

# Adversarial Training and Dilated Convolutions for Brain MRI Segmentation

Pim Moeskops<sup>(✉)</sup>, Mitko Veta, Maxime W. Lafarge, Koen A.J. Eppenhof, and Josien P.W. Pluim

Medical Image Analysis Group, Department of Biomedical Engineering,  
Eindhoven University of Technology, Eindhoven, The Netherlands  
p.moeskops@tue.nl

**Abstract.** Convolutional neural networks (CNNs) have been applied to various automatic image segmentation tasks in medical image analysis, including brain MRI segmentation. Generative adversarial networks have recently gained popularity because of their power in generating images that are difficult to distinguish from real images.

In this study we use an adversarial training approach to improve CNN-based brain MRI segmentation. To this end, we include an additional loss function that motivates the network to generate segmentations that are difficult to distinguish from manual segmentations. During training, this loss function is optimised together with the conventional average per-voxel cross entropy loss.

The results show improved segmentation performance using this adversarial training procedure for segmentation of two different sets of images and using two different network architectures, both visually and in terms of Dice coefficients.

**Keywords:** Adversarial networks · Deep learning · Convolutional neural networks · Dilated convolution · Medical image segmentation · Brain MRI

## 1 Introduction

Convolutional neural networks (CNNs) have become a very popular method for medical image segmentation. In the field of brain MRI segmentation, CNNs have been applied to tissue segmentation [13, 14, 20] and various brain abnormality segmentation tasks [3, 5, 8].

A relatively new approach for segmentation with CNNs is the use of dilated convolutions, where the weights of convolutional layers are sparsely distributed over a larger receptive field without losing coverage on the input image [18, 19]. Dilated CNNs are therefore an effective approach to achieve a large receptive field with a limited number of trainable weights and a limited number of convolutional layers, without the use of subsampling layers.

Generative adversarial networks (GANs) provide a method to generate images that are difficult to distinguish from real images [4, 15, 17]. To this end,

GANs use a discriminator network that is optimised to discriminate real from generated images, which motivates the generator network to generate images that look real. A similar adversarial training approach has been used for domain adaptation, using a discriminator network that is trained to distinguish images from different domains [2, 7] and for improving image segmentations, using a discriminator network that is trained to distinguish manual from generated segmentations [11]. Recently, such a segmentation approach has also been applied in medical imaging for the segmentation of prostate cancer in MRI [9] and organs in chest X-rays [1].

In this paper we employ adversarial training to improve the performance of brain MRI segmentation in two sets of images using a fully convolutional and a dilated network architecture.

## 2 Materials and Methods

### 2.1 Data

**Adult Subjects.** 35  $T_1$ -weighted MR brain images (15 training, 20 test) were acquired on a Siemens Vision 1.5T scanner at an age ( $\mu \pm \sigma$ ) of  $32.9 \pm 19.2$  years, as provided by the MICCAI 2012 challenge on multi-atlas labelling [10]. The images were segmented in six classes: white matter (WM), cortical grey matter (cGM), basal ganglia and thalami (BGT), cerebellum (CB), brain stem (BS), and lateral ventricular cerebrospinal fluid (lvCSF).

**Elderly Subjects.** 20 axial  $T_1$ -weighted MR brain images (5 training, 15 test) were acquired on a Philips Achieva 3T scanner at an age ( $\mu \pm \sigma$ ) of  $70.5 \pm 4.0$  years, as provided by the MRBrainS13 challenge [12]. The images were segmented in seven classes: WM, cGM, BGT, CB, BS, lvCSF, and peripheral cerebrospinal fluid (pCSF). Possible white matter lesions were included in the WM class.

### 2.2 Network Architecture

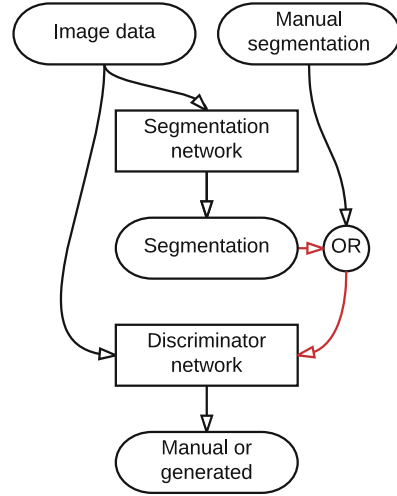
Two different network architectures are used to evaluate the hypothesis that adversarial training can aid in improving segmentation performance: a fully convolutional network and a network with dilated convolutions. The outputs of these networks are input for a discriminator network, which distinguishes between generated and manual segmentations. The fully convolutional nature of both networks allows arbitrarily sized inputs during testing. Details of both segmentation networks are listed in Fig. 1, left.

**Fully Convolutional Network.** A network with 15 convolutional layers of  $3 \times 3$  kernels is used (Fig. 1, left), which results in a receptive field of  $31 \times 31$  voxels. During training, an input of  $51 \times 51$  voxels is used, corresponding to an output of  $21 \times 21$  voxels. The network has 140,039 trainable parameters for  $C = 7$  classes (6 plus background; adult subjects) and 140,296 trainable parameters for  $C = 8$  classes (7 plus background; elderly subjects).

Fully convolutional network			
Kernel size	Dilation	Kernels	Layers
$3 \times 3$	1	32	15
$1 \times 1$	1	256	1
$1 \times 1$	1	$C$	1

Dilated network			
Kernel size	Dilation	Kernels	Layers
$3 \times 3$	1	32	2
$3 \times 3$	2	32	1
$3 \times 3$	4	32	1
$3 \times 3$	8	32	1
$3 \times 3$	16	32	1
$3 \times 3$	1	32	1
$1 \times 1$	1	$C$	1

(a) Segmentation networks

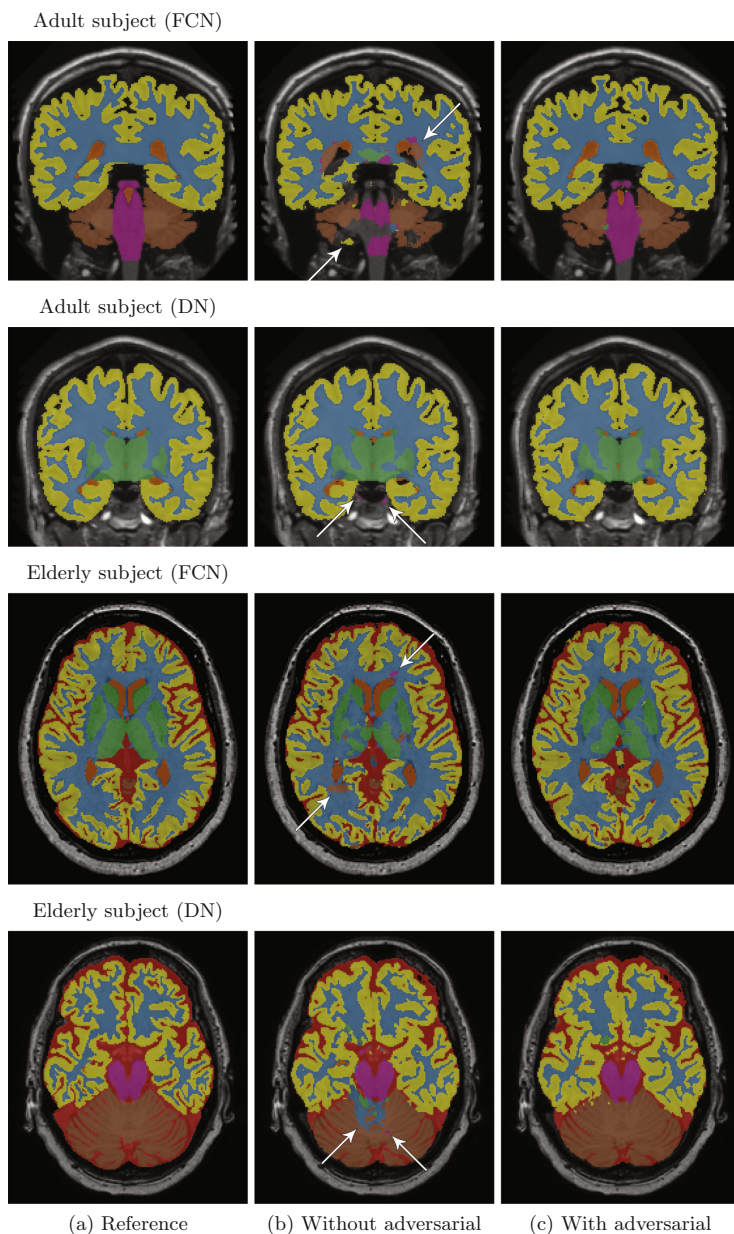


(b) Adversarial training

**Fig. 1.** Left: Segmentation network architectures for the 17-layer fully convolutional (top) and 8-layer dilated (bottom) segmentation networks. The receptive fields are  $67 \times 67$  for the dilated network and  $31 \times 31$  for the fully convolutional network. No subsampling layers are used in both networks. Right: Overview of the adversarial training procedure. The red connections indicate how the discriminator loss influences the segmentation network during backpropagation. (Color figure online)

**Dilated Network.** The dilated network uses the same architecture as proposed by Yu et al. [19], which uses layers of  $3 \times 3$  kernels with increasing dilation factors (Fig. 1, left). This results in a receptive field of  $67 \times 67$  voxels using only 7 layers of  $3 \times 3$  convolutions, without any subsampling layers. During training, an input of  $87 \times 87$  voxels is used, which corresponds to an output of  $21 \times 21$  voxels. In each layer 32 kernels are trained. The network has 56,039 trainable parameters for  $C = 7$  classes (6 plus background; adult subjects) and 56,072 trainable parameters for  $C = 8$  classes (7 plus background; elderly subjects).

**Discriminator Network.** The input to the discriminator network are the segmentation, as one-hot encoding or softmax output, and image data in the form of a  $25 \times 25$  patch. In this way, the network can distinguish real from generated combinations of image and segmentation patches. The image patch and the segmentation are concatenated after two layers of  $3 \times 3$  kernels on the image patch. The discriminator network further consists of three layers of  $32 \times 3 \times 3$  kernels, a  $3 \times 3$  max-pooling layer, two layers of  $32 \times 3 \times 3$  kernels, and a fully connected layer of 256 nodes. The output layer with two nodes, distinguishes between manual and generated segmentations.



**Fig. 2.** Example segmentation results in four of the test images. From top to bottom: an adult subject using the fully convolutional network (FCN), an adult subject using the dilated network (DN), an elderly subject using the fully convolutional network (FCN), and an elderly subject using the dilated network (DN). The colours are as follows: WM in blue, cGM in yellow, BGT in green, CB in brown, BS in purple, lvCSF in orange, and pCSF in red. The arrows indicate errors that were corrected when the adversarial training procedure was used. (Color figure online)

### 2.3 Adversarial Training

An overview of the adversarial training procedure is shown in Fig. 1, right.

Three types of updates for the segmentation network parameters  $\theta_s$  and the discriminator network parameters  $\theta_d$  are possible during the training procedure: (1) an update of only the segmentation network based on the cross-entropy loss over the segmentation map,  $L_s(\theta_s)$ , (2) an update of the discriminator network based on the discrimination loss using a manual segmentation as input,  $L_d(\theta_d)$ , and (3) an update of the whole network (segmentation and discriminator network) based on the discriminator loss using an image as input,  $L_a(\theta_s, \theta_d)$ . Only  $L_s(\theta_s)$  and  $L_a(\theta_s, \theta_d)$  affect the segmentation network. The parameters  $\theta_s$  are updated to maximise the discriminator loss  $L_a(\theta_s, \theta_d)$ , i.e. the updates for the segmentation network are performed in the direction to ascend the loss instead of to descend the loss.

The three types of updates are performed in an alternating fashion. The updates based on the segmentation loss and the updates based on the discriminator loss are performed with separate optimisers using separate learning rates. Using a smaller learning rate, the discriminator network adapts more slowly than the segmentation network, such that the discriminator loss does not converge too quickly and can have enough influence on the segmentation network.

For each network, rectified linear units are used throughout, batch normalisation [6] is used on all layers and dropout [16] is used for the  $1 \times 1$  convolution layers.

## 3 Experiments and Results

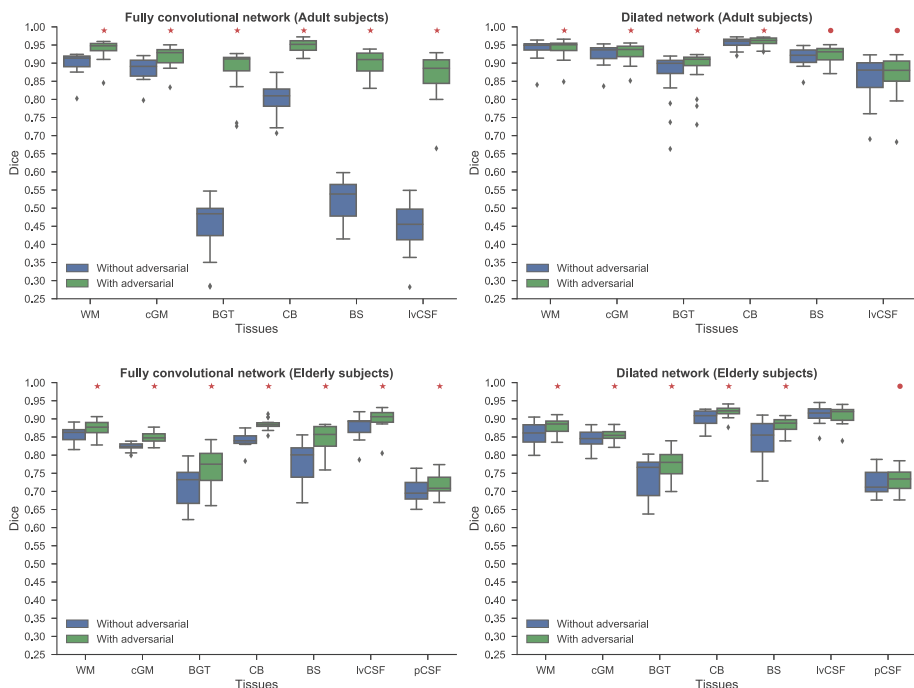
### 3.1 Experiments

As a baseline, the segmentation networks are trained without the adversarial network. The updates are performed with RMSprop using a learning rate of  $10^{-3}$  and minibatches of 300 samples. The networks are trained in 5 epochs, where each epoch corresponds to 50,000 training patches per class per image. Note that during this training sample balancing process, the class label corresponds to the label of the central voxel, even though a larger image patch is labelled.

The discriminator and segmentation network are trained using the alternating update scheme. The updates for both loss functions are performed with RMSprop using a learning rate of  $10^{-3}$  for the segmentation loss and a learning rate of  $10^{-5}$  for the discriminator loss. The updates alternate between the  $L_s$ ,  $L_d$  and  $L_a$  loss functions, using minibatches of  $300/3 = 100$  samples for each.

### 3.2 Evaluation

Figure 2 provides a visual comparison between the segmentations obtained with and without adversarial training, showing that the adversarial approach generally resulted in less noisy segmentations. The same can be seen from the total number of 3D components (including the background class) that compose the



**Fig. 3.** Dice coefficients for the adult subjects (top row) and the elderly subjects (bottom row) for white matter (WM), cortical grey matter (cGM), basal ganglia and thalami (BGT), cerebellum (CB), brain stem (BS), lateral ventricular cerebrospinal fluid (lvCSF), and peripheral cerebrospinal fluid (pCSF), without (blue) and with (green) adversarial training. Left column: fully convolutional network. Right column: dilated network. Red stars ( $p < 0.01$ ) and red circles ( $p < 0.05$ ) indicate significant improvement based on paired  $t$ -tests. (Color figure online)

segmentations. For the adult subjects, the number of components per image ( $\mu \pm \sigma$ ) decreased from  $1745 \pm 400$  to  $626 \pm 247$  using the fully convolutional network and from  $417 \pm 152$  to  $365 \pm 122$  using the dilated network. For the elderly subjects, the number of components per image ( $\mu \pm \sigma$ ) decreased from  $926 \pm 134$  to  $692 \pm 88$  using the fully convolutional network and from  $601 \pm 104$  to  $481 \pm 90$  using the dilated network.

The evaluation results in terms of Dice coefficients (DC) between the automatic and manual segmentations are shown in Fig. 3 as boxplots. Significantly improved DC, based on paired  $t$ -tests, were obtained for each of the tissue classes, in both image sets, and for both networks. The only exception was lvCSF in the elderly subjects using the dilated network. For the adult subjects, the DC averaged over all 6 classes ( $\mu \pm \sigma$ ) increased from  $0.67 \pm 0.04$  to  $0.91 \pm 0.03$  using the fully convolutional network and from  $0.91 \pm 0.03$  to  $0.92 \pm 0.03$  using the dilated network. For the elderly subjects, the DC averaged over all 7 classes ( $\mu \pm \sigma$ )

increased from  $0.80 \pm 0.02$  to  $0.83 \pm 0.02$  using the fully convolutional network and from  $0.83 \pm 0.02$  to  $0.85 \pm 0.01$  using the dilated network.

## 4 Discussion and Conclusions

We have presented an approach to improve brain MRI segmentation by adversarial training. The results showed improved segmentation performance both qualitatively (Fig. 2) and quantitatively in terms of DC (Fig. 3). The improvements were especially clear for the deeper, more difficult to train, fully convolutional networks as compared with the more shallow dilated networks. Furthermore, the approach improved structural consistency, e.g. visible from the reduced number of components in the segmentations. Because these improvements were usually small in size, their effect on the DC was limited.

The approach includes an additional loss function that distinguishes between real and generated segmentations and can therefore capture inconsistencies that a normal per-voxel loss averaged over the output does not capture. The proposed approach can be applied to any network architecture that, during training, uses an output in the form of an image patch, image slice, or full image instead of a single pixel/voxel.

Various changes to the segmentation network that might improve the results could be evaluated in future work, such as different receptive fields, multiple inputs, skip-connections, 3D inputs, etc. Using a larger output patch size or even the whole image as output could possibly increase the effect of the adversarial training by including more information that could help in distinguishing manual from generated segmentations. This could, however, also reduce the influence of local information, resulting in a too global decision. Further investigation is necessary to evaluate which of the choices in the network architecture and training procedure have most effect on the results.

**Acknowledgements.** The authors would like to thank the organisers of MRBrainS13 and the multi-atlas labelling challenge for providing the data. The authors gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan X Pascal GPU.

## References

1. Dai, W., Doyle, J., Liang, X., Zhang, H., Dong, N., Li, Y., Xing, E.P.: SCAN: structure correcting adversarial network for chest X-rays organ segmentation. arXiv preprint [arXiv:1703.08770](https://arxiv.org/abs/1703.08770) (2017)
2. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *J. Mach. Learn. Res.* **17**(59), 1–35 (2016)
3. Ghafoorian, M., Karssemeijer, N., Heskes, T., Bergkamp, M., Wissink, J., Obels, J., Keizer, K., de Leeuw, F.E., van Ginneken, B., Marchiori, E., Platel, B.: Deep multi-scale location-aware 3D convolutional neural networks for automated detection of lacunes of presumed vascular origin. *NeuroImage. Clin.* **14**, 391–399 (2017)

4. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS, pp. 2672–2680 (2014)
5. Havaei, M., Davy, A., Warde-Farley, D., Biard, A., Courville, A., Bengio, Y., Pal, C., Jodoin, P.M., Larochelle, H.: Brain tumor segmentation with deep neural networks. *Med. Image Anal.* **35**, 18–31 (2017)
6. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: ICML (2015)
7. Kamnitsas, K., et al.: Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. In: Niethammer, M., Styner, M., Aylward, S., Zhu, H., Oguz, I., Yap, P.-T., Shen, D. (eds.) IPMI 2017. LNCS, vol. 10265, pp. 597–609. Springer, Cham (2017). doi:[10.1007/978-3-319-59050-9\\_47](https://doi.org/10.1007/978-3-319-59050-9_47)
8. Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B.: Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Med. Image Anal.* **36**, 61–78 (2017)
9. Kohl, S., Bonekamp, D., Schlemmer, H.P., Yaqubi, K., Hohenfellner, M., Hadaschik, B., Radtke, J.P., Maier-Hein, K.: Adversarial networks for the detection of aggressive prostate cancer. arXiv preprint [arXiv:1702.08014](https://arxiv.org/abs/1702.08014) (2017)
10. Landman, B.A., Ribbens, A., Lucas, B., Davatzikos, C., Avants, B., Ledig, C., Ma, D., Rueckert, D., Vandermeulen, D., Maes, F., Erus, G., Wang, J., Holmes, H., Wang, H., Doshi, J., Kornegay, J., Manjon, J., Hammers, A., Akhondi-Asl, A., Asman, A.J., Warfield, S.K.: MICCAI 2012 Workshop on Multi-Atlas Labeling. CreateSpace Independent Publishing Platform, Neme (2012)
11. Luc, P., Couprie, C., Chintala, S., Verbeek, J.: Semantic segmentation using adversarial networks. In: NIPS Workshop on Adversarial Training (2016)
12. Mendrik, A.M., Vincken, K.L., Kuijf, H.J., Breeuwer, M., Bouvy, W.H., de Bresser, J., Alansary, A., de Bruijne, M., Carass, A., El-Baz, A., Jog, A., Katyal, R., Khan, A.R., van der Lijn, F., Mahmood, Q., Mukherjee, R., van Opbroek, A., Paneri, S., Pereira, S., et al.: MRBrainS challenge: online evaluation framework for brain image segmentation in 3T MRI scans. *Comput. Intel. Neurosci.* **2015** (2015). Article No. 813696
13. Moeskops, P., Viergever, M.A., Mendrik, A.M., de Vries, L.S., Benders, M.J., Išgum, I.: Automatic segmentation of MR brain images with a convolutional neural network. *IEEE Trans. Med. Imag.* **35**(5), 1252–1261 (2016)
14. Moeskops, P., Wolterink, J.M., Velden, B.H.M., Gilhuijs, K.G.A., Leiner, T., Viergever, M.A., Išgum, I.: Deep learning for multi-task medical image segmentation in multiple modalities. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 478–486. Springer, Cham (2016). doi:[10.1007/978-3-319-46723-8\\_55](https://doi.org/10.1007/978-3-319-46723-8_55)
15. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. In: ICLR (2016)
16. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**(1), 1929–1958 (2014)
17. Wolterink, J.M., Leiner, T., Viergever, M.A., Išgum, I.: Generative adversarial networks for noise reduction in low-dose CT. *IEEE Trans. Med. Imag.* (2017). <https://doi.org/10.1109/TMI.2017.2708987>



18. Wolterink, J.M., Leiner, T., Viergever, M.A., Išgum, I.: Dilated convolutional neural networks for cardiovascular MR segmentation in congenital heart disease. In: Zuluaga, M.A., Bhatia, K., Kainz, B., Moghari, M.H., Pace, D.F. (eds.) RAMBO/HVSMR -2016. LNCS, vol. 10129, pp. 95–102. Springer, Cham (2017). doi:[10.1007/978-3-319-52280-7\\_9](https://doi.org/10.1007/978-3-319-52280-7_9)
19. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. In: ICLR (2016)
20. Zhang, W., Li, R., Deng, H., Wang, L., Lin, W., Ji, S., Shen, D.: Deep convolutional neural networks for multi-modality isointense infant brain image segmentation. *NeuroImage* **108**, 214–224 (2015)