Vascular Segmentation in TOF MRA Images of the Brain Using a Deep Convolutional Neural Network

Renzo Phellan[1(✉)], Alan Peixinho[2], Alexandre Falcão[2], and Nils D. Forkert[1]

[1] Department of Radiology and Hotchkiss Brains Institute,
University of Calgary, Calgary, AB, Canada
{phellan.renzo,nils.forkert}@ucalgary.ca
[2] Laboratory of Image Data Science, Institute of Computing,
University of Campinas, Campinas, SP, Brazil
peixinho@lids.ic.unicamp.br, afalcao@ic.unicamp.br

Abstract. Cerebrovascular diseases are one of the main causes of death and disability in the world. Within this context, fast and accurate automatic cerebrovascular segmentation is important for clinicians and researchers to analyze the vessels of the brain, determine criteria of normality, and identify and study cerebrovascular diseases. Nevertheless, automatic segmentation is challenging due to the complex shape, inhomogeneous intensity, and inter-person variability of normal and malformed vessels. In this paper, a deep convolutional neural network (CNN) architecture is used to automatically segment the vessels of the brain in time-of-flight magnetic resonance angiography (TOF MRA) images of healthy subjects. Bi-dimensional manually annotated image patches are extracted in the axial, coronal, and sagittal directions and used as input for training the CNN. For segmentation, each voxel is individually analyzed using the trained CNN by considering the intensity values of neighboring voxels that belong to its patch. Experiments were performed with TOF MRA images of five healthy subjects, using varying numbers of images to train the CNN. Cross validations revealed that the proposed framework is able to segment the vessels with an average Dice coefficient ranging from 0.764 to 0.786 depending on the number of images used for training. In conclusion, the results of this work suggest that CNNs can be used to segment cerebrovascular structures with an accuracy similar to other high-level segmentation methods.

Keywords: Vessel segmentation · Deep learning · Cerebrovascular segmentation · Convolutional neural networks

1 Introduction

Vascular diseases have led the ranking of major causes of death in the last fifteen years, according to reports of the world health organization (WHO) [1]. In particular, cerebrovascular diseases that lead to stroke were responsible for more

than six million deaths only in 2015 [1]. Consequently, clinicians and researchers require fast and accurate tools, which aid them to detect, analyze, and treat cerebrovascular diseases, such as aneurysms, arteriovenous malformations (AVMs), and stenoses.

Segmentation of the vascular system in medical images allows clinicians to identify and isolate vessels from other surrounding types of tissue, thus, allowing better visualization and quantitative analysis. However, manual vessel segmentation is a time-consuming, error-prone task, which is subject to inter-observer variability. Consequently, research has been focused on developing faster and more accurate automatic vessel segmentation methods.

Lesage et al. [2] review paper lists a considerable number of automatic vessel segmentation approaches. The referenced methods range from approaches that are based on computing Hessian-based features of vessels, proposed by Frangi et al. [3] and Sato et al. [4], to atlas or model-based approaches, other feature-based methods, and extraction schemes, such as level-sets [5]. In all cases, different handcrafted features are used to guide the segmentation process, such as image intensities, Hessian eigenvalues, curvature values, gradient flow, and many others. It is the researcher who decides, based on experiments related to each particular application, which features are used to extract the to the most accurate segmentation results.

Deep convolutional neural networks (CNN) is a recent and popular strategy, with successful results solving different medical image analysis problems [6], which proposes to let the computer learn in an automatic and supervised manner, and decide which features are relevant to generate accurate segmentation results. Automatic vessel segmentation methods that use deep CNN have been used to segment 2D images of the retina [7], ultrasound images of the femoral region of the body [8], and computed tomography (CT) volumes of the liver [9], with a high performance in all cases.

To our knowledge, no study has been performed yet to adapt and apply deep CNN to segment the vessels of the brain, mainly due to the technical difficulties to obtain manually segmented brain datasets, the novelty of deep learning methods, and its associated long execution times. However, given the successful performance of CNN, as it has permeated the entire field of medical image analysis [6], this paper presents an initial strategy to apply deep CNN to segment the vascular system in time-of-flight magnetic resonance angiography (TOF MRA) images of the brain.

2 Vascular Segmentation of TOF MRA Images Using Deep CNN

TOF MRA is a medical imaging modality, which allows the acquisition of non-contrast enhanced images of the brain vascular system with a high spatial resolution [10]. TOF MRA images are affected by noise artifacts that do not allow the establishment of fixed intensity values to identify different types of tissue. Additionally, the intricate shape of the vascular tree, and its high inter-person

variability, make it hard to define a common atlas that can be used for segmentation, as often conducted for different organs [11].

Given the variability of the intensity profile and complex shape of the vascular system in TOF MRA and other imaging modalities, defining suitable characteristics to identify and segment vessels represents a challenging problem. In order to solve this problem, deep CNN approaches have been used to let the computer discover and learn those characteristics by itself, in a supervised manner [8,12].

Traditionally, three-dimensional patches are extracted from a set of training images and used to optimize a deep CNN, which is then used to segment the vascular system, but they have not been tested for the purpose of segmenting cerebrovascular structures from 3D TOF MRA datasets yet.

2.1   CNN Architecture

Complex deep CNN architectures can lead to a possible over-fitting in the model learning, as well as significantly increasing the processing time, when considering TOF MRA images. For this reason, we propose a CNN architecture composed of only two convolutional layers (C1 and C2) and two fully connected layers (FC3 and FC4). This architecture is shown in Fig. 1.

The first convolutional layer, C1, contains 32 filters with 5×5 voxels receptive field, in a 2 voxels stride sliding (S1), sub sampled in a 3×3 voxels max-pooling (P1), in order to reduce translation variance. The next convolutional layer, C2, has a receptive field of 3×3 voxels, with 64 filters, and no sub sampling. In order to reduce the impact of the backpropagation vanishing problem, both convolutional layers are followed by a rectified linear activation (Relu).

After the convolutional layers, two more fully connected layers are added. The first fully connected layer, FC3, reduces the dimensionality from 256 ($2 \times 2 \times 64$) to 100 neurons, and FC4 can be seen as a decision layer that determines the likelihood of belonging to a vessel or not. These layers have hyperbolic tangent (Tanh) and sigmoid (Sigm) activation functions, respectively.

2.2   CNN Training

In order to identify the best weights for our model, we selected a balanced number of patches from vessel and non-vessel regions in our training dataset. In particular, we used a number of vessel and non-vessel patches equal to half the number of voxels in the vessel region of each dataset, in the axial, coronal, and sagittal planes. Through a mini-batch gradient descent approach, the squared error over the entire training set was minimized, considering a mini-batch of 50 elements. This learning approach is applied through 40 epochs, while considering a learning rate of 0.001 and a gradient momentum of 0.9.
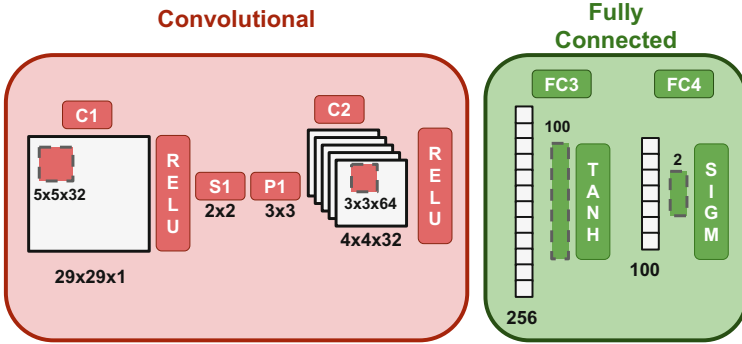
Fig. 1. Network architecture, composed of two convolutional layers, C1 and C2, and two fully connected layers, FC1 and FC2. After C1, we include a stride S1 of two voxels. All layers are also followed by a Relu, Tanh or Sigmoid function as indicated.

3   Materials and Methods

3.1   Data Acquisition and Image Preprocessing

Five TOF MRA datasets of healthy subjects were used to analyze and evaluate the proposed deep learning cerebrovascular segmentation method. The datasets were acquired on a 3T Intera MRI scanner (Philips, Eindhoven, the Netherlands) without application of contrast agent using a TE = 2.68 ms, a TR = 15.72 ms, a 20∘ flip angle, and a spatial resolution of $0.35 \times 0.35 \times 0.65 \, mm^3$. The datasets size is $512 \times 512 \times 120$ voxels.

For preprocessing, slab boundary artefact correction was performed using the method described by Kholmovski et al. [13] followed by intensity non-uniformity correction using the N3 algorithm [14]. A skull stripping algorithm [15] was also applied to mask the brain images and their corresponding binary segmentations. The vessels were manually segmented in each dataset by a medical expert based on the preprocessed TOF MRA datasets.

3.2   Classification

For all voxels inside the brain region, we define a cubic region of $29 \times 29 \times 29$ around this voxel, where the axial, coronal, and sagittal patches are extracted, as in [9]. All patches have $29 \times 29 \times 1$ voxels, as they are a bi-dimensional slice of each axis. Each patch is fed to the CNN, which calculates the vessel likelihood, so that three probability maps (for each orientation) are available after application of the CNN. A voxel is defined as a vessel voxel if at least one of the probability values (for the three directions) is above a threshold t and defined background otherwise. In this work, an empirically defined threshold of $t = 0.95$ was used for all experiments.

3.3   CNN Evaluation

The performance of the deep CNN is evaluated by selecting random sets of TOF MRA images for training. The number of images used for training is increased from one to four images, to evaluate if increasing the number of training images generates more accurate results. Initially, one TOF MRA image is randomly selected to train the CNN, which is used to segment the test image. The selected training image is different for each test image. Then, the number of training images is consecutively incremented up to four, always guaranteeing that the training set does not contain the test image.

The Dice similarity coefficient (DSC) [16] is used to compare the CNN-based segmentation and ground-truth manual segmentations, as it has been used in other cerebrovascular segmentation methods, thus, allowing an easier comparison. It is defined as $DSC = 2|A \cap B|/(|A| + |B|)$, where A and B represent the ground-truth and CNN segmentations, respectively.

A standard one-way analysis of variance (ANOVA) is applied to determine if the segmentation accuracies using an increasing number of images are statistically different, followed by the Tukey's honest significant difference procedure. The Statistical Package for the Social Sciences version 16.0 (SPSS Inc., Chicago, IL, USA) was used for this statistical analysis, and the criterion of statistical significance was set at $p < 0.05$.

3.4   Hardware Settings

Our deep CNN is implemented using version 2.7 of the python language, and the Theano 0.9.0 library [17]. Experiments are executed on a desktop computer with eight Intel(R) Core(TM) i7-4790 CPU @ 3.60 GHz processors, 32 GB of RAM memory, and graphic card GeForce GTX 745 (NVIDIA corp., United States), with 4 GB RAM memory. Testing and training were done using the graphic card and cuDNN extensions for faster processing [18].

4   Results

The results of the CNN approach for segmentation of vessels in TOF MRA images are reported in Table 1. The values correspond to the DSC when comparing the CNN segmentation with the manual ground-truths available for evaluation. Each row corresponds to a different dataset, and each column to the corresponding DSC when using the indicated number of images to train the deep neural network. The numbers in parenthesis identify the datasets used for training. Average DSC values, training and testing times are reported in the final rows.

The average DSC values for our deep learning approach vary between 0.764 and 0.786, depending on the number of images used for training. According to the ANOVA analysis, there is not enough evidence to guarantee that the resulting DSC values are significantly different (p >= 0.05). As expected, training times increase with the number of images used. The testing times are independent of the number of training images.

**Table 1.** DSC values for each dataset, when using a different number of images to train the deep neural network.

Dataset	1 image	2 images	3 images	4 images
1	0.774 (2)	0.763 (4, 5)	0.770 (2, 3, 4)	0.767
2	0.758 (5)	0.780 (3, 5)	0.759 (1, 3, 4)	0.765
3	0.769 (4)	0.784 (1, 4)	0.730 (1, 2, 4)	0.771
4	0.770 (2)	0.795 (1, 2)	0.804 (1, 3, 5)	0.781
5	0.751 (1)	0.809 (1, 3)	0.742 (1, 2, 4)	0.788
Average	**0.764±0.010**	**0.786±0.017**	**0.761±0.028**	**0.774±0.010**
Train (min)	40	65	92	120
Test (min)	30	30	30	30

5   Discussion

This paper presents a feasibility analysis of a deep CNN vessel segmentation method for TOF MRA images of the brain, with promising accuracy results. The CNN analyzes only in-plane neighboring voxels in the axial, coronal, and sagittal planes, and not full three-dimensional patches. Additionally, the Theano library with cuDNN extensions, and graphic card were used in the CNN implementation.

According to the executed statistical tests, using more images for training did not lead to a significant increase in segmentation accuracy. This result clearly highlights the benefit that a simple CNN, as described here, only needs very few well segmented ground truth datasets to achieve proper results, making an application in research or clinical settings more feasible.

Figure 2 shows 3D visualizations of the segmentation results for dataset 1, using the indicated number of images to train the proposed deep CNN. Visually, no considerable difference can be depicted between the segmentation results, confirming by the quantitative analysis. In general, it can be noted that large vessels are correctly identified. On the other hand, small vessels are partially affected by noise, such that their shape is not correctly delineated.

Manual   Train 1   Train 2   Train 3   Train 4
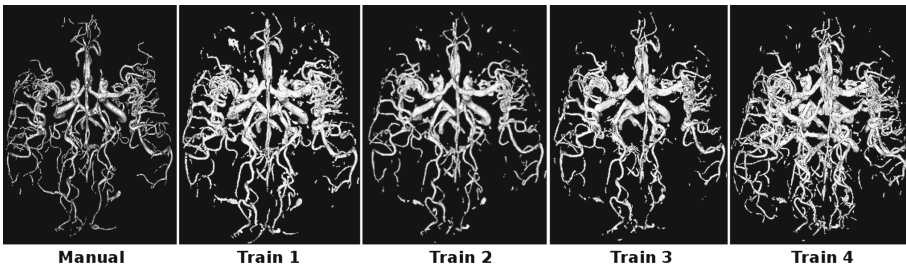
Fig. 2. 3D renderings of segmentation results for dataset 1 using a deep CNN trained with the indicated number of images.

As a limitation, it has to be noted that a small sample size with five TOF MRA images may not be enough to support general conclusions about the most suitable deep CNN architecture for vessel segmentation. However, the promising results of this initial analysis (as seen in Fig. 2) motivates further developments and analyses of this approach.

6   Conclusion

This paper presents a first feasibility analysis to apply deep CNN for automatic segmentation of the cerebrovascular system. Processing times were optimized by using bi-dimensional patches to identify vessels, and by taking advantage of the Theano library with cuDNN extensions, and graphic card of the system. No significant accuracy differences were found when using different numbers of images to train the deep CNN. The developed program calculates axial, coronal, and sagittal vessel probability maps and applies a fixed threshold to determine which voxels belong to vessels. It is expected that more complex approaches based on the calculated probability maps would lead to more accurate results.

References

1. World Health Organization: The top 10 causes of death (2015)
2. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: A review of 3D vessel lumen segmentation techniques: models, features and extraction schemes. Med. Image Anal. **13**(6), 819–845 (2009)
3. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. In: Wells, W.M., Colchester, A., Delp, S. (eds.) MICCAI 1998. LNCS, vol. 1496, pp. 130–137. Springer, Heidelberg (1998). doi:10.1007/BFb0056195
4. Sato, Y., Nakajima, S., Atsumi, H., Koller, T., Gerig, G., Yoshida, S., Kikinis, R.: 3D multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. In: Troccaz, J., Grimson, E., Mösges, R. (eds.) CVRMed/MRCAS -1997. LNCS, vol. 1205, pp. 213–222. Springer, Heidelberg (1997). doi:10.1007/BFb0029240
5. Sethian, J.A.: Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision, and Materials Science. Cambridge University Press, Cambridge (1996)
6. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A., van Ginneken, B., Sánchez, C.I.: A survey on deep learning in medical image analysis. Med. Image Anal. **42**, 60–88 (2017)
7. Fu, H., Xu, Y., Wong, D.W.K., Liu, J.: Retinal vessel segmentation via deep learning network and fully-connected conditional random fields. In: IEEE 13th International Symposium on Biomedical Imaging, pp. 698–701. IEEE (2016)

8. Smistad, E., Løvstakken, L.: Vessel detection in ultrasound images using deep convolutional neural networks. In: Carneiro, G., Mateus, D., Peter, L., Bradley, A., Tavares, J.M.R.S., Belagiannis, V., Papa, J.P., Nascimento, J.C., Loog, M., Lu, Z., Cardoso, J.S., Cornebise, J. (eds.) LABELS/DLMIA -2016. LNCS, vol. 10008, pp. 30–38. Springer, Cham (2016). doi:10.1007/978-3-319-46976-8_4

9. Kitrungrotsakul, T., Han, X.H., Iwamoto, Y., Foruzan, A.H., Lin, L., Chen, Y.W.: Robust hepatic vessel segmentation using multi deep convolution network. In: SPIE Medical Imaging, International Society for Optics and Photonics, pp. 1013711–1013711 (2017)

10. Saloner, D.: The AAPM/RSNA physics tutorial for residents. an introduction to MR angiography. Radiographics **15**(2), 453–465 (1995)

11. Forkert, N., Fiehler, J., Suniaga, S., Wersching, H., Knecht, S., Kemmling, A., et al.: A statistical cerebroarterial atlas derived from 700 MRA datasets. Methods Inf. Med. **52**(6), 467–474 (2013)

12. Wu, A., Xu, Z., Gao, M., Buty, M., Mollura, D.J.: Deep vessel tracking: a generalized probabilistic approach via deep learning. In: IEEE 13th International Symposium on Biomedical Imaging, pp. 1363–1367. IEEE (2016)

13. Kholmovski, E.G., Alexander, A.L., Parker, D.L.: Correction of slab boundary artifact using histogram matching. J. Magn. Reson. Imaging **15**(5), 610–617 (2002)

14. Sled, J.G., Zijdenbos, A.P., Evans, A.C.: A nonparametric method for automatic correction of intensity nonuniformity in MRI data. IEEE Trans. Med. Imaging **17**(1), 87–97 (1998)

15. Forkert, N., Säring, D., Fiehler, J., Illies, T., Möller, D., Handels, H., et al.: Automatic brain segmentation in time-of-flight MRA images. Methods Inf. Med. **48**(5), 399–407 (2009)

16. Dice, L.R.: Measures of the amount of ecologic association between species. Ecology **26**(3), 297–302 (1945)

17. Bergstra, J., Breuleux, O., Bastien, F., Lamblin, P., Pascanu, R., Desjardins, G., Turian, J., Warde-Farley, D., Bengio, Y.: Theano: a CPU and GPU math compiler in Python. In: Proceedings of the 9th Python in Science Conference, pp. 1–7 (2010)

18. Chetlur, S., Woolley, C., Vandermersch, P., Cohen, J., Tran, J., Catanzaro, B., Shelhamer, E.: cuDNN: efficient primitives for deep learning. arXiv preprint arXiv:1410.0759 (2014)