

ADVERSARIAL NORMALIZATION FOR MULTI DOMAIN IMAGE SEGMENTATION

Pierre-Luc Delisle, Benoit Anctil-Robitaille, Christian Desrosiers, Herve Lombaert

ETS Montreal, Canada

ABSTRACT

Image normalization is a critical step in medical imaging. This step is often done on a per-dataset basis, preventing current segmentation algorithms from the full potential of exploiting jointly normalized information across multiple datasets. To solve this problem, we propose an adversarial normalization approach for image segmentation which learns common normalizing functions across multiple datasets while retaining image realism. The adversarial training provides an optimal normalizer that improves both the segmentation accuracy and the discrimination of unrealistic normalizing functions. Our contribution therefore leverages common imaging information from multiple domains. The optimality of our common normalizer is evaluated by combining brain images from both infants and adults. Results on the challenging iSEG and MRBrainS datasets reveal the potential of our adversarial normalization approach for segmentation, with Dice improvements of up to 59.6% over the baseline.

Index Terms— Task-driven intensity normalization, brain segmentation.

1. INTRODUCTION

In medical imaging applications, datasets with annotated images are rare and often composed of few samples. This causes an accessibility problem for developing supervised learning algorithms such as those based on deep learning. Although these algorithms have helped in automating image segmentation, notably in the medical field [1, 2], they need a massive number of training samples to obtain accurate segmentation masks that generalize well across different sites. One possible approach to alleviate this problem would be to use data acquired from multiple sites to increase the generalization performance of the learning algorithm. However, medical images from different datasets or sites can be acquired using various protocols. This leads to a high variance in image intensities and resolution, increasing the sensitivity of segmentation algorithms to raw images and thus impairing their performance.

Recently, the problems of image normalization and learned pre-processing of medical images have generated a growing interest. In [3], it was shown that two consecutive fully-convolutional deep neural networks (FCN) can normalize an input prior to segmentation. However, the intermediary synthetic images produced by forward pass on the first fully-




Fig. 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.

convolutional network lack interpretability as there is no constraint on the realism of produced images (see Fig. 1). Also, a limitation of this previous work is the separate processing of 2-D slices of volumetric data, which does not take into account the valuable 3-D context of this data. The study in [4] analyzed multi-site training of deep learning algorithms and compared various traditional normalization techniques such as histogram matching [5] or Gaussian standardization across different datasets. However, these standard normalization techniques are not learned for a specific task, certainly leading to suboptimal results compared to a task-driven normalization approach. Gradient reversal layer based domain adaptation and layer-wise domain-adversarial networks are used in [6] but this method is limited to two domains.

This paper extends medical image normalization so that the normalized images remain realistic and interpretable by clinicians. Our method also leverages information from multiple datasets by learning a joint normalizing transformation accounting for large image variability. We propose a task-and-data-driven adversarial normalization technique that constrains the normalized image to be realistic and optimal for the image segmentation task. Our approach exploits two fully-convolutional 3-D deep neural networks [7]. The first acts as a normalized image generator, while the second serves as segmentation network. This paper also includes a 3-D discriminator [8] network that constrains the generator to produce interpretable images. Standard generative adversarial networks (GANs) [9] aim at classifying generated images as fake or real. Our discriminator rather acts as a domain classifier by distinguishing images between all input domains (i.e. a dataset which has been sampled with a specific protocol at a specific location) including an additional “generated” one.




Fig. 2. 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 segmentation 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.

Hence, produced images are both realistic and domain invariant. The parameters of all three networks are learned end-to-end.

Our contributions can be summarized as follows: 1) an adversarially-constrained 3-D pre-processing and segmentation technique using fully-convolutional neural networks which can train on more than one dataset, 2) a learned normalization network for medical images which produces images that are realistic and interpretable by clinicians. The proposed method yields a significant improvement in segmentation performance over using a conventional segmentation model trained and tested on two different data distributions which haven't been normalized by a learning approach. To the best of our knowledge this is the first work using purely task-and-data driven medical image normalization while keeping the intermediary image medically usable.

2. METHOD

Let $\mathbf{x} \in \mathbb{R}^{|\Omega|}$ be a 3-D image, where Ω is the set of voxels, and $\mathbf{y} \in \{0, 1\}^{|\Omega| \times C}$ be its segmentation ground truth with pixel labels in $\{1, \dots, C\}$. The training set $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i, z_i) \mid i = 1, \dots, M\}$ contains M examples, each composed of an image \mathbf{x}_i , a manual expert segmentation \mathbf{y}_i and an image domain label $z_i \in \{1, \dots, K\}$. As shown in Fig. 2, the proposed model is composed of three networks. The first network G is a fully-convolutional network. A 3-D U-Net architecture, without loss of generality, has been chosen for its simplicity. This network transforms an input image \mathbf{x} into a cross-domain normalized image $\hat{\mathbf{x}} = G(\mathbf{x})$. The second network S , which is also a 3-D FCN, receives the normalized image as input and outputs the segmentation map $S(\hat{\mathbf{x}})$. The third network D is the discriminator which receives both raw images and normalized images as input and predicts their domain. Network D learns a $(K+1)$ -class

classification problem, with one class for each raw image domain and a $(K+1)$ -th class for generated images of any domain. As mentioned before, the discriminator is used to ensure that images produced by G are both realistic and domain invariant.

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

$$\min_{G, S} \max_D \mathcal{L}(G, S, D) = \sum_{i=1}^M \mathcal{L}_{\text{seg}}(S(G(\mathbf{x}_i)), \mathbf{y}_i) - \lambda \left[\sum_{i=1}^M (\mathcal{L}_{\text{dis}}(D(\mathbf{x}_i), z_i) + \mathcal{L}_{\text{dis}}(D(G(\mathbf{x}_i)), \text{fake})) \right]. \quad (1)$$

For training the segmentation network, we use the weighted Dice loss defined as

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

where $s_{v,c} \in [0, 1]$ is the softmax output of S for voxel v and class c , and ϵ is a small constant to avoid zero-division. For the discriminator classification loss, we employ we use the standard log loss. Let p_D be the output class distribution of D following the softmax. For raw (unnormalized) images, the loss is given by

$$\mathcal{L}_{\text{dis}}(D(\mathbf{x}), z) = -\log p_D(Z = z \mid \mathbf{x}). \quad (3)$$

On the other hand, in the case of generated (normalized) images, the loss becomes

$$\begin{aligned} \mathcal{L}_{\text{dis}}(D(G(\mathbf{x})), \text{fake}) &= -\log p_D(Z = \text{fake} \mid G(\mathbf{x})) \\ &= -\log [1 - p_D(Z \leq K \mid G(\mathbf{x}))] \end{aligned} \quad (4)$$

As in standard adversarial learning methods, we train our model in two alternating steps, first updating the parameters of G and S , and then updating the parameters of D .

Exp. #	Method	Train dataset	Test dataset	Dice		
				CSF	GM	WM
1	No adaptation	iSEG	iSEG	0.906	0.868	0.863
2	No adaptation	MRBrainS	MRBrainS	0.813	0.789	0.839
3	No adaptation, Cross-testing	iSEG	MRBrainS	0.401	0.354	0.519
4	No adaptation, Cross-testing	MRBrainS	iSEG	0.293	0.082	0.563
5	Standardized	iSEG + MRBrainS Standardized	iSEG + MRBrainS Standardized	0.849	0.808	0.809
6	Without constraint	iSEG + MRBrainS	iSEG + MRBrainS	0.834	0.859	0.885
7	Adversarially normalized (ours)	iSEG + MRBrainS	iSEG + MRBrainS	0.919	0.902	0.905

Table 1. Dice score in function of the model architecture and data. The proposed method yielded a significant performance improvement over training and testing on single-domain or on standardized inputs.

3. EXPERIMENTS AND RESULTS

3.1. Data

To evaluate the performance of our method, two databases have been carefully retained for their drastic difference in intensity profile and nature. The first one, iSEG [10], is a set of 10 T1 MRI images of infant. The ground truth is the segmentation of the three main structures of the brain: white matter (WM), gray matter (GM) and cerebrospinal fluids (CSF), all three being critical for detecting abnormalities in brain development. Images are sampled into an isotropic 1.0 mm^3 resolution. The second dataset is MRBrainS [11] which contains 5 adult subjects with T1 modality. The dataset also contains the same classes as ground truth. Images were acquired following a voxel size of $0.958 \text{ mm} \times 0.958 \text{ mm} \times 3.0 \text{ mm}$.

While both datasets have been acquired with a 3 Tesla scanner, iSEG is a set of images acquired on 6-8 month-old infants while MRBrainS is a set of adult images. This implies strong structural differences. Moreover, iSEG is particularly challenging because the images have been acquired during subject’s isointense phase in which the white matter and gray matter voxel intensities greatly overlap, thus leading to a lower tissue contrast. This lower contrast is known to misleads common classifiers.

3.2. Pre-processing and implementation details

Only T1 images were used in our experiments. Since images in MRBrainS dataset are a full head scan, skull stripping was performed using the segmentation map. Resampling to isotropic 1.0 mm^3 as been done to match iSEG sampling. Overlapping patches of 32^3 voxels centered on foreground were extracted from full volumes with stride of 8^3 which yielded 27,647 3-D patches in total for all training images. Each training took 8 hours using distributed stochastic gradient descent on two NVIDIA RTX 2080 Ti GPUs.

3.3. Experiment setup

We split 60% of the patches for training, while the other 40% are split in half for validation and testing. Split is done

	Input data	Normalized images
JSD	3.509	3.504

Table 2. Jensen-Shannon divergence (JSD) of input and normalized images from the generator. A lower value corresponds to more similar distributions.




Fig. 3. Mixed iSEG and MRBrainS inputs (**left**) and the generated images with adversarial normalization with $\lambda = 1.0$ (**right**). Notice the improved homogeneity of intensities in the normalized images, making their analysis easier.

following a stratified shuffle split strategy based on center voxel class. Seven different experimental settings were tested. Baselines for segmentation performance on respectively training and testing on each dataset are done (Exp. 1-2 in Table 1). Cross-domain testing is then done for both dataset (Exp. 3-4). For comparison purpose, we trained a segmentation network with both datasets on which we applied Gaussian standardization (Exp. 5). We also trained two pipelined FCN networks with both datasets at the same time (Exp. 6). We then trained using our method with both datasets (Exp. 7). For each experiment involving the generator network, a Mean-Square Error (MSE) loss is used to initialize weights and produce an output close to real input. Optimum is reached after three epochs. Segmentation Dice loss and discriminator’s cross-entropy loss is then added at the beginning of the fourth epoch. Each experiment has been




Fig. 4. Histograms of generator outputs (blue) and unnormalized inputs (red). The intensity range of the generated images for grey matter (GM), white matter (WM) and CSF is more compact, showing a reduced variance in voxel intensities.

trained with a Stochastic Gradient Descent optimizer with momentum of 0.9 and weight decay of 0.1. All networks have been initialized using Kaiming initialization [12]. Generator uses a learning rate of 0.00001 while segmentation and discriminator networks use 0.0001. A learning rate scheduler with patience of 3 epochs on validation losses reduces the learning rate of all optimizers by a factor of 10. No data augmentation was applied.

3.4. Normalization performance

To evaluate the normalization performance, we used the Jensen-Shannon divergence (JSD) between intensity histograms of images in the validation set. JSD measures the mean KL divergence between a distribution (i.e., histogram) and the average of distributions. Table 2 gives the JSD between input images and between images normalized by the generator. We see a decrease in JSD for normalized images showing that the intensity profiles of generated images are more similar to each other. The normalization effect of our method can be better appreciated in Fig. 4. This figure shows a narrower distribution that is more centered around a single mode therefore reducing the intra-class variance and increasing segmentation accuracy. Another benefit of our method is the contrast enhancement it provides to the generated images (Fig. 3). This is mainly brought by our task-driven approach, where minimizing the segmentation loss helps at increasing the contrast along region boundaries.

3.5. Segmentation performance

Our method relies on the segmentation task while performing online normalization. Since the Dice loss is being used, structural elements are kept when generating the intermediate image, while cross-entropy aims at keeping the global features of the image while reducing the differences across domains. The main advantage of our method is the ability to train on drastically different datasets regarding to structures and intensity distributions while still maintaining a good segmentation performance. We are effectively performing segmentation of structures on both adults and infant brains at the same time while still achieving 90.8% of mean Dice score across all classes. This demonstrates the relevance of adding adversarial normalization, increasing the Dice score of up to 59.6% up in

mean segmentation performance over all classes against training on a single dataset and testing on the other one. As seen in Fig. 3, our method is also able to normalize the images while maximizing the segmentation and keeping the image interpretable and realistic. This is achieved by the discriminator’s loss which aims at minimizing the cross-entropy between real inputs and generated input up to the point it cannot differentiate anymore from which domain (real iSEG input, real MR-BrainS input or generated input) the generator’s output comes from.

4. CONCLUSION

We proposed a novel task-and-data-driven normalization technique to improve a segmentation task using two datasets. Our method drastically improves performance on test sets by finding the optimal intermediate representation using the structural elements of the image brought by Dice loss and global features brought by cross-entropy. We believe this work is an important contribution to biomedical imaging as it would unlock the training of deep learning models with data from multiple sites, thus reducing the strain on data accessibility. This increase in data availability will help computing models with higher generalization performance. We also demonstrated that it is still possible with our method to train with different datasets while one of these is particularly difficult because of the lower image contrasts. Future work will aim at using a better model to act as a discriminator or using other training methods such as Wasserstein GAN [13] or CycleGAN [14] to compare performance difference on the image normalization. Our method is also theoretically capable of handling more than two datasets, finding automatically a common intermediate representation between all input domains and adapting voxel intensities to maximize the segmentation performance. Further work will also demonstrate the architecture’s performance on new datasets and on the gain this task-driven normalization can provide in case of segmentation tasks with greatly imbalanced data such as brain lesion segmentation.

Acknowledgment – This work was supported financially by the Research Council of Canada (NSERC), the Fonds de Recherche du Quebec (FQRNT), ETS Montreal, and NVIDIA with the donation of a GPU.

5. REFERENCES

- [1] K. Kamnitsas, C. Ledig, V. Newcombe, J. Simpson, A. Kane, D. Menon, D. Rueckert, and B. Glocker, "Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation," *Medical Image Analysis*, 2017.
- [2] J. Dolz, K. Gopinath, J. Yuan, H. Lombaert, C. Desrosiers, and I. Ben Ayed, "HyperDense-Net: A Hyper-Densely Connected CNN for Multi-Modal Image Segmentation," *IEEE Transactions on Medical Imaging*, 2019.
- [3] Michal Drozdal, Gabriel Chartrand, Eugene Vorontsov, Mahsa Shakeri, Lisa Di Jorio, An Tang, Adriana Romero, Yoshua Bengio, Chris Pal, and Samuel Kadoury, "Learning normalized inputs for iterative estimation in medical image segmentation," *Medical Image Analysis*, 2018.
- [4] J. Onofrey, D. Casetti-Dinescu, A. Lauritzen, S. Sarkar, R. Venkataraman, R. Fan, G. Sonn, P. Sprenkle, L. Staib, and X. Papademetris, "Generalizable Multi-Site Training and Testing Of Deep Neural Networks Using Image Normalization," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2019.
- [5] D. Shapira, S. Avidan, and Y. Hel-Or, "Multiple histogram matching," in *IEEE International Conference on Image Processing (ICIP)*, 2013.
- [6] O. Ciga, J. Chen, and A. Martel, *Multi-layer Domain Adaptation for Deep Convolutional Networks*, 2019.
- [7] Ozgun Cicek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger, "3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2016.
- [8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Networks," in *International Conference on Machine Learning (ICML)*, 2014.
- [10] L. Wang, D. Nie, G. Li, É. Puybureau, J. Dolz, Q. Zhang, F. Wang, J. Xia, Z. Wu, J. Chen, K. Thung, T. D. Bui, J. Shin, G. Zeng, G. Zheng, V. S. Fonov, A. Doyle, Y. Xu, P. Moeskops, J. P. W. Pluim, C. Desrosiers, I. B. Ayed, G. Sanroma, O. M. Benkarim, A. Casamitjana, V. Vilaplana, W. Lin, G. Li, and D. Shen, "Benchmark on Automatic Six-Month-Old Infant Brain Segmentation Algorithms: The iSeg-2017 Challenge," *IEEE Transactions on Medical Imaging*, 2019.
- [11] A. Mendrik, K. Vincken, H. Kuijf, M. Breeuwer, W. Bouvy, J. de Bresser, A. Alansary, M. de Bruijne, A. Carass, A. El-Baz, Amod Jog, R. Katyal, A. Khan, F. Lijn, Q. Mahmood, R. Mukherjee, A. Opbroek, S. Paneri, S. Pereira, and M. Viergever, "MRBrainS Challenge: Online Evaluation Framework for Brain Image Segmentation in 3T MRI Scans," *Computational Intelligence and Neuroscience*, 2015.
- [12] K. He, X. Zhang, S. Ren, and J. Sun, "Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification," in *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [13] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein Generative Adversarial Networks," in *International Conference on Machine Learning (ICML)*, 2017.
- [14] J. Zhu, T. Park, P. Isola, and A. Efros, "Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks," in *IEEE International Conference on Computer Vision (ICCV)*, 2017.