Stroke Thrombus Segmentation on SWAN with Muliti-Directional U-Nets

J. Kobold, V. Vigneron,
H. Maaref and D. Fourer
IBISC, EA 4526
Univ Evry, Université Paris-Saclay, France

M. Aghasaryan, C. Alecu, N. Chausson,
Y. L'Hermitte and D. Smadja
Centre Hospitalier Sud-Francilien
Evry-Corbeil, France

E. Lang
CIML Group
Universität Regensburg, Germany

A.M. Tome
IEETA
Universidade de Aveiro, Portugal

Abstract—The thrombus causing a stroke can be seen on the susceptibility weighted angiography (SWAN) magnetic resonance imaging (MRI) sequence. But it is very small and hard to detect by humans. Up to date the thrombus is identified by trained human experts. But as stroke needs quick treatment, an automatic detection of the thrombus would be useful to speed up the diagnosis of acute stroke. We propose a method for automatic thrombus detection from SWAN using three separate U-Nets which work on the axial, coronal and sagittal planes.

Index Terms—Stroke, Thrombus, Deep Learning, MRI, Automatic Segmentation.

I. Introduction

According to the World Health Organisation, stroke is the second leading cause of death worldwide. This disease of the brain is so common, and at the same time so difficult to treat, that even an entire journal was created exclusively for research on stroke[1]. Stroke can be divided into two subtypes which have a different cause: hemorrhagic stroke and ischemic stroke. Only the more common ischemic stroke is considered in the following. In ischemic stroke a blood clot, the thrombus, blocks an artery in the brain. As a consequence a part of the brain is cut off from the blood supply and quickly dies, mainly within the first six hours after stroke onset. The damaged area is called *lesion*.

One of the vital parts of stroke diagnosis is the identification and localisation of the thrombus. Its size and location determines the prospects of success of possible treatment strategies [1].

Some semi-automatic methods have been proposed to segment the thrombus on computed tomography (CT) images of the brain [2], [3], [4], but they require the user to indicate the spatial location of the thrombus in the images. To the best of our knowledge, there is currently no publication on thrombus segmentation or identification on stroke MRI. Manually identifying the thrombus on stroke MRI is extremely difficult, due to its small size and high variability in spacial location and colour, and thus requires longstanding clinical experience. Furthermore the treatment decision for a stroke patient arriving at a hospital needs to be taken fast, as any

[1]See http://stroke.ahajournals.org/

delay reduces the patient's chances of survival. A computer guided assistance in thrombus identification would therefore be helpful to gain some time in the diagnosis pipeline.

In this paper we propose a method for *automatic thrombus segmentation*.

II. Method and Materials

A. Data-Set

The data-set for this retrospective study consists of MRIs of 61 patients which were treated for stroke at the centre hospitalier sud-francilien (CHSF). The MRIs were taken upon the patients' arrival at the hospital before treatment, which means in the hyper acute phase of stroke. All patients have an occlusion in the middle cerebral artery visible on SWAN. Multiple MRI sequences were available from the CHSF stroke MRI protocol but as the thrombus can only be seen on SWAN only the latter was used in the following. SWAN is a susceptibility weighted imaging sequence tied to the manufacturer General Electrics. Very similar sequences are available from all major MRI manufacturers (SWI by Siemens, SWIp by Phillips, BSI by Hitachi and FSBB by Canon) which produce almost identical images. Thus we expect our method to generalise well to MRIs of different manufacturers.

The images are from $1.5\,\mathrm{T}$ and $3\,\mathrm{T}$ General Electrics MRI machines. The SWAN was acquired with an echo time of $20\,\mathrm{ms}$, a repetition time of $48\,\mathrm{ms}$ and a flip angle of $20°$. The slice thickness is $2.8\,\mathrm{mm}$ and the pixel size $0.61\,\mathrm{mm} \times 1.07\,\mathrm{mm}$.

The Thrombus was segmented manually on each SWAN by at least two experienced neurologists. In total 5 neurologists carried out segmentations. Their inter-observer agreement was measured using the dice coefficient [5] and they achieved a dice of 0.7 with a mean size of 1166 voxels. For such small volumes this is a very good inter observer agreement.

MRI is a qualitative and no quantitative measurement [6], *i.e.* the intensity values and range of repeated measurements of the same patient are different even though it appears to be the same image to human eyes. Across patients this variability is even higher (compare figures 1 a) and c)). Therefore a
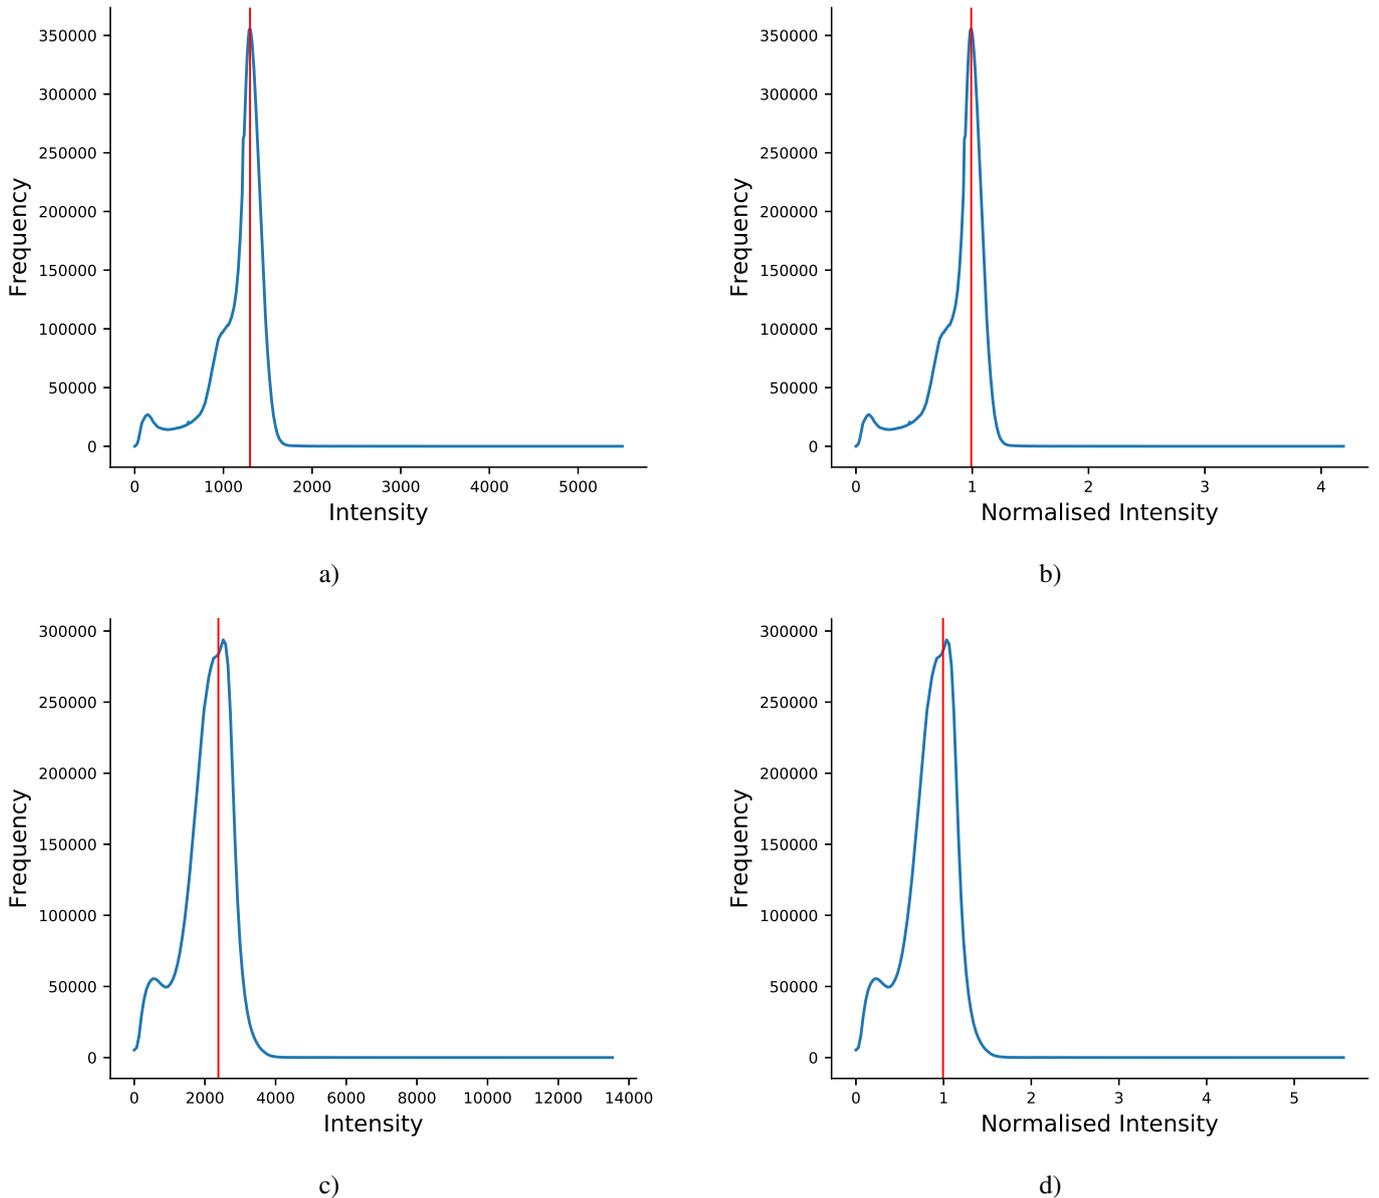
Figure 1. Two examples of the intensity normalisation of the SWAN histogram. In the original histograms a) and c) the location of the peak corresponding to the normal brain tissue (red line) is noted. Both histograms show this peak, but its location is different. This illustrates why intensity normalisation of MRI is important. Then the intensity values are divided by the peak intensity value and the normalised histograms b) and d) are obtained. Note that this histogram shows intensity values of the brain only, the very large peak around zero of the empty space around the head has been removed for readability.

normalisation of the images is necessary to obtain repeatable results with automatic processing of the images. We normalise the images such that the value of the normal brain tissue has always the same value of 1, assuming that intensities of other parts of the brain behave the same relative to the normal brain tissue across patients. This is a simplified version of the method by Nyul et al. [7]. Figure 1 gives an example of the normalisation process, the location of the peaks was determined by fitting a mixture of two Gaussians to the histogram.

B. Objective

Segmentation of the thrombus is a very difficult task. It cannot be seen directly on the MRI images. Instead its presence is indicated by a secondary effect. The thrombus is constituted by clotted blood, which contains a lot of densely packed haemoglobin proteins which in turn contain iron atoms. This high local iron concentration causes a defect in the magnetic field of the MRI. Most MRI sequences compensate for magnetic field defects which renders the thrombus invisible. But on susceptibility weighted imaging sequences this defect is visible as a dark spot. Usual clinical T2* sequences are of too low quality, due to time constraints in a stroke scenario, to see

anything but the largest thrombi. Other susceptibility weighted imaging sequences, for example the SWAN, produce higher quality images in the same time where even the magnetic defects of small thrombi are visible.

Thus the objective of our work is to automatically segment the thrombus on SWAN with the application in the acute phase of stroke in mind. That means our method should be fast as well. The task is complicated by the fact that a lot of dark spots – which look like the thrombus – are visible on SWAN (see Figure 3.a for a comparison).

C. U-Net

Three years after the breakthrough of convolutional neural networks for image classification tasks [8], the U-Net [9] set a new standard for the segmentation of bio-medical images. It was not the first to be applied to bio-medical image segmentation tasks, but it combined multiple ideas which made it outperform its predecessors.

As any classic convolutional neuronal network the U-Net consecutively down-samples the input image in a compressing branch. At this stage features at multiple resolutions are extracted (Figure 2from the left to the middle). The resolution is halved with a maximum pooling layer and the number of features at each resolution is doubled with a convolutional layer.

This is followed by an up-sampling branch which increases the resolution back to the input space (Figure 2 on the right). Notably transposed 2D convolutions are used for up-sampling, *i.e.* the up-sampling operation is learned. The other notable part is the use of *skip* connections which are represented with grey arrows (Figure 2 and correspond to *copy* operation). The idea behind the skip connections is to use the features which were extracted in the down-sampling branch again during the up-sampling in order to preserve fine-grained details. This is one of the key-points which allows this model to produce segmentations with pixel precision.

And finally the classification layer uses 1×1 convolutions adopting the principle of the all convolutional network [10]. This avoids fully connected layers and their drawbacks all together.

The U-Net has been successful in a number of bio-medical image segmentation tasks and works well with small data-sets. We use the (Python) `Tensorflow` implementation of the U-Net by Akeret *et al.* [11].

D. Multi-directional U-Net

The drawback of the native U-Net is that it works on 2D images only. From our experience a single U-Net alone is not able to identify the thrombus (compare figure 3, d)). As confirmed by our medical experts, the 3D context is indeed very important for identifying a thrombus. The U-Net has been extended to 3D [12] but the number of parameters and the evaluation time increases significantly in this case. As we are looking for a fast method we decided to introduce the 3D context by using three distinct U-Nets which are trained on axial, coronal and sagittal slices respectively.

E. Training

The thrombus segmentation task proved to be very difficult, which made a rather long training scheme necessary. First up, the models are trained on 64×64 patches with a batch size of 20. To counter the class imbalance, every second image in a batch shows a thrombus and every other a non thrombus part of the brain. *Data augmentation* was done through image rotations of π, 2π and 3π, random crops and by adding a small random constant to the pixel values. To further increase the available training data with this rather small data-set the models were trained in a *leave-one-out* cross validation scheme [13]. That means for each patient a separate model was trained using all other patients.

The models achieved with this approach detected still too many false positives. To enhance the models they were re-trained on an improved data-set, created as follows: the MRIs of the patients in the training set were classified with the trained model. Then the non thrombus patches in the training set are replaced by non thrombus patches which have been erroneously classified as thrombus by the trained model. The model was retrained on the improved data-set. Thus the model is forced to focus on the differences between true thrombus and false positives. This re-training step greatly reduced the number of false positives.

According to this schedule three U-Nets were trained for each patient. The first one was trained on patches from axial slices, the second was trained on coronal slices and the last was trained on sagittal slices. The models were trained on a NVIDIA GTX 1070 GPU and training time was roughly 1 hour per model (183h in total).

F. Label Merging

The three models provide three different scores s_1, s_2, s_3 for each voxel which need to be merged for a final decision. A simple majority vote would underestimate the size of the thrombus. Thus the label merging is split into two parts. At first possible thrombi are extracted and in the second part these candidates are classified into thrombus and non thrombus.

For the first part a new volume $\mathcal{V}$ is created where each voxel is assigned the maximum $\max(s_1, s_2, s_3)$ out of the three scores. $\mathcal{V}$ is then thresholded at a value of 0.4 and the resulting binary map is divided into connected components.

To decide whether one component should be a thrombus or not, the scores (s_1, s_2, s_3) are used again. If all three scores are bigger than 0.7 for one voxel in the component, *i.e.* $s_1 > 0.7 \wedge s_2 > 0.7 \wedge s_3 > 0.7$, then the component is validated as a thrombus.

III. RESULTS AND DISCUSSION

We have proposed and evaluated a method for automatic thrombus segmentation on SWAN capable to find 77.4% of all thrombi (see table I). The mean dice of the found thrombi of 0.415 lies below the inter observer agreement of 0.7. But the dice coefficient is very sensitive to small differences for objects of the size of a thrombus. For example if the size of a thrombus is overestimated by displacing the border by one
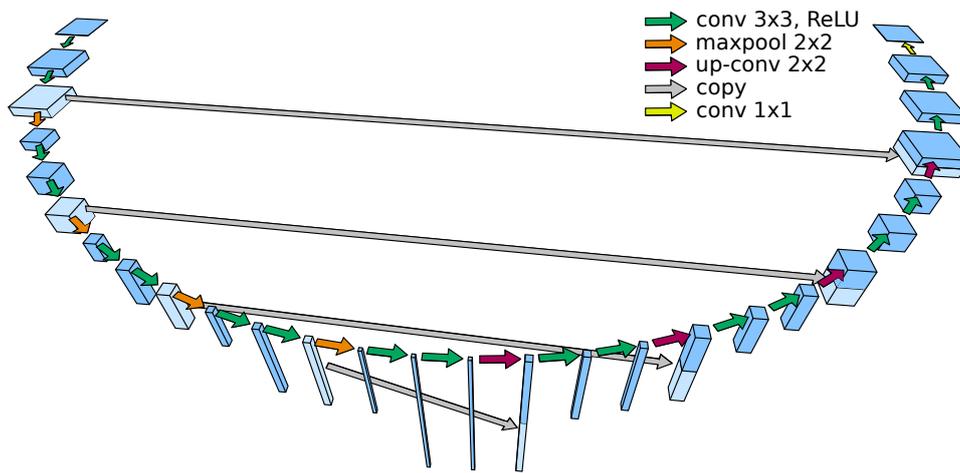
Figure 2. The U-Net architecture from [9]. Arrows represent operations and cubes represent feature maps where the height of the cube stands for the number of feature maps and the width and depth of the cubes for the size of the feature maps.
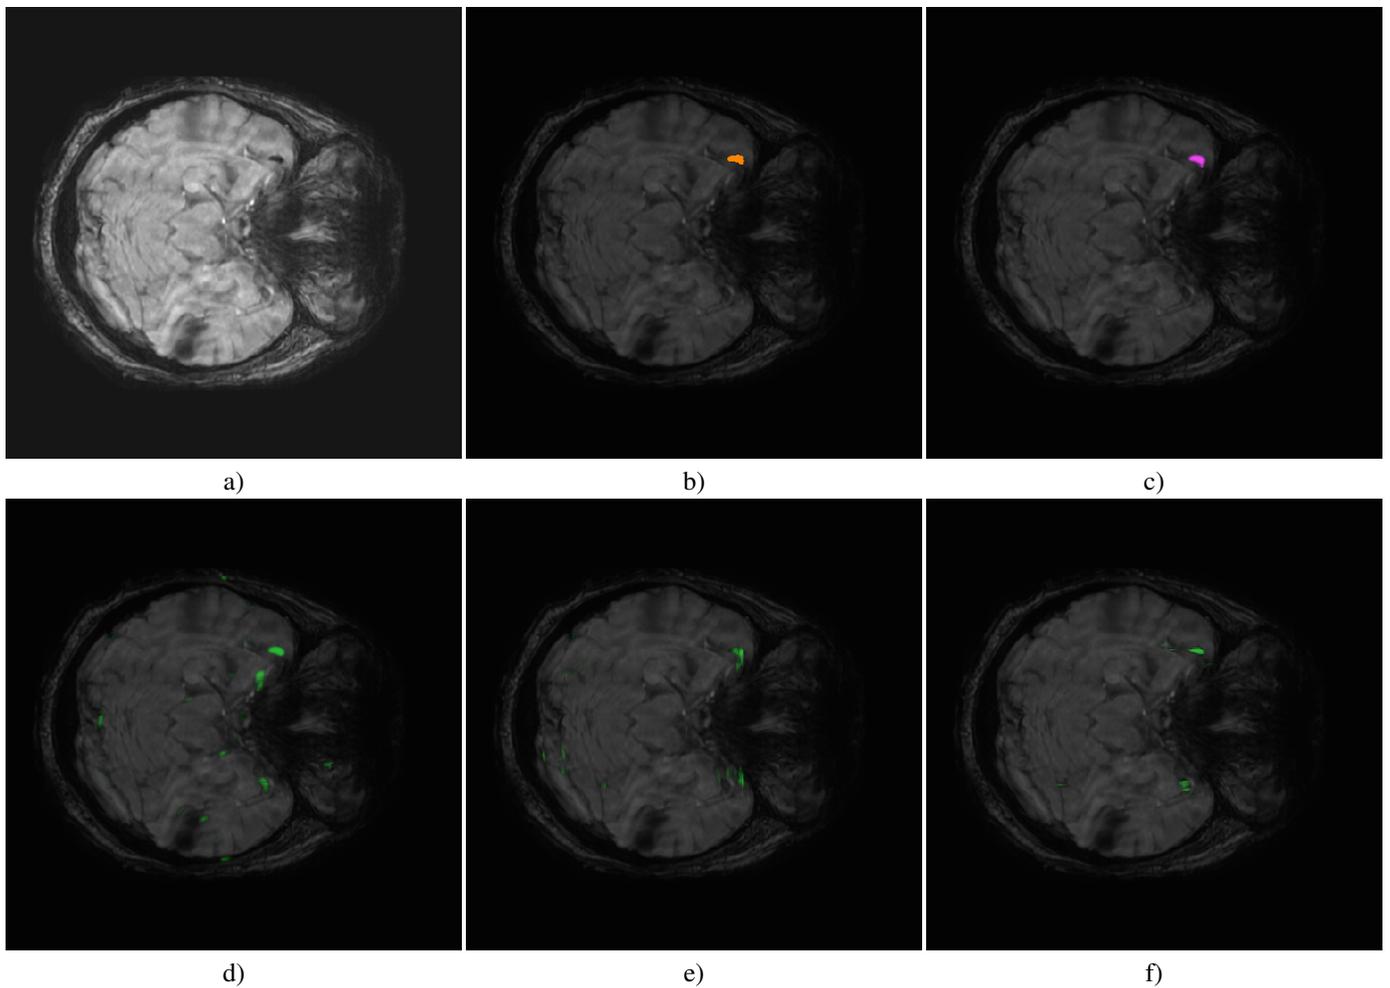

Figure 3. A SWAN image a) showing a thrombus. The thrombus location is given by the manual segmentation b). The segmentation from the merged predictions c) is very close to the ground truth. d), e) and f) are the predictions from the U-Nets in axial, coronal and sagittal direction. Each of the U-Net identifies a couple of the dark spots as a thrombus. But only the true thrombus is found by all three U-Nets, the other spots are always discarded by at least one U-Net. It is also visible that the size of the thrombus is underestimated by two of the U-Nets which made the first part of the label merging step necessary. Figure is best viewed in colour.

Table I
SEGMENTATION RESULTS

Detection Rate	0.774
Mean Dice	0.415
Dice Standard Deviation	0.243
Mean Objects	2.26
Objects Standard Deviation	3.36

pixel outside, the overestimated thrombus is already a few hundred voxels larger. Given the average size of the thrombi of 1166, this would be already a significant change in the dice. So the dice of 0.415 actually means that our method delivers segmentations which are close to the ground truth, up to a few pixels.

Given that the thrombus segmentation task is very difficult and up to now no result on thrombus segmentation from MRI has been published, this is a good first result. For the future we plan to improve the detection rate by incorporating additional clinical data and information about the lesion location into the model.

REFERENCES

[1] Powers William J., Rabinstein Alejandro A., Ackerson Teri, Adeoye Opeolu M., Bambakidis Nicholas C., Becker Kyra, Biller José, Brown Michael, Demaerschalk Bart M., Hoh Brian, Jauch Edward C., Kidwell Chelsea S., Leslie-Mazwi Thabele M., Ovbiagele Bruce, Scott Phillip A., Sheth Kevin N., Southerland Andrew M., Summers Deborah V., and Tirschwell David L., "2018 guidelines for the early management of patients with acute ischemic stroke: A guideline for healthcare professionals from the american heart association/american stroke association," *Stroke*, vol. 49, no. 3, Jan. 2018. [Online]. Available: http://stroke.ahajournals.org/content/early/2018/01/23/STR.0000000000000158.abstract

[2] E. Qazi, A. T. Wilson, C. McDougall, M. Boesen, F. S. Al-Ajlan, P. Pordeli, C. Batchelor, K. Khan, T. T. Sajobi, T.-Y. Lee, M. D. Hill, A. M. Demchuk, M. Goyal, C. D. d'Esterre, B. K. Menon, and N. D. Forkert, "Abstract tp45: One threshold does not fit all: Hounsfield unit thresholds to segment clot on ncct are patient specific," *Stroke*, vol. 47, no. Suppl 1, pp. ATP45–ATP45, 2016. [Online]. Available: http://stroke.ahajournals.org/content/47/Suppl_1/ATP45

[3] E. M. M. Santos, H. A. Marquering, O. A. Berkhemer, W. H. van Zwam, A. van der Lugt, C. B. Majoie, W. J. Niessen, and on behalf of the MR CLEAN investigators, "Development and validation of intracranial thrombus segmentation on ct angiography in patients with acute ischemic stroke," *PLOS ONE*, vol. 9, no. 7, pp. 1–8, 07 2014. [Online]. Available: https://doi.org/10.1371/journal.pone.0101985

[4] C. H. Riedel, U. Jensen, A. Rohr, M. Tietke, K. Alfke, S. Ulmer, and O. Jansen, "Assessment of thrombus in acute middle cerebral artery occlusion using thin-slice nonenhanced computed tomography reconstructions," *Stroke*, vol. 41, no. 8, pp. 1659–1664, 2010. [Online]. Available: http://stroke.ahajournals.org/content/41/8/1659

[5] Dice Lee R., "Measures of the amount of ecologic association between species," *Ecology*, vol. 26, no. 3, pp. 297–302, 1945. [Online]. Available: http://www.jstor.org/stable/1932409

[6] C. Pierpaoli, "Quantitative brain mri," *Topics in Magnetic Resonance Imaging*, vol. 21, no. 2, 2010. [Online]. Available: https://journals.lww.com/topicsinmri/Fulltext/2010/04000/Quantitative_Brain_MRI.1.aspx

[7] L. G. Nyul, J. K. Udupa, and X. Zhang, "New variants of a method of mri scale standardization," *IEEE Transactions on Medical Imaging*, vol. 19, no. 2, pp. 143–150, Feb 2000.

[8] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 6, pp. 84–90, May 2017. [Online]. Available: http://doi.acm.org/10.1145/3065386

[9] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," *CoRR*, vol. abs/1505.04597, 2015. [Online]. Available: http://arxiv.org/abs/1505.04597

[10] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. A. Riedmiller, "Striving for simplicity: The all convolutional net," *CoRR*, vol. abs/1412.6806, 2014. [Online]. Available: http://arxiv.org/abs/1412.6806

[11] J. Akeret, C. Chang, A. Lucchi, and A. Refregier, "Radio frequency interference mitigation using deep convolutional neural networks," *Astronomy and Computing*, vol. 18, pp. 35–39, 2017.

[12] O. Cicek, A. Abdulkadir, S. Lienkamp, T. Brox, and O. Ronneberger, "3d u-net: Learning dense volumetric segmentation from sparse annotation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, S. Ourselin, L. Joskowicz, M. R. Sabuncu, G. Unal, and W. Wells, Eds. Cham: Springer International Publishing, 2016, pp. 424–432.

[13] C. Zhang and Y. Ma, *Ensemble Machine Learning: Methods and Applications*. Springer Publishing Company, Incorporated, 2012.