

# Comparison of Automatic Vessel Segmentation Techniques for Whole Body Magnetic Resonance Angiography with Limited Ground Truth Data

Andrew McNeil<sup>1</sup>(✉), Giulio Degano<sup>1</sup>, Ian Poole<sup>2</sup>, Graeme Houston<sup>3</sup>,  
and Emanuele Trucco<sup>1</sup>

<sup>1</sup> CVIP, Computing, School of Science and Engineering,  
University of Dundee, Dundee, UK  
aymcneil@dundee.ac.uk

<sup>2</sup> Toshiba Medical Visualization Systems Europe, Edinburgh, UK

<sup>3</sup> School of Medicine, Ninewells Hospital and Medical School, Dundee, UK

**Abstract.** This work is part of a project aimed at automatically detecting vascular disease in whole body magnetic resonance angiograms (WBMRA). Here we present a comparison of four techniques for automatic artery segmentation in WBMRA data volumes; active contours, two “vesselness” filter approaches (the Frangi filter and Optimally Oriented Flux (OOF)) and a convolutional neural network (Convnet) trained for voxel-wise classification. Their performance was assessed on three manually segmented WBMRA datasets, comparing the maximum Dice Similarity Coefficient (DSC) achieved by each method. Our results show that, in the presence of limited training data, OOF performs best for our three patients, achieving a mean DSC of 0.71 across all patients. By comparison, the 3D Convnet achieved a mean DSC of 0.63. We discuss the potential reasons for these differences, and the implications it has for the automated segmentation of arteries in large WBMRA datasets, where ground truth data is often limited and there are currently no pre-trained 3D Convnet models available, requiring models to be trained from scratch. To the best of our knowledge this is the first comparison of these automated vessel segmentation techniques for WBMRA data, and the first quantitative results of applying a Convnet to vessel segmentation in WBMRA, for which no public sets of manually annotated vascular networks currently exist.

## 1 Introduction

Contrast-enhanced Whole-Body Magnetic Resonance Angiography (WBMRA) is performed by injecting a contrast agent and acquiring images using an MRI scanner as it passes through the arteries of interest [14]. This technique generates high contrast in the lumen (the channel where the blood is flowing), providing a non-invasive, comprehensive imaging method for assessing cardiovascular disease (CVD) throughout the entire body [19]. Analysing these large datasets is

very labour-intensive however, and thus there is a great need for automated, quantitative analysis tools to help stage the disease from these scans.


The first stage of any such system is to locate and segment the arteries of interest. The segmentation of vascular structures is a common task to many medical applications [3, 8, 13], and is a fundamental step in the the quantification of pathologies such as stenoses.

Many vessel segmentation techniques have been proposed in the literature, as explored in [8, 12]. In this work we examine three commonly used techniques—active contours and two “vesselness” filters—comparing their results against a more recent approach using a convolutional neural network (Convnet), structured as a voxel-wise binary classifier following the network structures explored in [21].

## 2 Materials and Methods

### 2.1 Patient Data and Ground Truth

The data used in this study consists of three whole-body datasets, each of which are split into four “stations”; station one comprised the head and neck, station 2 the thorax and abdomen, station 3 the pelvis and thighs, and station 4 the feet (see Fig. 1). These were acquired at Ninewells Hospital in Dundee, UK, using a 3.0 Tesla MRI scanner (Siemens Magnetom Trio).



**Fig. 1.** Maximum intensity projections of the four stations of patient 1, shown after digital subtraction of the pre-contrast from the post-contrast volumes.

The arteries in each individual station were manually segmented by a trained observer using the software package “3DSlicer” [4]. The following pre-processing steps were applied to aid visualisation during the manual segmentation; firstly all volumes were scaled in one direction due to varying slice thicknesses, giving an isotropic voxel size of  $0.98\text{ mm}^3$  (the raw slice thicknesses varied from  $0.98\text{ mm}$ – $1.3\text{ mm}$ ). Next the pre-contrast volume was registered to the post-contrast volume using the mutual information similarity measure. Subtracting the pre-contrast from the post-contrast volume was done to suppress static tissues. To aid with visualisation, an intensity equalisation step was applied in the

axial direction, ensuring a consistent vessel intensity across the entirety of each volume. A final artefact correction step was applied, masking border voxels to remove MR artefacts and tissues not covered by the pre-contrast volume.

## 2.2 Active Contours

Segmentation using active contours, where an initial curve is evolved using a cost function depending on local gradients (external forces) and shape constraints (internal forces), was first proposed in [7] and has been successfully applied to many segmentation problems [1, 12]. In general, active contour methods are based on image gradient, detecting edges and generating well-defined boundaries on which to evaluate the internal and external energy.

The level set model tries to solve the optimisation problem by embedding the active contour as a constant set (zero level) in a function  $\phi$  that evolves in time with speed  $S$ .

For our comparative study we chose the classic Chan-Vese model [1]. It has been applied to the segmentation of objects whose edges are not well defined by the gradient, and has a well defined implementation for 3D segmentation as described in [25].

The Chan-Vese model is formulated as a “mean-curvature flow”-like evolving active contour, where the stopping term depends not on the gradient of the image, as in classical active contour models, but is instead related to a particular segmentation of the image [1].

For 3D data, we define the bounded domain  $\Omega \in \mathbb{R}^n$  (in our case  $n = 3$ ), and the bounded image function  $I : \Omega \rightarrow \mathbb{R}$ .  $\Omega$  can be divided into a set of connected domains by a curve  $C$  by  $\Omega - C = \cup_{i \in I} \Omega_i$ . We then define two different regions  $R_1 = \cup_{i \in I_1} \Omega_i$  and  $R_2 = \cup_{i \in I_2} \Omega_i$  that represent the object support and the background support respectively. The final energy functional  $E(\phi, \mu_1, \mu_2)$  is then given by

$$E(\phi, \mu_1, \mu_2) = \lambda_1 \int_{\Omega} (I - \mu_1)^2 H(\phi) d\Omega + \lambda_2 \int_{\Omega} (I - \mu_2)^2 (1 - H(\phi)) d\Omega + \alpha \int_{\Omega} H(\phi) d\Omega + \beta \int_{\Omega} |\nabla H(\phi)| d\Omega \quad (1)$$

where  $\mu_1$  and  $\mu_2$  represent the mean value of the object support region and background support region of image  $I$  respectively, and  $H(\phi)$  is the Heaviside function. Here, the first two terms measure the variations inside and outside the active contour, the third term measures the area inside the contour and the fourth term measures the length of the contour [25].

## 2.3 Vessel Enhancement Filters

**Frangi Filter.** This method of enhancing vessel-like structures is based on calculate the local curvature by analysing the Hessian function, so as to extract the main directions in which the local structure can be decomposed [5].

To derive the “vesselness” function we first define  $\lambda_k$  as being the eigenvalue with the  $k$ -th smallest magnitude, i.e. ( $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$ ). Therefore, for an ideal tubular structure in a 3D image

$$|\lambda_1| \approx 0, \quad |\lambda_1| \ll |\lambda_2|, \quad |\lambda_2| \approx |\lambda_3| \quad (2)$$

In other words, the curvature should be large in the two directions (around the circumference of the vessel), and very small along the length of the vessel.

The final vesselness function to be evaluated was defined in [5] as being

$$V_0(s) = \begin{cases} 0 & \text{if } \lambda_2 > 0 \text{ or } \lambda_3 > 0 \\ \left(1 - \exp\left(-\frac{R_A^2}{2\alpha^2}\right)\right) \exp\left(-\frac{R_B^2}{2\beta^2}\right) \left(1 - \exp\left(-\frac{S^2}{2c^2}\right)\right) & \text{otherwise} \end{cases} \quad (3)$$

where  $R_A = |\lambda_2|/|\lambda_3|$ ,  $R_B = |\lambda_1|/\sqrt{|\lambda_2\lambda_3|}$ ,  $S = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}$ . Here,  $\alpha$ ,  $\beta$  and  $c$  are thresholds which control the sensitivity of the filter to the measures  $R_A$ ,  $R_B$  and  $S$ .

For the application of enhancing vessels in a 3D angiographic dataset, the vesselness measure in Eq. (3) is analysed at different scales,  $s$ , corresponding to the sigma of the Gaussian kernel used in the construction of the Hessian matrix. It logically follows that the response of the filter will be maximised at the scale which approximately matches the size of the vessel in that region. Therefore, the final estimate of vesselness is obtained by integrating the vesselness measure provided by the filter response at different scales,

$$V_0(\gamma) = \max_{s_{min} \leq s \leq s_{max}} V_0(s, \gamma) \quad (4)$$

where  $s_{min}$  and  $s_{max}$  are the minimum and maximum scales at which relevant structures are expected to be found, chosen so that they cover the range of relevant vessel widths. The vesselness map can then be thresholded to provide a binary vessel tree.

**Optimally Oriented Flux.** The “optimally oriented flux” filter, first published in [10], evaluates a scalar measure of the flux flowing through a spherical surface. Before computing this value, directional information is extracted by projecting the gradient along “optimal” axes, and the flux measure then evaluated. For each voxel a sphere with variable radius is built, centred on the voxel, which produces an “OOF response” when touching an object edge. If the voxel is inside the curvilinear structure the response will be positive, otherwise it will be negative.

The outwardly oriented flux along the direction  $\hat{\rho}$  is firstly computed by projecting the gradient of the image  $\mathbf{v}$  along  $\hat{\rho}$ , with the flux then evaluated through the spherical region  $S_r$  with radius  $r$  using the definition

$$f(\mathbf{x}; \mathbf{r}, \hat{\rho}) = \int_{\delta S_r} ((\mathbf{v}(\mathbf{x} + \mathbf{h}) \cdot \hat{\rho})\hat{\rho}) \cdot \hat{n} dA \quad (5)$$

where  $dA$  is the infinitesimal area of  $S_r$ ,  $\hat{n}$  is the unit normal to the surface at position  $h = r\hat{n}$ .

As before, the goal is to obtain the principal eigenvalues for each voxel. Inside the vessel, when the local spherical region with surface  $S_r$  touches the boundaries of the object  $\mathbf{v}$  is aligned opposite to the direction of  $\hat{n}$ , therefore the eigenvalues  $\lambda_1 \leq \lambda_2 \ll 0$ . The gradient of the image will be perpendicular to the direction of the curvilinear structure, with a value of  $\lambda_3 \approx 0$ . In the case where the voxel is in the background,  $\mathbf{v}$  will have the same direction as  $\hat{n}$ , and therefore  $\lambda_3 \gg 0$ .

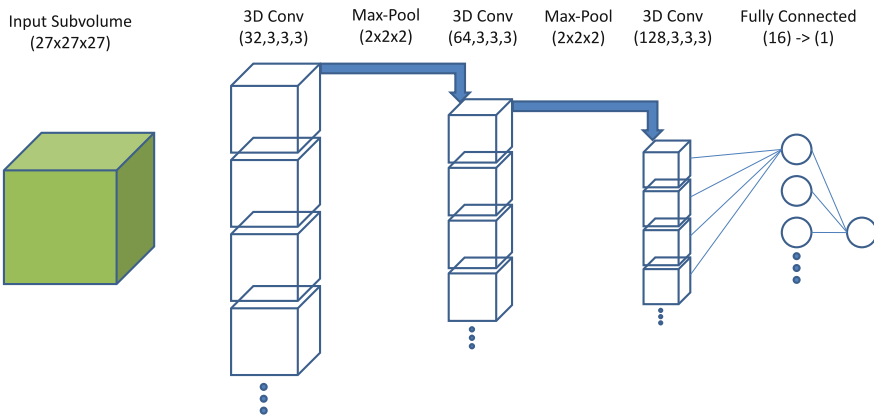
To obtain the maximum response to the OOF while changing the radius  $r$ , we evaluate of the geometric mean of the eigenvalues, as

$$M(\mathbf{x}; s) = \begin{cases} \sqrt{|\lambda_1(\mathbf{x}, s)\lambda_2(\mathbf{x}, s)|} & \lambda_1(\mathbf{x}, s) \leq \lambda_2(\mathbf{x}, s) < 0 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where  $s$  represents the scale factor. Similar to the Frangi approach, evaluating the maximum response over an appropriate range of scales generates the final map, which can be thresholded to produce a final segmentation.

## 2.4 Convolutional Neural Network

In recent years, deep Convnet approaches have been driving advances in many computer vision tasks, such as image classification [9, 21] and image segmentation [18, 20]. Many network models have been developed for these tasks, and it is a very active area of research [11]. The network structure we chose was inspired by those explored in [21], and recently applied to segmentation tasks in MRI [2, 15, 17]. To the best of our knowledge, this is the reported results of applying a Convnet to vessel segmentation in WBMRA, for which no public sets of manually annotated vascular networks currently exist.



**Fig. 2.** Structure of the 3D Convnet segmentation network. All layer activation functions were “ReLU” except the final output node, which was “sigmoid”. The “Adam” optimiser was used during training, with “binary cross-entropy” selected as the loss function

The final network structure is shown in Fig. 2, consisting of five layers; 3 sets of convolutional and max-pooling layers, followed by two fully connected layers. The output node of the final layer gives a single binary output of vessel/non-vessel for the central voxel of the input patch.

Our network was implemented using Keras v1.1.0 and Theano v0.8.2. All layer activation functions were “ReLU” except the final output node, which was “sigmoid”. The “Adam” optimiser was used during training, with “binary cross-entropy” selected as the loss function [6].

While this approach of a voxel-wise classifier has been shown to be less computationally efficient than a fully convolutional network [20], it allowed fine control over dataset balancing for our limited amount of ground truth data, and the fully connected layer gave additional flexibility to the network without increasing the required input volume size, which is inherent to the operation of convolutional layers.

## 2.5 Comparison Metric

There are many metrics used for evaluating the quality of segmentation in medical images [24]. For our data, we have selected the Dice Similarity Coefficient (DSC—also referred to as the F1-Measure). This is given by

$$DSC = \frac{2|X \cap Y|}{|X| + |Y|} \quad (7)$$


where  $|X|$  is the number of all the vessel voxels in the segmentation obtained by the tested method and  $|Y|$  is the number of all the vessel voxels in the ground truth.

## 2.6 Pre-processing and Parameter Optimisation

**Pre-processing.** As can be seen in Fig. 1, the subtraction of the pre-contrast from the post-contrast volume still leaves some tissues and non-arterial structures behind, particularly in stations 1 and 2. The most problematic of these are the lungs in station 1, which contain vessels which were not included in the manual ground truth. For this reason, the small region around the lungs and heart were masked out in both the original volumes and the ground truth data, excluding this area from our analyses.

Another artefact which remained after subtraction was the variation of lumen intensity along the length of vessels. These may arise due to poor timing of the contrast agent during acquisition, or by inhomogeneities in the magnetic field (such as surface coil artefacts). A simple procedure was followed to correct these variations in each station, in the form of a local intensity normalisation.

First we make the assumption that each axial slice contains an artery, and they are the brightest objects present (which holds true for all regions except for slices above the head and below the feet—these slices were simply masked to zero after the procedure was applied). We then applied a 7-slice sliding window



**Fig. 3.** Results of intensity equalisation on station 3 of patient 1. The MIP of the raw volume is show on the left, and the equalised volume on the right.

axially, in which the local vessel intensity was estimated from its histogram by choosing the highest frequency bin above 70% of the maximum intensity, with this value corresponding to the vessel intensity estimate for the central slice. Once calculated for the entire volume, Gaussian smoothing of the values was applied and then each slice divided by it's corresponding estimate. An example of the results of this processing is show in Fig. 3.

**Active Contours.** For the active contour method, the Toolbox implementation provided by [25] was used. Values for  $\beta$ ,  $\Delta t$ ,  $\lambda$  were fixed using a grid search optimisation procedure across all patients. In our case, the optimal values were found to be 0.08 for the smoothing weight term, 0.0002 for the image weight term, and 2.72 for the time step.

The final step was the initialisation of  $\phi_0$ . This choice was critical as it affects the time and the speed of the evolution of the curve. So again under the hypothesis that the highest intensity voxels belong to the vessels, we took a set of seed points with high grey levels as  $\phi_0$ . To keep the process completely automatic we used Otsu's method for generating thresholds from grey-level histograms [16].

For each station we generated 10 thresholds, which served as 10 different sets of seed points. The active contour method was then applied using the above parameters, and the highest Dice score recorded.

**Enhancement Filters.** The optimal parameters for the enhancement filters are shown in Table 1. These were optimised for each station across all patients using a grid search, with a fixed segmentation threshold.

The final segmentations acquired by calculating the vesselness map using the parameters in Table 1, then 20 thresholds automatically calculated using Otsu's

**Table 1.** Enhancement filter parameters. The scale factor and radii values are written in the form *minimum:step:maximum*.

Station number	SF/Radii	$\alpha$ (Frangi)	$\beta$ (Frangi)	$\sigma$ (OOF)
1	1:1:15	0.5	0.5	0.4
2	1:1:10	0.5	0.5	0.5
3	3:0.25:5	0.5	0.5	0.5
4	1:0.5:4	0.5	0.5	0.4

method [16]. The highest Dice score achieved from all 20 segmentation maps was then recorded.

**3D Convolutional Neural Network.** A number of network structures were explored during optimisation of the network structure. Inspired by models discussed in [20, 21], we trained models consisting of 2–6 convolutional layers with 16–128  $3 \times 3 \times 3$  kernels, 1–3 max-pooling layers, and 1–2 fully-connected layers. To help combat overfitting,  $l_2$  weight regularisation was used for each convolutional layer [6], and 20% dropout used on the fully connected layers [22]. All layer weights were initialised from a scaled Gaussian distribution.

A single network was trained for each station, with training patches extracted from two patients and the trained model applied to the held-out third patient in a 3-fold cross-validation setup.

The cubic patches were varied in size according to the network structure used, based on ensuring that the deepest layer still received a patch large enough to perform meaningful calculations on. A minimum side length of 15 voxels was needed to capture the thickest vessels, leading to a side length range of 15–50 voxels for the network structures we explored. For our final network, a patch size of  $27 \times 27 \times 27$  was found to be optimal.

Finally, the number of training patches was chosen to maximise the available data. For each station, the minimum number of ground truth vessel voxels across all 3 patients was calculated, and this used as the number of positive samples to be extracted from each patient.

The data was balanced by extracting an equal amount of background samples. The position of the background samples were weighted to have two-thirds from regions within 5 voxels of a vessel and one third sampled randomly from the rest of the volume. This was found to improve the networks tendency to over estimate the diameter of the vessels when the background patches were sampled completely at random.


The total number of training samples used for each station were 136000 for station 1, 160000 for station 2, 28000 for station 3, and finally 24000 for station 4. During training, 5% of the training was data held out for validation, and the best network weights saved as those giving the highest validation accuracy score after 20 epochs (the network performance was found to typically converge after 8–12 epochs).



The networks were trained using an Nvidia Titan X Pascal GPU, with training times of between 3–4 h for each model (depending on the station and number of training samples used).

### 3 Results

The results of applying the four automated strategies described in Sect. 2 to all stations and patients are shown in Table 2, with the corresponding segmentation results of patient 1 shown in Fig. 4.



**Fig. 4.** Segmentation results for patient 1, shown as coronal projections.

It can be seen from Table 2 that for our case of only three patients, the OOF filter achieved the greatest mean DSC of 0.705. The 3D Convnet typically

**Table 2.** Dice coefficients for each method

Patient (station)	Level set	Frangi	OOF	3D convnet
1 (1)	0.785	0.697	<b>0.795</b>	0.727
1 (2)	<b>0.837</b>	0.817	0.812	0.618
1 (3)	0.647	<b>0.702</b>	0.657	0.605
1 (4)	0.672	0.780	<b>0.843</b>	0.806
2 (1)	<b>0.803</b>	0.689	0.794	0.752
2 (2)	0.712	<b>0.722</b>	0.674	0.552
2 (3)	<b>0.353</b>	0.323	0.281	0.303
2 (4)	0.662	0.776	<b>0.844</b>	0.692
3 (1)	0.572	<b>0.665</b>	0.663	0.567
3 (2)	<b>0.504</b>	0.488	0.469	0.422
3 (3)	0.803	0.747	0.835	<b>0.845</b>
3 (4)	0.618	<b>0.843</b>	0.791	0.697
Mean DSC	0.660	0.690	<b>0.705</b>	0.632

outperforms at least one of the other techniques, except for station 2. The main reason for this appears to be because of the additional artefacts left over from the imperfect volume registration and subtraction procedure (particularly the kidneys and bladder). The network approach had the most difficulty distinguishing between these artefacts and the arteries, causing it to over-segment station 2, resulting in a lower Dice score.

Looking at the segmentation results in Fig. 4, a number of observations can be made. The active contour method often produces broken vessels, such as the right branch in station 2, and has the most difficulty segmenting the finest vessels in station 4.

The Frangi and OOF enhancement filters produce visually similar results, though the OOF performs better at rejecting non-vessel artefacts (most noticeably in the brain and abdomen of station 1). Both filters do exhibit difficulties segmenting the finest vessels, such as at the bottom of station 2, and in different cases tend to either underestimate (Frangi in station 2, OOF in station 4) or overestimate (Frangi and OOF in station 3) the true diameter of the vessels as compared to the ground truth.

The Convnet performs poorest at rejecting non-arterial artefacts in station 2, but it also has the highest sensitivity to extracting fine vessels. Indeed, some of the finest vessels in the lower half of stations 2 and 3 were not present in the ground truth, having been either overlooked or rejected due to low contrast.

## 4 Conclusions

In this paper we have presented a quantitative comparison between four automated vessel segmentation techniques for whole-body MRA data, using three manually segmented patient datasets.

In this regime of having limited ground truth data, it has been found that the Optimally Oriented Flux filter provides the best average DSC of 0.705. Visually, the Convnet approach segments vessels most consistently, with the least number of breaks, picking up finer vessels, and having the most consistently accurate diameters when compared with the ground truth. However it performed poorest at rejecting non-arterial artefacts, resulting in a lower DSC overall. It was also noted that some of the fine vessels segmented by the Convnet were not present in our ground truth. Due to having ground truth from a single observer, we are unable to estimate the quality and reliability of the ground truth data, and therefore the impact of this on the DSC results cannot be easily estimated for our data. We are not aware of any publicly available sets of manually annotated vascular networks for WBMRA volumes.

The Convnet approach appears to be mainly limited by the lack of training data. Other deep learning approaches which integrate large amounts of data augmentation, such as U-Net [18], may achieve better results at rejecting non-arterial artefacts, however those techniques have not been explored here. Another approach often used is to fine-tune a previously trained network such as GoogLeNet, a 22 layer network trained on a database of 1 million natural images [23]. Given that currently no pre-trained 3D networks exist for medical data, the training of deeper networks will require a larger database of ground truth segmentations.

## References

1. Chan, T.F., Vese, L.A.: Active contours without edges. *IEEE Trans. Image Process.* **10**(2), 266–277 (2001)
2. de Brbisson, A., Montana, G.: Deep neural networks for anatomical brain segmentation. In: 2015 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 20–28, June 2015
3. Dehkordi, M.T., Sadri, S., Doosthoseini, A.: A review of coronary vessel segmentation algorithms. *J. Med. Sig. Sens.* **1**(1), 49 (2011)
4. Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., Fillion-Robin, J.-C., Pujol, S., Bauer, C., Jennings, D., Fennessy, F., Sonka, M., Buatti, J., Aylward, S., Miller, J., Pieper, S., Kikinis, R.: 3D slicer as an image computing platform for the quantitative imaging network. *Magn. Reson. Imaging* **30**(9), 1323–1341 (2012)
5. 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](https://doi.org/10.1007/BFb0056195)
6. Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press, Cambridge (2016). <http://www.deeplearningbook.org>
7. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: active contour models. *Int. J. Comput. Vis.* **1**(4), 321–331 (1988)

8. Kirbas, C., Quek, F.: A review of vessel extraction techniques and algorithms. *ACM Comput. Surv.* **36**(2), 81–121 (2004)
9. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: Pereira, F., Burges, C.J.C., Bottou, L., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 25, pp. 1097–1105. Curran Associates Inc. (2012)
10. Law, M.W.K., Chung, A.C.S.: Three dimensional curvilinear structure detection using optimally oriented flux. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008*. LNCS, vol. 5305, pp. 368–382. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-88693-8\\_27](https://doi.org/10.1007/978-3-540-88693-8_27)
11. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015)
12. 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)
13. Lupaşcu, C.A., Tegolo, D., Trucco, E.: Accurate estimation of retinal vessel width using bagged decision trees and an extended multiresolution hermite model. *Med. Image Anal.* **17**(8), 1164–1180 (2013)
14. McRobbie, D.W., Moore, E.A., Graves, M.J., Prince, M.R.: *MRI from Picture to Proton*, 2nd edn. Cambridge University Press, Cambridge (2007)
15. Moeskops, P., Viergever, M.A., Mendrik, A.M., de Vries, L.S., Benders, M.J.N.L., Igum, I.: Automatic segmentation of MR brain images with a convolutional neural network. *IEEE Trans. Med. Imaging* **35**(5), 1252–1261 (2016)
16. Otsu, N.: A threshold selection method from gray-level histograms. *IEEE Trans. Syst. Man Cybern.* **9**(1), 62–66 (1979)
17. Pereira, S., Pinto, A., Alves, V., Silva, C.A.: Brain tumor segmentation using convolutional neural networks in MRI images. *IEEE Trans. Med. Imaging* **35**(5), 1240–1251 (2016)
18. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation, pp. 234–241. Springer International Publishing, Cham (2015)
19. Ruehm, S.G., Goehde, S.C., Goyen, M.: Whole body MR angiography screening. *Int. J. Cardiovasc. Imaging* **20**(6), 587–591 (2004)
20. Shelhamer, E., Long, J., Darrell, T.: Fully convolutional networks for semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(4), 640–651 (2017)
21. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556 (2014)
22. Srivastava, N., Hinton, G.E., 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)
23. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1–9, June 2015
24. Taha, A.A., Hanbury, A.: Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool. *BMC Med. Imaging* **15**(1), 29 (2015)
25. Zhang, Y., Matuszewski, B.J., Shark, L.K., Moore, C.J.: Medical image segmentation using new hybrid level-set method. In: *2008 5th International Conference BioMedical Visualization: Information Visualization in Medical and Biomedical Informatics*, pp. 71–76, July 2008