot|z')$, $\forall z, z'$. This in turn gives

$$q(\mathbf{x}|z) = \frac{1}{K+1} \Big[p_g(\mathbf{x}|z) + \sum_{z'=1}^{K} p_r(\mathbf{x}|z') \Big]. \qquad (A.12)$$

Using the fact that $p_g(\mathbf{x}|z) = q(\mathbf{x}|z)$, we finally obtain

$$p_g(\mathbf{x}|z) = \frac{1}{K} \sum_{z'=1}^{K} p_r(\mathbf{x}|z') = \bar{p}_r(\mathbf{x}). \qquad (A.13)$$

Hence, G^* will produce outputs from the average of real image distributions over the different data sources.

References

Arjovsky, M., Chintala, S., Bottou, L., 2017. Wasserstein generative adversarial networks, in: International Converence on Machine Learning (ICML).

Baid, U., Mahajan, A., Talbar, S., SiddheshThakur, S.R., Moiyadi, A., Thakur, M., Gupta, S., 2018. GBM Segmentation with 3D U-Net and Survival Prediction with Radiomics, in: MICCAI BraTS Challenge.

Birenbaum, A., Greenspan, H., 2017. Multi-view longitudinal CNN for multiple sclerosis lesion segmentation. Engineering Applications of Artificial Intelligence 65, 111 – 118.

Carass, A., Roy, S., Gherman, A., Reinhold, J.C., Jesson, A., Arbel, T., Maier, O., Handels, H., Ghafoorian, M., Platel, B., Birenbaum, A., Greenspan, H., Pham, D.L., Crainiceanu, C.M., Calabresi, P.A., Prince, J.L., Roncal, W.R.G., Shinohara, R.T., Oguz, I., 2020. Evaluating White Matter Lesion Segmentations with Refined Sørensen-Dice Analysis. Scientific Reports 10, 8242.

Casamitjana, A., Puch, S., Aduriz, A., Vilaplana, V., 2017. 3D Convolutional Neural Networks for Brain Tumor Segmentation: A Comparison of Multiresolution Architectures, in: MICCAI Workshop on Brain Lesion.

Chen, H., Dou, Q., Yu, L., Qin, J., Heng, P.A., 2018. VoxResNet: Deep voxelwise residual networks for brain segmentation from 3D MR images. NeuroImage 170, 446–455.

Cheplygina, V., de Bruijne, M., Pluim, J.P.W., 2019. Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. Medical Image Analysis .

Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O., 2016. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation, in: Medical Image Computing and Computer-Assisted Intervention (MICCAI).

Ciga, O., Chen, J., Martel, A., 2019. Multi-layer domain adaptation for deep convolutional networks, in: MICCAI Workshop on DART.

Dewey, B.E., Zhao, C., Reinhold, J.C., Carass, A., Fitzgerald, K.C., Sotirchos, E.S., Saidha, S., Oh, J., Pham, D.L., Calabresi, P.A., van Zijl, P.C., Prince, J.L., 2019. Deepharmony: A deep learning approach to contrast harmonization across scanner changes. Magnetic Resonance Imaging 64, 160 – 170.

Di Martino et al., A., Yan, C.G., Li, Q., Denio, E., Castellanos, F.X., Alaerts, K., Anderson, J.S., Assaf, M., Bookheimer, S.Y., Dapretto, M., Deen, B., Delmonte, S., Dinstein, I., Ertl-Wagner, B., Fair, D.A., Gallagher, L., Kennedy, D.P., Keown, C.L., Keysers, C., Lainhart, J.E., Lord, C., Luna, B., Menon, V., Minshew, N.J., Monk, C.S., Mueller, S., Müller, R.A., Nebel, M.B., Nigg, J.T., O'Hearn, K., Pelphrey, K.A., Peltier, S.J., Rudie, J.D., Sunaert, S., Thioux, M., Tyszka, J.M., Uddin, L.Q., Verhoeven, J.S., Wenderoth, N., Wiggins, J.L., Mostofsky, S.H., Milham, M.P., 2014. The autism brain imaging data exchange: Towards a large-scale evaluation of the intrinsic brain architecture in autism. Molecular Psychiatry 19, 659–667.

Dolz, J., Gopinath, K., Yuan, J., Lombaert, H., Desrosiers, C., Ben Ayed, I., 2019. HyperDense-Net: A Hyper-Densely Connected CNN for Multi-Modal Image Segmentation. IEEE Transactions on Medical Imaging 38, 1116–1126.

Drozdzal, M., Chartrand, G., Vorontsov, E., Shakeri, M., Di Jorio, L., Tang, A., Romero, A., Bengio, Y., Pal, C., Kadoury, S., 2018. Learning normalized inputs for iterative estimation in medical image segmentation. Medical Image Analysis .

Feng, X., Tustison, N.J., Patel, S.H., Meyer, C.H., 2019. Brain tumor segmentation using an ensemble of 3d u-nets and overall survival prediction using radiomic features. Frontiers of Computational Neuroscience 14.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets, in: Advances in Neural Information Processing Systems (NeurIPS).

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

Hesse, L.S., Kuling, G., Veta, M., Martel, A., 2020. Intensity augmentation to improve generalizability of breast segmentation across different MRI scan protocols. IEEE Transactions on Biomedical Engineering , 1–1.

Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B., 2017. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Medical Image Analysis 36, 61–78.

Kingma, D.P., Ba, J., 2015. Adam: A method for stochastic optimization, in: International Conference on Learning Representations (ICLR).

Kochunov, P., Jahanshad, N., Sprooten, E., Nichols, T.E., Mandl, R.C., Almasy, L., Booth, T., Brouwer, R.M., Curran, J.E., de Zubicaray, G.I., Dimitrova, R., Duggirala, R., Fox, P.T., Elliot Hong, L., Landman, B.A., Lemaitre, H., Lopez, L.M., Martin, N.G., McMahon, K.L., Mitchell, B.D., Olvera, R.L., Peterson, C.P., Starr, J.M., Sussmann, J.E., Toga, A.W., Wardlaw, J.M., Wright, M.J., Wright, S.N., Bastin, M.E., McIntosh, A.M., Boomsma, D.I., Kahn, R.S., den Braber, A., de Geus, E.J., Deary, I.J., Hulshoff Pol, H.E., Williamson, D.E., Blangero, J., van 't Ent, D., Thompson, P.M., Glahn, D.C., 2014. Multi-site study of additive genetic effects on fractional anisotropy of cerebral white matter: Comparing meta and megaanalytical approaches for data pooling. NeuroImage 95, 136–150.

Logue, M.W., van Rooij, S.J., Dennis, E.L., Davis, S.L., Hayes, J.P., Stevens, J.S., Densmore, M., Haswell, C.C., Ipser, J., Koch, S.B., Korgaonkar, M., Lebois, L.A., Peverill, M., Baker, J.T., Boedhoe, P.S., Frijling, J.L., Gruber, S.A., Harpaz-Rotem, I., Jahanshad, N., Koopowitz, S., Levy, I., Nawijn, L., O'Connor, L., Olff, M., Salat, D.H., Sheridan, M.A., Spielberg, J.M., van Zuiden, M., Winternitz, S.R., Wolff, J.D., Wolf, E.J., Wang, X., Wrocklage, K., Abdallah, C.G., Bryant, R.A., Geuze, E., Jovanovic, T., Kaufman, M.L., King, A.P., Krystal, J.H., Lagopoulos, J., Bennett, M., Lanius, R., Liberzon, I., McGlinchey, R.E., McLaughlin, K.A., Milberg, W.P., Miller, M.W., Ressler, K.J., Veltman, D.J., Stein, D.J., Thomaes, K., Thompson, P.M., Morey, R.A., 2018. Smaller hippocampal volume in posttraumatic stress disorder: A multisite enigma-pgc study: Subcortical volumetry results from posttraumatic stress disorder consortia. Biological Psychiatry 83, 244 – 253.

Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

Mendrik, A., Vincken, K., Kuijf, H., Breeuwer, M., Bouvy, W., de Bresser, J., Alansary, A., de Bruijne, M., Carass, A., El-Baz, A., Jog, A., Katyal, R., Khan, A., Lijn, F., Mahmood, Q., Mukherjee, R., Opbroek, A., Paneri, S., Pereira, S., Viergever, M., 2015. MRBrainS Challenge: Online Evaluation Framework for Brain Image Segmentation in 3T MRI Scans. Computational Intelligence and Neuroscience .

Milletari, F., Navab, N., Ahmadi, S., 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation, in: International Conference on 3D Vision (3DV).

Modanwal, G., Vellal, A., Buda, M., Mazurowski, M.A., 2020. MRI image harmonization using cycle-consistent generative adversarial network, in: SPIE Medical Imaging, SPIE.

Mårtensson, G., Ferreira, D., Granberg, T., Cavallin, L., Oppedal, K., Padovani, A., Rektorova, I., Bonanni, L., Pardini, M., Kramberger, M.G., Taylor, J.P., Hort, J., Snædal, J., Kulisevsky, J., Blanc, F., Antonini, A., Mecocci, P., Vellas, B., Tsolaki, M., Kłoszewska, I., Soininen, H., Lovestone, S., Simmons, A., Aarsland, D., Westman, E., 2020. The reliability of a deep learning model in clinical out-of-distribution MRI data: a multicohort study. Medical Image Analysis .

Nie, D., Trullo, R., Lian, J., Wang, L., Petitjean, C., Ruan, S., Wang, Q., Shen, D., 2018. Medical image synthesis with deep convolutional adversarial networks. IEEE Transactions on Biomedical Engineering 65, 2720–2730.

Oguz, I., Malone, J.D., Atay, Y., Tao, Y.K., 2020. Self-fusion for OCT noise reduction, in: SPIE Medical Imaging.

Onofrey, J.A., Casetti-Dinescu, D.I., Lauritzen, A.D., Sarkar, S., Venkataraman, R., Fan, R.E., Sonn, G.A., Sprenkle, P.C., Staib, L.H., Papademetris, X., 2019. Generalizable multi-site training and testing of deep neural networks using image normalization, pp. 348–351.

van Opbroek, A., Ikram, M.A., Vernooij, M.W., de Bruijne, M., 2015. Transfer learning improves supervised image segmentation across imaging protocols. IEEE transactions on medical imaging 34, 1018–1030.

Pomponio, R., Erus, G., Habes, M., Doshi, J., Srinivasan, D., Mamourian, E., Bashyam, V., Nasrallah, I.M., Satterthwaite, T.D., Fan, Y., Launer, L.J., Masters, C.L., Maruff, P., Zhuo, C., Völzke, H., Johnson, S.C., Fripp, J., Koutsouleris, N., Wolf, D.H., Gur, R., Gur, R., Morris, J., Albert, M.S., Grabe, H.J., Resnick, S.M., Bryan, R.N., Wolk, D.A., Shinohara, R.T., Shou, H., Davatzikos, C., 2020. Harmonization of large MRI datasets for the analysis of brain imaging patterns throughout the lifespan. NeuroImage 208.

Radford, A., Metz, L., Chintala, S., 2016. Unsupervised representation learning with deep convolutional generative adversarial networks, in: International Conference on Learning Representations (ICLR).

Shinohara, R.T., Oh, J., Nair, G., Calabresi, P.A., Davatzikos, C., Doshi, J., Henry, R.G., Kim, G., Linn, K.A., Papinutto, N., Pelletier, D., Pham, D.L., Reich, D.S., Rooney, W., Roy, S., Stern, W., Tummala, S., Yousuf, F., Zhu, A., Sicotte, N.L., Bakshi, R., 2017. Volumetric analysis from a harmonized multisite brain MRI study of a single subject with multiple sclerosis. American Journal of Neuroradiology 38, 1501–1509.

Song, S., Zheng, Y., He, Y., 2017. A review of methods for bias correction in medical images. Biomedical Engineering Review 1.

Wang, L., Nie, D., Li, G., Puybareau, É., Dolz, J., Zhang, Q., Wang, F., Xia, J., Wu, Z., Chen, J., Thung, K., Bui, T.D., Shin, J., Zeng, G., Zheng, G., Fonov, V.S., Doyle, A., Xu, Y., Moeskops, P., Pluim, J.P.W., Desrosiers, C., Ayed, I.B., Sanroma, G., Benkarim, O.M., Casamitjana, A., Vilaplana, V., Lin, W., Li, G., Shen, D., 2019. Benchmark on Automatic Six-Month-Old Infant Brain Segmentation Algorithms: The iSeg-2017 Challenge. IEEE Transactions on Medical Imaging .

Xue, Y., Farhat, F.G., Boukrina, O., Barrett, A., Binder, J.R., Roshan, U.W., Graves, W.W., 2020. A multi-path 2.5 dimensional convolutional neural network system for segmenting stroke lesions in brain MRI images. NeuroImage: Clinical 25.

13