

An anatomically-informed 3D CNN for brain aneurysm classification with weak labels*

Tommaso Di Noto¹[0000-0002-5161-055X], Guillaume Marie¹[0000-0002-2447-1056], Sébastien Tourbier¹[0000-0002-4441-899X], Yasser Alemán-Gómez^{1,2}[0000-0001-6067-8639], Guillaume Saliou¹[0000-0003-3832-7976], Meritxell Bach Cuadra^{1,3,4}[0000-0003-2730-4285], Patric Hagmann¹[0000-0002-2854-6561], and Jonas Richiardi¹[0000-0002-6975-5634]

¹ Department of Radiology, Lausanne University Hospital and University of Lausanne, Lausanne, Switzerland

² Center for Psychiatric Neuroscience, Department of Psychiatry, Lausanne University Hospital and University of Lausanne, Lausanne, Switzerland

³ Medical Image Analysis Laboratory, Center for Biomedical Imaging, Lausanne, Switzerland

⁴ Signal Processing Laboratory (LTS5), Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

Abstract. A commonly adopted approach to carry out detection tasks in medical imaging is to rely on an initial segmentation. However, this approach strongly depends on voxel-wise annotations which are repetitive and time-consuming to draw for medical experts. An interesting alternative to voxel-wise masks are so-called “weak” labels: these can either be coarse or oversized annotations that are less precise, but noticeably faster to create. In this work, we address the task of brain aneurysm detection as a patch-wise binary classification with weak labels, in contrast to related studies that rather use supervised segmentation methods and voxel-wise delineations. Our approach comes with the non-trivial challenge of the data set creation: as for most focal diseases, anomalous patches (with aneurysm) are outnumbered by those showing no anomaly, and the two classes usually have different spatial distributions. To tackle this frequent scenario of inherently imbalanced, spatially skewed data sets, we propose a novel, anatomically-driven approach by using a multi-scale and multi-input 3D Convolutional Neural Network (CNN). We apply our model to 214 subjects (83 patients, 131 controls) who underwent Time-Of-Flight Magnetic Resonance Angiography (TOF-MRA) and presented a total of 111 unruptured cerebral aneurysms. We compare two strategies for negative patch sampling that have an increasing level of difficulty for the network and we show how this choice can strongly affect the results. To assess whether the added spatial information helps improving performances, we compare our anatomically-informed CNN with a baseline, spatially-agnostic CNN. When considering the more realistic and challenging scenario including vessel-like negative patches, the former model

* This is the author version of the accepted paper at the 3rd Workshop on Machine Learning in Clinical Neuroimaging (MLCN), 2020. The published version can be found at https://doi.org/10.1007/978-3-030-66843-3_6.

attains the highest classification results (accuracy \simeq 95%, AUROC \simeq 0.95, AUPR \simeq 0.71), thus outperforming the baseline.

Keywords: 3D-CNN · Negative Sampling · Weak Labels · Magnetic Resonance Angiography · Aneurysm Detection.

1 Introduction

Cerebral aneurysms (CA) are abnormal focal dilatations in brain arteries caused by a weakness in the blood vessel wall. The overall population prevalence of CA ranges from 5% to 8% [1] and CA rupture is the predominant cause of nontraumatic subarachnoid hemorrhages (SAH) [2]. The mortality rate of aneurysmal SAH is around 40% and only half of post-SAH patients return to independent life [3,4]. Considering that the workload of radiologists is steadily increasing [5,6] and the detection of CAs is deemed a non-trivial task (especially for small aneurysms) [7], the development of an automatic tool able to detect aneurysms before they become symptomatic would be highly beneficial, both to reduce false negative cases, and to speed up the daily workflow in radiology departments.

Nowadays, non-enhanced Time-Of-Flight Magnetic Resonance Angiography (TOF-MRA) is routinely used for CA detection because of its high sensitivity (\simeq 95%) and pooled specificity of 89% [8]. Also, it has the advantage of being non-invasive and without radiation exposure, as opposed to Digital Subtraction Angiography (DSA) or Computed Tomography Angiography (CTA).

In the last few years, several medical imaging tasks such as classification, detection and segmentation have been profoundly revolutionized by the application of deep learning (DL) algorithms [9] which have shown a noteworthy potential. However, DL has to deal with the recurrent challenge of limited availability of (labelled) training examples, for building predictive algorithms that do not suffer from overfitting [10]. This is especially true in radiology where the voxel-wise manual annotation of medical images is commonly considered a tedious and time-consuming task [11] which often takes away precious time from experts.

The task of automated brain aneurysm detection with DL algorithms has already been addressed by several research groups. For instance, [12] used 2D patches and a ResNet-like model to detect aneurysms from TOF-MRA. Similarly, 2D Maximum Intensity Projection (MIP) patches with Convolutional Neural Network (CNN) have been proposed by [7,13]. In [14], 2D nearby projection (NP) images extracted from 3D CTA are fed as input to a Region-CNN (R-CNN) to detect aneurysms. Other works rather use 3D patches to perform aneurysm detection either in MRA or CTA imaging [15,16]. Though many of these works present encouraging results for the development of a Computer-Assisted Diagnosis (CAD) system for aneurysm detection, most of them [7,13,15,16] build their supervised models starting from voxel-wise manual annotations. From these annotations, they either carry out plain aneurysm segmentation [16], or they first perform a segmentation and then refine it with post-processing steps [7,13,15], in order to obtain detection bounding boxes.

Differently from previous approaches, our work investigates the task of brain aneurysm classification exploiting “weak” labels. In our application, these correspond to manual annotations which are not drawn with voxel-wise precision, but rather consist of spheres enclosing the aneurysms which are faster to create for the expert with respect to a slice-by-slice labelling. The concept of “lazy” or “weak” labels has already been used in previous works, in particular for segmentation [17,18], or cell type concentration prediction [19] where full labelling would be infeasible. However, while [17,19] exploited under-labelled data, we use over-labelled (more labelled voxels than actual true positives) data to perform aneurysm classification.

The goal of this study is three-fold: first, we assess the capability of a custom CNN to distinguish 3D TOF-MRA patches positive/negative for aneurysms, using weak labels. Second, we show the substantial impact that negative sampling can have on classification performances. Lastly, we propose an anatomically-driven solution to mitigate the problem of negative sampling for our dataset and for similar medical imaging tasks.

2 Materials and Methods

2.1 Data set

A retrospective cohort of 214 subjects who underwent clinically-indicated TOF-MRA between 2010 and 2012 was used. Out of these 214 subjects, 83 had one (or more) aneurysm(s), while 131 did not present any. For the former group, patients with one or more unruptured intracranial aneurysms were included, while patients with treated and ruptured aneurysms were excluded. Different aneurysms of the same patient were treated as independent, but most patients (81%) had only one aneurysm. Similarly, for patients with multiple sessions, we treated each session independently. The overall number of aneurysms included in the study is 111 and their anatomical location distribution is shown in Table 1. A 3D gradient recalled echo sequence with Partial Fourier technique was used for all subjects (see MR acquisition parameter details in Table 2). Aneurysms were annotated by one radiologist with 4 years of experience in neuroimaging. The Mango software was used to create the aforementioned weak labels which correspond to spheres that enclose the whole aneurysm, regardless of the shape (i.e. saccular, fusiform or multilocular). All TOF-MRA subjects included in the study were double checked by a senior neuroradiologist with over 14 years of experience, in order to exclude potential false positives or false negatives that might have been present in the original medical reports. The data set was organized according to the Brain Imaging Data Structure (BIDS) standard [20].

2.2 Image processing

Two preprocessing steps were carried out for each subject. First, we performed skull-stripping with the Brain Extraction Tool [21] to remove regions such as

Table 1. Spatial distribution of aneurysms. MCA = Middle Cerebral Artery, ACOM = Anterior Communicating Artery, PCOM = Posterior Communicating Artery.

| | Count | % |
|---------------------------------|-------|------|
| MCA | 22 | 19.8 |
| ACOM | 20 | 18.0 |
| Intradural carotid other | 13 | 11.7 |
| Carotid extra | 13 | 11.7 |
| MC other | 8 | 7.2 |
| Carotid tip | 8 | 7.2 |
| Pericallosal | 8 | 7.2 |
| PCOM | 7 | 6.3 |
| Basilar tip | 5 | 4.5 |
| Ophthalmic | 4 | 3.6 |
| Post other | 3 | 2.7 |

Table 2. MR acquisition parameters of TOF-MRA scans used for the study population.

| # scans | Vendor | Model | Field Strength [T] | TR [ms] | TE [ms] | Pixel spacing [mm ²] | Slice Thickness [mm] | Slice Gap [mm] |
|---------|----------------------|----------|--------------------|---------|---------|----------------------------------|----------------------|----------------|
| 81 | Philips | Intera | 3.0 | 18.3 | 3.40 | 0.41x0.41 | 1.1 | 0.55 |
| 10 | Siemens Healthineers | Aera | 1.5 | 24.0 | 7.0 | 0.35x0.35 | 0.5 | 0.09 |
| 21 | Siemens Healthineers | Skyra | 3.0 | 21.0 | 3.43 | 0.27x0.27 | 0.5 | 0.08 |
| 35 | Siemens Healthineers | Symphony | 1.5 | 39.0 | 5.02 | 0.39x0.39 | 1 | 0.25 |
| 28 | Siemens Healthineers | TrioTim | 3.0 | 23.0 | 4.18 | 0.46x0.46 | 0.69 | 0.14 |
| 61 | Siemens Healthineers | Verio | 3.0 | 22.0 | 3.95 | 0.46x0.46 | 0.70 | 0.13 |

the skull or the eyes. Second, a probabilistic vessel atlas built from multi-center MRA data sets [22] was co-registered to each patient’s TOF-MRA using the Advanced Neuroimaging Tools (ANTs) [23]. Specifically, we first registered the probabilistic atlas to the T1-weighted anatomical scan of each patient through a symmetric diffeomorphic registration. Second, we registered the obtained warped volume to the TOF subject space through an affine registration. The registered atlas was used only to provide prior information about vessel locations for the patch sampling strategy (see 2.4 below). As for most of the previously mentioned studies, we adopt a patch-based approach for the classification of aneurysms: we use 3D TOF-MRA patches as input samples to our network, rather than the entire volumes.

2.3 An anatomically-informed 3D-CNN

The task of aneurysm classification is extremely spatially constrained, since not only aneurysms solely occur in arteries, but they also occur in precise locations of the vasculature that have higher probability than others. Inspired by previous works in neuroimaging [24,25], we decided to include this strong anatomical prior

into our model. This was achieved in two ways: first, we designed a two-channel CNN which analyzes patches of the input volume at different spatial scales, in order to provide anatomical context on the vascular tree surrounding the patch of interest. Second, we computed for each input sample a numerical vector \mathbf{D} containing tailored spatial features which is integrated in the fully connected layers of the network; namely, the vector \mathbf{D} includes the $[x,y,z]$ coordinates of the center of the input patch in MNI space, and the Euclidean distances from this center to several coordinate and landmark points (also in MNI space). The coordinate points belong to a cubic ($6 \times 6 \times 6$), uniformly-sampled grid superimposed on the TOF-MRA volume. The landmark points correspond to 24 arterial locations where aneurysms are most recurrent; these were selected basing on the literature [26] and on our data. The final dimension of \mathbf{D} is 243. A 3D visual representation of these distances and of the creation of \mathbf{D} are provided in Figure 1.




Fig. 1. Computation of spatial features: for each input patch, we extract a feature vector \mathbf{D} which is composed of the $[x,y,z]$ coordinates of the patch center (in purple), the distances from this center to the points of a uniform grid (light green) and the distances to some landmarks (red dots) recurrent for aneurysms (light blue). The main brain arteries are segmented from the vessel atlas and are depicted in yellow.

Network architecture - We designed a custom CNN with building blocks inspired by the VGG-16 network [27]. Figure 2 illustrates in detail the structure of our CNN. As already proposed in [28,29], we performed late fusion of the features extracted from the input channels. Essentially, the same stack of convolutional layers is applied in parallel both to the small-scale and large-scale TOF-MRA patches. Then, feature vectors are merged, passed to a stack of fully connected layers, and further concatenated with vector \mathbf{D} . The rationale behind the multi-scale approach is that the convolutions over the small-scale patches produce aneurysm-specific features, whereas the large-scale patches provide con-




Fig. 2. CNN architecture: features are extracted in parallel from small-scale and large-scale TOF-MRA patches through a stack of convolutional layers. Then, they are merged into a single fully connected layer. Later, the spatial information vector \mathbf{D} is concatenated to a fully-connected layer.

text/spatial descriptors. Since most of the aneurysms in our dataset (92%) had an equivalent diameter smaller than 30 voxels, we decided to fix the side of the small-scale input patches to 30. Instead, a side of 80 was set for the large-scale patches in order to include even the largest aneurysm (equivalent diameter = 58 voxels) and some context around it. All patches were standardized to have mean 0 and variance 1 before being fed to the CNN as suggested by [30]. The standardization was also performed to mitigate intensity differences which are inherently present across different patients and scanners [31]. A kernel size of $3 \times 3 \times 3$ was used in all convolutional layers, with padding and a stride=1 in all directions. We applied the Rectified Linear Unit (ReLU) activation function for all layers, except for the last fully connected layer which is followed by a sigmoid function. To fit the model, the Adam optimization algorithm [32] was applied with variable learning rate, together with the binary cross-entropy loss. Moreover, we used the Xavier initialization [33] for all the layers of the CNN. Biases were initialized to 0 and a batch size of 4 was chosen. The final output of the CNN is simply the class probability p of the input sample: positive (patch with aneurysm) if $p > 0.5$ or negative (without aneurysm) if $p < 0.5$.

To elucidate whether the injection of anatomical and spatial information into the model can improve classification results, we compare two distinct architectures: the **anatomically-informed 3D-CNN** illustrated in Figure 2 and a **baseline CNN** which is identical to the previous one, but has one input channel (i.e. only the small-scale TOF-MRA) and no spatial features (i.e. no vector \mathbf{D}). The former has 884,529 trainable parameters, while the latter has 448,417.

2.4 Patch sampling strategy

In addition to the comparison between models (anatomically-informed and baseline), we also investigated the influence that negative sampling can have on classification performances. Indeed, while for the minority class (i.e. patches with aneurysm) the sampling is restricted by the availability of positive cases, extraction of negative samples (majority-class sampling) necessarily entails the choice of one (or more) extraction criteria. Therefore, we chose two different sampling criteria for extracting negative patches:

1) Random sampling: negative samples were extracted randomly within the skull-stripped brain, without overlapping with the positive patches.

2) Intensity-matched sampling: we imposed intensity constraints for the extraction basing on the co-registered vessel atlas. More specifically, with an iterative search, we only extract the negative sample when the corresponding (i.e. same center coordinates) vessel patch has both a local (patch-wise) and global (volume-wise) brightness that are higher than some specific thresholds. These thresholds, in turn, were chosen empirically according to the local and global brightness of all positive patches in the vessel atlas. First, this sampling strategy avoids extracting patches that are too dark with respect to positive ones. Second, it allows us to extract patches which always include part of the vasculature. Needless to say, this sampling creates both more realistic and more difficult negative samples for the CNN.

Instead, positive patches were extracted around the aneurysms in a non-centered fashion, but ensuring that the aneurysm mask was always completely included in the small-scale patch. As last step, regardless of the sampling strategy, we combined the samples (negative and positive) of all subjects into a unique dataset that was fed as input to the CNN. We decided to extract 8 negative samples per subject. This led to a final dataset composed of 1808 negative and 111 positive samples (ratio $\simeq 1 : 16$). During training, a series of data augmentation techniques were applied on positive patches: namely, rotations ($90^\circ, 180^\circ, 270^\circ$), horizontal flip, elastic deformation and contrast adjustment. Training and evaluation were performed with Tensorflow 2.0 and a GeForce RTX 2080TI GPU.

2.5 Evaluation approach

We evaluated all different scenarios through a nested stratified Cross Validation (CV), with 5 external folds and 3 internal folds. This ensured that the patches in every test fold were always unseen samples with respect to the training set and to the chosen hyperparameters. The only hyperparameter that was tuned in the internal CV is the learning rate: 0.00001, 0.0001 and 0.001 were tested. All other hyperparameters were fixed, so as not to excessively lengthen training time. We set a dropout rate of 0.2, a sample weight factor of 3 to give more importance to the minority class and we trained the CNN for 50 epochs with an early stopping condition on the validation/test set. To statistically compare classification results, Wilcoxon signed-rank tests were performed [35]. For simplicity, the tests only accounted for the area under the PR (AUPR) curve of the

classifiers, since this metric is particularly suited when working with imbalanced data sets [34]. A significance threshold level $\alpha = 0.05$ was set for comparing P values. First, we saved the best hyperparameters for each experiment. Then, we re-ran the training/test of the CNN 10 times. For each of the 10 realizations we always changed the patient order and ensured that the two models (baseline and anatomically-informed) were evaluated against the exact same samples.

3 Results

Overall, **4 experiments** were carried out: the two networks (anatomically-informed and baseline) were evaluated first against the dataset with random negative samples and then against the dataset with intensity-matched negative samples. Classification results of the four experiments are reported in Table 3. Training the model took about 2 hours for the baseline model and 3 hours for the anatomically-informed one. The most frequent learning rate across the external test folds was 0.0001. The Wilcoxon tests performed on the AUPRs distributions highlighted two main findings: first, both for the baseline and the anatomically-informed model, AUPRs were statistically higher when random negative sampling was used with respect to the intensity-matched sampling ($P = 0.01$). This proves how one task is evidently easier than the other. Second, when comparing the baseline model and the anatomically-informed model against the intensity-matched dataset (difficult scenario), AUPR distributions were again significantly different ($P = 0.01$), suggesting that the proposed anatomically-informed CNN indeed outperforms the baseline.

Table 3. Classification results of baseline and anatomically-informed models both with random negative patches and with intensity-matched (IM) ones. Acc=accuracy, Sens=sensitivity, Spec=specificity, PPV=positive predictive value, NPV=negative predictive value, AUC=Area Under ROC Curve, AUPR=Area Under PR curve.

| Network | Negative Samples | Acc (%) | Sens (%) | Spec (%) | PPV (%) | NPV (%) | AUC | AUPR |
|-----------------------|------------------|---------|----------|----------|---------|---------|------|------|
| Baseline | Random | 95.4 | 84.7 | 96.2 | 59.7 | 99.0 | .961 | .779 |
| | IM | 93.2 | 77.5 | 94.1 | 45.8 | 98.6 | .949 | .608 |
| Anatomically Informed | Random | 97.3 | 91.9 | 97.7 | 72.6 | 99.5 | .979 | .875 |
| | IM | 94.7 | 77.5 | 95.7 | 53.6 | 98.6 | .946 | .714 |

4 Discussion

This work presented an alternative approach for performing cerebral aneurysm detection when voxel-wise annotations are not available. To this end, we proposed a binary classification method, making use of weak labels enclosing the aneurysms of interest. In addition, we shed light over the recurrent problem of negative sampling in imbalanced and spatially-skewed data sets, showing how

this step can dramatically alter final results. Lastly, we devised a tailored CNN able to mitigate the negative sampling problem, by incorporating spatial anatomical information. This CNN was able to outperform its baseline counterpart despite the small sample size and having about twice the number of parameters. We believe this general principle is applicable to several other brain diseases with sparse spatial extent.

Our work is limited by the relatively high number of false positive cases even for the anatomically-informed CNN (see low PPV in Table 3). In addition, a separate analysis should be performed to understand whether the added distances of vector \mathbf{D} are indeed helpful: these might be redundant with respect to the $[x,y,z]$ center coordinates of the patches, which could already be informative enough. Though the presented patch-wise analysis is useful to gain insights on the network performances, it cannot be easily exploited in a clinical scenario. Thus, future work will aim at shifting towards a patient-wise analysis. Lastly, we acknowledge that the dataset size is still limited and it should be increased.

References

1. Rinkel, Gabriel JE, et al. "Prevalence and risk of rupture of intracranial aneurysms: a systematic review." *Stroke* 29.1 (1998): 251-256.
2. Jaja, Blessing NR, et al. "Clinical prediction models for aneurysmal subarachnoid hemorrhage: a systematic review." *Neurocritical care* 18.1 (2013): 143-153.
3. Frösen, Juhana, et al. "Saccular intracranial aneurysm: pathology and mechanisms." *Acta neuropathologica* 123.6 (2012): 773-786.
4. Xu, Zhen, et al. "Intracranial Aneurysms: Pathology, Genetics, and Molecular Mechanisms." *Neuromolecular medicine* (2019): 1-19.
5. Rao, Balaji, et al. "Utility of Artificial Intelligence Tool as a Prospective Radiology Peer Reviewer—Detection of Unreported Intracranial Hemorrhage." *Academic Radiology* (2020).
6. McDonald, Robert J., et al. "The effects of changes in utilization and technological advancements of cross-sectional imaging on radiologist workload." *Academic radiology* 22.9 (2015): 1191-1198.
7. Nakao, Takahiro, et al. "Deep neural network-based computer-assisted detection of cerebral aneurysms in MR angiography." *Journal of Magnetic Resonance Imaging* 47.4 (2018): 948-953.
8. Chen, Xiaodan, et al. "Meta-analysis of computed tomography angiography versus magnetic resonance angiography for intracranial aneurysm." *Medicine* 97.20 (2018).
9. Syeda-Mahmood, Tanveer. "Role of big data and machine learning in diagnostic decision support in radiology." *Journal of the American College of Radiology* 15.3 (2018): 569-576.
10. Shen, Dinggang, Guorong Wu, and Heung-Il Suk. "Deep learning in medical image analysis." *Annual review of biomedical engineering* 19 (2017): 221-248.
11. Razzak, Muhammad Imran, Saeeda Naz, and Ahmad Zaib. "Deep learning for medical image processing: Overview, challenges and the future." *Classification in BioApps*. Springer, Cham, 2018. 323-350.
12. Ueda, Daiju, et al. "Deep learning for MR angiography: automated detection of cerebral aneurysms." *Radiology* 290.1 (2019): 187-194.

13. Stemmer, Joseph N., et al. "Convolutional neural networks for the detection and measurement of cerebral aneurysms on magnetic resonance angiography." *Journal of digital imaging* 32.5 (2019): 808-815.
14. Dai, Xilei, et al. "Deep learning for automated cerebral aneurysm detection on computed tomography images." *International Journal of Computer Assisted Radiology and Surgery* (2020): 1-9.
15. Sichtermann, T., et al. "Deep Learning-Based Detection of Intracranial Aneurysms in 3D TOF-MRA." *American Journal of Neuroradiology* 40.1 (2019): 25-32.
16. Park, Allison, et al. "Deep Learning-Assisted Diagnosis of Cerebral Aneurysms Using the HeadXNet Model." *JAMA network open* 2.6 (2019): e195600-e195600.
17. Ke, Rihuan, et al. "A multi-task U-net for segmentation with lazy labels." *arXiv preprint arXiv:1906.12177* (2019).
18. Ezhov, Matvey, Adel Zakirov, and Maxim Gusarev. "Coarse-to-fine volumetric segmentation of teeth in cone-beam ct." *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*. IEEE, 2019.
19. Abousamra, Shahira, et al. "Weakly-Supervised Deep Stain Decomposition for Multiplex IHC Images." *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2020.
20. Gorgolewski, Krzysztof J., et al. "The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments." *Scientific data* 3.1 (2016): 1-9.
21. Smith, Stephen M. "Fast robust automated brain extraction." *Human brain mapping* 17.3 (2002): 143-155.
22. Mouches, Pauline, and Nils D. Forkert. "A statistical atlas of cerebral arteries generated using multi-center MRA datasets from healthy subjects." *Scientific data* 6.1 (2019): 1-8.
23. Avants, Brian B., et al. "A reproducible evaluation of ANTs similarity metric performance in brain image registration." *Neuroimage* 54.3 (2011): 2033-2044.
24. Ghafoorian, Mohsen, et al. "Location sensitive deep convolutional neural networks for segmentation of white matter hyperintensities." *Scientific Reports* 7.1 (2017): 1-12.
25. Ganaye, Pierre-Antoine, Michaël Sdika, and Hugues Benoit-Cattin. "Towards integrating spatial localization in convolutional neural networks for brain image segmentation." *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. IEEE, 2018.
26. Brown Jr, Robert D., and Joseph P. Broderick. "Unruptured intracranial aneurysms: epidemiology, natural history, management options, and familial screening." *The Lancet Neurology* 13.4 (2014): 393-404.
27. Simonyan, Karen, and Andrew Zisserman. "Very deep convolutional networks for large-scale image recognition." *arXiv preprint arXiv:1409.1556* (2014).
28. Huang, Yechong, et al. "Diagnosis of Alzheimer's disease via multi-modality 3D convolutional neural network." *Frontiers in Neuroscience* 13 (2019): 509.
29. Nie, Dong, et al. "Multi-channel 3D deep feature learning for survival time prediction of brain tumor patients using multi-modal neuroimages." *Scientific reports* 9.1 (2019): 1-14.
30. Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
31. Zhuge, Ying, and Jayaram K. Udupa. "Intensity standardization simplifies brain MR image segmentation." *Computer vision and image understanding* 113.10 (2009): 1095-1103.

32. Kingma, Diederik P., and Jimmy Ba. "Adam: A method for stochastic optimization." arXiv preprint arXiv:1412.6980 (2014).
33. Glorot, Xavier, and Yoshua Bengio. "Understanding the difficulty of training deep feedforward neural networks." Proceedings of the thirteenth international conference on artificial intelligence and statistics. 2010.
34. Davis, Jesse, and Mark Goadrich. "The relationship between Precision-Recall and ROC curves." Proceedings of the 23rd international conference on Machine learning. 2006.
35. Wilcoxon, Frank. "Individual comparisons by ranking methods." Breakthroughs in statistics. Springer, New York, NY, 1992. 196-202.