Using a Siamese Network to Accurately Detect Ischemic Stroke in Computed Tomography Scans

Ana Beatriz Vieira[1](✉), Ana Catarina Fonseca[2], José Ferro[2], and Arlindo L. Oliveira[1]

[1] INESC-ID/Instituto Superior Técnico, University of Lisbon, Lisbon, Portugal
{anabeatrizvieira,arlindo.oliveira}@tecnico.ulisboa.pt

[2] Department of Neurosciences and Mental Health, Neurology, Hospital de Santa Maria, CHULN, Lisbon, Portugal
{acfonseca,jmferro}@medicina.ulisboa.pt

Abstract. The diagnosis of stroke, a leading cause of death in the world, using computed tomography (CT) scans, makes it possible to assess the severity of the incident and to determine the type and location of the lesion. The fact that the brain has two hemispheres with a high level of anatomical similarity, exhibiting significant symmetry, has led to extensive research based on the assumption that a decrease in symmetry is directly related to the presence of pathologies. This work is focused on the analysis of the symmetry (or lack of it) of the two brain hemispheres, and on the use of this information for the classification of computed tomography brain scans of stroke patients. The objective is to contribute to a more precise diagnosis of brain lesions caused by ischemic stroke events. To perform this task, we used a siamese network architecture that receives a double two-dimensional image of a CT slice (the original and a mirrored version) and a label that reflects the existence or not of a visible stroke event. The network then extracts the relevant features and can be used to classify brain-CT slices taking into account their perceived symmetry. The best performing network exhibits an average accuracy and F1-score of 72%, when applied to CT slices of previously unseen patients, significantly outperforming two state-of-the-art convolutional network architectures, which were used as baselines. When applied to slices chosen randomly, that may or may not be from the same patient, the network exhibits an accuracy of 97%, but this performance is due in part to overfitting, as the system is able to learn specific features of each patient brain.

Keywords: Ischemic stroke · Computed tomography · Symmetry detection · Image classification · Siamese network

1 Introduction

According to the *World Stroke Organization*, stroke is the second leading cause of death and disability in the world, with about 13 million cases annually [9].

G. Marreiros et al. (Eds.): EPIA 2022, LNAI 13566, pp. 159–170, 2022.
https://doi.org/10.1007/978-3-031-16474-3_14

In this work only ischemic strokes were considered, since they represent a large fraction of all stroke events. Ischemic stroke occurs when a vessel supplying blood to the brain is obstructed, leading to damage or death of brain cells. It is usually caused by blood clots in a brain vessel or by narrowing of the blood vessels that irrigate the brain due to a process of atherosclerosis [15].

The diagnosis is made based on the evaluation of the symptoms and images resulting from Magnetic Resonance Imaging (MRI) or Computed Tomography (CT - used in this paper) exams. CT is most often chosen due to its widespread availability and short imaging time, but the resulting images are harder to interpret by automated means. The exams should be performed immediately upon admission, in order to identify and evaluate as rapidly as possible the existing lesions. The time elapsed since the onset of the stroke is crucial for the treatment and recovery of the patient. Recovery depends on the severity of the stroke, and the faster the diagnosis and treatment, the greater the chance of recovering the penumbra, which is the area around the ischemic region that can be recovered [15]. CT scans, obtained shortly after the incident, can detect the presence of ischemic or hemorrhagic lesions as well as the location and extent of the lesion, and exclude situations that can be confused with strokes [16].

There are several types of CT scans that can detect different changes in brain structures. The first exam most commonly performed is the Non-Contrast Computed Tomography (NCCT), since it enables the quick identification of the stroke type (ischemic or hemorrhagic) [14]. Depending on the type, other exams may follow. In our case, since only patients with ischemic stroke are considered, the exams that are usually performed later include Computed Tomography Perfusion (CTP) and Computed Tomography Angiography (CTA).

A human brain has two roughly symmetrical hemispheres, separated by the longitudinal fissure, which is a membrane filled with Cerebro-Spinal Fluid (CSF) [19]. This is usually known as the Mid-Sagittal Plane (MSP), because it is a virtual plane perpendicular to the brain, which divides it into left and right halves [19,21]. It is well known that the two hemispheres display both anatomical and functional asymmetries and this has been a significant topic of research [13]. Several results have shown that the asymmetry of the hemispheres is correlated with the presence of brain injuries [10,20], tumors [2,23], or mental illness [18].

The results presented in this paper are based on an approach that uses the level of symmetry between the two brain hemispheres to detect the presence or absence of stroke. The approach is based on the use of a siamese network, a particular neural network architecture developed to identify similarities between images, in this case symmetric ones[1].

After introducing the problem, we present some related work and describe the methodology used in the construction of the dataset, including image preprocessing and data augmentation. We then present and discuss the results obtained with the different architectures and the conclusions obtained from this work.

[1] Code available at: https://github.com/anagilvieira/siamese_network.git.

2 Related Work

One of the most effective methods for stroke segmentation was developed by Chen *et al.* [3], who proposed a framework consisting of two convolutional neural networks (CNNs) to segment stroke lesions using diffusion-weighted imaging (DWI), a specific MRI technique. One of the networks was a combination of two DeconvNets and the other one was a multiscale convolutional label evaluation network whose purpose is to evaluate the lesions detected by the first network, removing potential false positives.

Chin *et al.* [4] proposed a method that was used to assist neurologists in the diagnosis to improve the chances of recovery. This work developed a CNN for the early detection of ischemic stroke, which accepted as input CT images of the brain enriched with MRI data. Skull bone and other structures that might mislead the model were removed, since these structures exhibited similar pixel intensities as the ischemic area. The authors reported an accuracy of 93% in the classification task.

Recently, Herzog and Magoulas [8] proposed a method based on brain asymmetry to identify early dementia and its diverse stages, such as amnestic early mild cognitive impairment and Alzheimer's disease. They analyzed the structural and functional cerebral changes in both hemispheres using supervised machine learning algorithms and convolutional neural networks. The dataset was composed of brain asymmetries features, extracted from MRI scans from the Alzheimer's disease neuroimaging initiative database. The proposed pipeline achieved an accuracy that ranged from 75% to 93% for the mentioned diseases. Furthermore, this method offers a promising low-cost alternative for the classification of dementia and could potentially be useful in other brain degenerative disorders that are accompanied by visible changes in brain symmetry.

Inspired by the characteristics of the siamese network, Barman *et al.* [1] proposed a siamese neural network (DeepSymNet) for the detection of ischemic stroke from CTA images. This method enabled them to detect the changes in symmetry of vascular and brain tissue texture of the two brain hemispheres in parallel. The model was tested on a clinical dataset of 217 patients, and an AUC (Area Under the Curve) greater than 0.89 was obtained.

To perform brain symmetry studies on large neuroimaging archives, reliable and automatic detection of the mid-sagittal plane (MSP) is required to extract the brain hemispheres. However, traditional planar estimation techniques fail when the MSP presents a curvature caused by existing pathology or a natural phenomenon known as brain torque. As a result, midline estimates can be inaccurate. To address this matter, Gibicar *et al.* [6] suggested an unsupervised midline estimation method that consisted of three main stages: head angle correction, control point estimation, and midline generation. The technique was applied on a slice-by-slice basis for more accurate results and is able to provide accurate delineation of the midline even in the septum pellucidum (exactly in the middle of the brain), which is a source of failure for traditional approaches.

Our approach differs from these and other related works in that we use only CT data obtained at the time of hospital admission. This makes the problem harder, since the changes in the brain images are more subtle, but also potentially more relevant in a clinical setting.

3 Dataset and Methods

The data used in this work was collected from ischemic stroke patients at *Hospital de Santa Maria*, in Lisbon. The dataset is composed of non-contrast computed tomography images, in DICOM format. The data was properly anonymized to preserve participant privacy and to meet all ethical requirements. In order to use only the most relevant slices of each NCCT, those in which the ischemic stroke can be detected, the 3D data was converted to 2D slices, each normalized to a size of 512×512 pixels. The final dataset used in the tests consists of slices chosen from different brain regions of different patients. The dataset is balanced and includes 340 images of each class, where one class corresponds to the absence of visible effects of stroke (and therefore expected symmetry of the hemispheres) and the other class corresponds to visible stroke effects (leading to a lack of symmetry). The images were annotated by one of the authors, Ana Catarina Fonseca, an expert neurologist.

Given the limited size of the dataset, we applied data augmentation techniques, namely Gaussian blur, to the data. In the experimental procedure, we used stratified 5-fold cross-validation where, besides making sure that all the folds contained approximately the same number of images for each class, we also ensured that all the slices corresponding to one patient were included in a single fold, in order to avoid overfitting and label information leakage.

3.1 Image Preprocessing

We used image preprocessing techniques to transform raw image data into a clean image data. This process was one of the most time-consuming since CT scan images are very complex and present several undesired characteristics.

Windowing. A CT scan consists of a set of points called voxels. The voxels represent elementary three-dimensional tissue volumes. The 'color' of each voxel is conventionally displayed as a shade of grey, corresponding to the density of brain tissue at that location and is called the attenuation value, which is given in Hounsfield Units (HU) [7]. It is a relative quantitative measurement of radio density used by radiologists in the interpretation of CT images, and is calculated as

$$HU = \frac{(\mu_v - \mu_w)}{\mu_w} \times K, \tag{1}$$

where μ_v is the calculated voxel attenuation coefficient, μ_w is the water attenuation coefficient, and K is an integer constant, standardized as taking the value 1000 [7]. To set the appropriate window for the color range, we define two parameters: windows level (WL), also known as windows center, and windows width

(WW). The windows level is the midpoint of the range of the CT numbers and is responsible for the image brightness, while the windows width is the range of the CT numbers that an image contains and is responsible for the image contrast.

All values above the upper level appear as white and all values below the lower level appear as black [22]. In our case we want to see the details of the brain, so we need to choose a window that maximizes contrast while not losing information. If the range is chosen poorly some gradations may poorly visible [22]. We found that picking a range between 0 and 80 for CTs, i.e., WL = 40 and WW = 80 provided best results, which coincides with the window most commonly selected for stroke analysis.

Figure 1 shows two sample images from an NCCT (WL:40, WW:80) and CTA (WL:60, WW:360), respectively, where a lesion caused by ischemic stroke can be seen near the left ventricle. The lesion is marked by a red circle. As is clear from the figure, the choice of the correct window enables us to see more details and, consequently, makes it easier to detect the asymmetry between the two hemispheres.

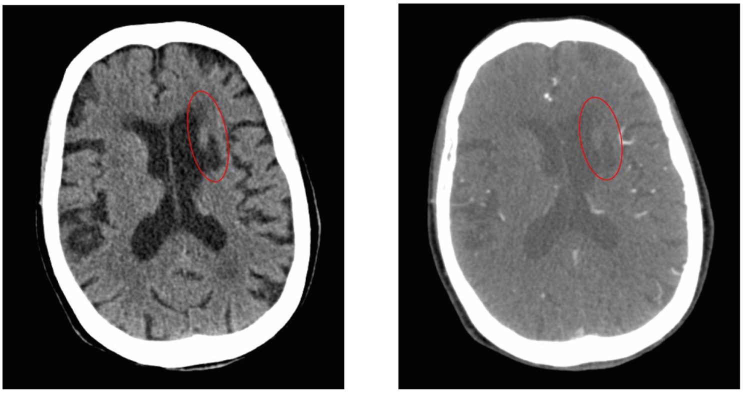

(a) Noncontrast Computed Tomography

(b) Computed Tomography Angiography

Fig. 1. Ischemic stroke lesions after two different windowing techniques were applied.

Head Tilt Correction. A common phenomenon in medical imaging devices is that they produce distorted brain images under certain circumstances, which can mislead visual inspection and lead to false clinical interpretation. The main reasons for this include the lack of a correct immobilization of the patients, inexperience of the technicians, and deficient calibration [10]. In order to correct the head tilt it is necessary to find the correct rotation angle, in order to align the MSP with the y-axis [8,19]. For that, the external contours of the brain were determined and an ellipse that best matches these contours was computed. Finally, the angle between the ellipse and the vertical axis was used to align the image. Figure 2 illustrates the same brain slice with and without head tilt correction.

Since the design of the external contours does not always fit the head optimally, there are images where the orientation of the head is not correctly aligned with the y-axis. This happens because of the slight (white) lines around the brain that causes interference in determining external contours.

Skull Stripping. Skull stripping is a fairly common process used with CT imaging because it is possible to focus only on the brain tissue, which is where the lesions are, and obtain a better segmentation of the different brain areas [4]. In fact, without the skull stripping step, the network attention was commonly focused on the eye region or on the bones around the brain.

Skull stripping was performed by subtracting from the original image an image where only the bone was visible, creating a resulting image with only the brain tissue. Figure 2 illustrates a slice before and after skull stripping.

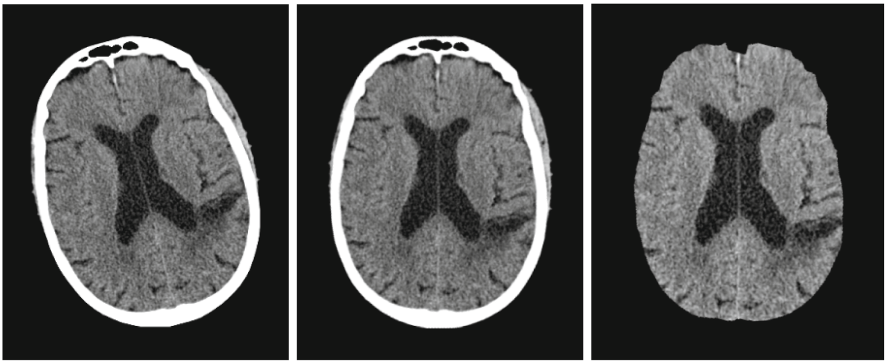

(a) Original slice **(b)** Head tilt correction **(c)** Skull stripping

Fig. 2. Main steps in CT images processing.

3.2 Symmetry Detection

Originally proposed for signature verification, similarity detection has been used extensively in various computer vision applications [1,5,11,12,17]. Similarity detection can be used to perform symmetry detection by using as inputs the original image and a mirrored image. We decided to use the siamese network (SiameseNet) to process two images, the original and a mirrored one. Unlike many other CNNs used in image recognition, this network does not assign a class to images, it only assesses the similarity between them.

The name is due to the fact that it is composed of two exactly equal convolutional neural networks that share the same parameters and weights. This network takes two images as input and calculates the similarity between them. Figure 3 represents the main stages of our siamese network. The first step is

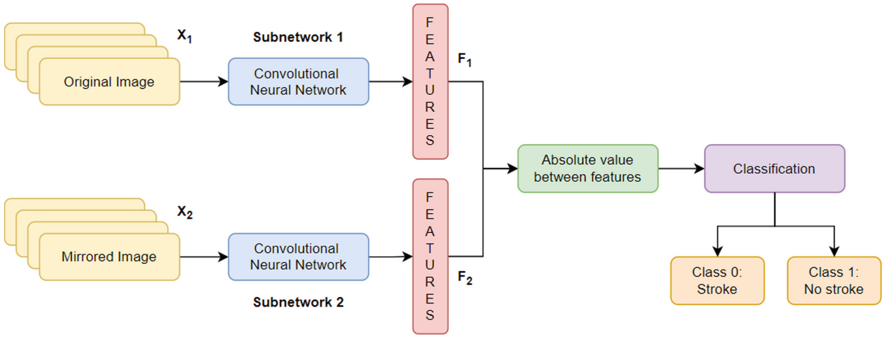

Fig. 3. Important stages of siamese network.

the feature extraction, where we used two CNNs (subnetwork 1 and subnetwork 2) to extract the features of each image. The network is trained with a label that reflects whether the two input images correspond to a visible stroke event (and are therefore less symmetric) or not. The output of each subnetwork was subsequently flattened (F_1 and F_2). The second step is the comparison between the two feature vectors. The difference of each subnetwork output is computed and the absolute value is used to generate a new feature vector D. Finally, the network output is the probability of the two input images being symmetric or not, reflecting the probability of stroke. The smaller the value of D, the more likely the images are to be symmetric and, therefore, to not exhibit evidence of stroke.

3.3 Proposed Architecture

Since the siamese network is made up of two equal networks, the network is trained, in each iteration, with a set of double images as input, the set size given by the batch size. One of the networks, subnetwork 1, takes as input the 'original' slices, with a resolution of 512×512, while subnetwork 2 takes as input the 'mirrored' slices, as shown in Fig. 3.

Table 1 shows in detail the characteristics of the different layers in our architecture. The feature extraction process is composed by 2D convolutional layers, 2D batch normalization layers, max pooling layers and a flatten layer, generating a feature vector with size 8192. The entire network uses ReLU activations, except for the prediction layer that uses softmax activation. After feature extraction and subtraction, two fully connected layers were used, the first generating a vector of size 1024 and the second generating a vector of size 2, one output unit for each class.

The model was trained and evaluated using stratified 5-fold cross-validation, for 200 epochs with a batch size of 32, using the ADAM optimizer with an initial learning rate of $\eta = 0.00001$. Since high weight values increase the chances of overfitting, L2 regularization with a weight factor $\lambda = 0.0005$ was used, applying a penalty for higher weight values. The model used 11M trainable parameters.

Table 1. Siamese network architecture.

Operation layer	Number of filters	Filter size	Stride value	Padding value	Output
Input layer					$1 \times 512 \times 512$
2D convolution	64	5×5	2	2×2	$64 \times 256 \times 256$
ReLU					$64 \times 256 \times 256$
Batch normalization 2D					$64 \times 256 \times 256$
Max pooling	64	2×2	2	0	$64 \times 128 \times 128$
2D convolution	128	5×5	2	2×2	$128 \times 64 \times 64$
ReLU					$128 \times 64 \times 64$
Batch normalization 2D					$128 \times 64 \times 64$
Max pooling	128	2×2	2	0	$128 \times 32 \times 32$
2D convolution	256	3×3	2	1×1	$256 \times 16 \times 16$
ReLU					$256 \times 16 \times 16$
Batch normalization 2D					$256 \times 16 \times 16$
Max pooling	256	2×2	2	0	$256 \times 8 \times 8$
2D convolution	512	3×3	2	1×1	$512 \times 4 \times 4$
ReLU					$512 \times 4 \times 4$
Batch normalization 2D					$512 \times 4 \times 4$
Flatten					8192
Fully connected					1024
Fully connected					2

3.4 Loss Function

When working with siamese networks, the most commonly used loss function is the contrastive loss, which is commonly used to compute a similarity score by calculating the Euclidean distance between the feature vectors [5]. However, since this is a classification task, we choose to use the cross entropy loss, given by $L(y_i, \hat{y_i}) = -(y_i \log \hat{y_i} + (1 - y_i) \log(1 - \hat{y_i}))$ where y_i and $\hat{y_i}$ are the actual class and the predicted class, respectively.

4 Results

To assess the performance of the siamese network architecture on this dataset, we performed an empirical comparison with two alternative baseline architectures, the SimResNet-18 and the ResNet-50. All models were trained using a PowerEdge C41402 server with an NVIDIA 32GB Tesla V100S.

ResNet-50. ResNet is a commonly used deep convolutional neural network architecture that normally exhibits good performance in image classification tasks. Other types of CNNs have difficulty optimizing the parameters during training because of vanishing gradients. ResNet minimizes these problem by using residual connections in parallel with the convolutional layers. We tested ResNets with 18 layers, 34 layers and 50 layers (three common variations of

ResNet) and found that the best performance was obtained using ResNet-50. With this model, we achieved 62% accuracy, using 23M parameters and 174 min of training time.

SimResNet-18. Initially motivated by the idea of combining ResNet and the siamese network architecture [17], we also tested an alternative approach where the ResNet is used only for the feature extraction step. The siamese network consists of two ResNets for feature extraction, where one of them accepts the original CT slice and the other accepts the mirrored slice. Since the ResNet uses a Global Average Pooling (GAP) layer after the last convolution layer, the resulting output is a feature vector with a size of 512, which is used to calculate the absolute value of the difference. The remaining components of the architecture are similar to those described for the SiameseNet. The feature vectors are subtracted, the absolute value of the difference is computed, and the images are classified into one of the two existing classes using a fully connected layer.

In this architecture, we tested different ResNet models, and after analyzing the number of parameters, the training time, and the results, opted for the 18-layer model. The hyperparameters used were the same as for the siamese network and the architecture used 11M parameters. This approach reached 64% accuracy and required 154 min to train.

Experimental Results. Table 2 summarizes the classification performance of the proposed architecture, SiameseNet, compared with the SimResNet-18 and ResNet-50 baselines. The results presented in the table show that the siamese network clearly outperforms the other two models, with an accuracy and F1-values of 72%. The runtime for training the SiameseNet in the proposed dataset was 144 min. Although the overall accuracy was not very high in absolute terms, it is the best result known to us for this specific problem, ischemic stroke classification from CT scans.

Table 2. Model classification performance.

Model	Accuracy	Precision	Recall	F1-score
ResNet-50	0.62 ± 0.09	0.71 ± 0.14	0.47 ± 0.27	0.57 ± 0.17
SimResNet-18	0.64 ± 0.04	0.73 ± 0.12	0.48 ± 0.12	0.58 ± 0.07
SiameseNet	0.72 ± 0.13	0.75 ± 0.17	0.69 ± 0.25	0.72 ± 0.14

The low accuracy results obtained with ResNet-50 are most likely due to the fact that CT images have several features that make model learning difficult. When ResNet models are added to the Siamese network (SimResNet-18), we see an improvement in performance, although not very significant. Finally, the siamese network architecture exhibits a considerable improvement in performance over the two baselines. However, the overall accuracy is still relatively low, mainly due to the fact that the dataset is small and the task very challenging.

The relatively high variance of the performance measures is essentially due to two factors. First, since the dataset is small and it cannot be guaranteed that all patients have the same number of samples, there is relevant variation on the fold size. In fact there are patients with 45 labeled slices and others with only 3 labeled slices. Imposing the restriction that all slices from the same patient were in only one of the folds is probably a significant source of performance variation between different folds.

In other experiments, when stratification was not used to guarantee that all the slices of one patient were in a single fold, we observed significant information leakage between the training and test sets, leading to the model not learning the features relevant to the problem but focusing on features that relate different slices of the same patient. In this case, when the dataset is split randomly, we obtained the results shown in Table 3. The results correspond to a very high accuracy for the SiameseNet (97%) but also for the two baseline networks. However, these results are not reliable nor reproducible in a clinical environment, since the network is using features that are not relevant to the problem, such as the dimension of the slices, or some other features that are only correlated with specific characteristics of the patient and not of the stroke effects.

Table 3. Model classification performance with random dataset splitting.

Model	Accuracy	Precision	Recall	F1-score
ResNet-50	0.89 ± 0.04	0.87 ± 0.08	0.92 ± 0.07	0.89 ± 0.05
SimResNet-18	0.91 ± 0.03	0.89 ± 0.04	0.93 ± 0.040	0.91 ± 0.03
SiameseNet	0.97 ± 0.02	0.96 ± 0.02	0.97 ± 0.03	0.97 ± 0.02

5 Conclusions and Future Work

We proposed to use a siamese network architecture for the detection of ischemic stroke, based on the fact that detected asymmetry of brain hemispheres is strongly correlated with evidence of stroke. The architecture obtained 72% accuracy and F1-score, on average, on an independent test set. Although the absolute value of the accuracy remains low and insufficient for clinical use, this architecture may become the starting point for other studies that use the lack of symmetry between the brain hemispheres to diagnose pathologies from CT scans.

During the development of the work, we took special care to avoid information leakage between the training and test sets, a phenomenon that leads to overfitting and limits reproducibility in a clinical setting. In some experiments, we achieved unexpectedly high performance values, until we realized that the model was learning from features uncorrelated with the target task.

There are some aspects of the presented work that can be improved. One of these aspects is image preprocessing. Although the skull stripping process is mostly successful there are slices where small portions of non-brain tissue are mistakenly removed. This happens because all the white pixels in the image are removed when only the white pixels around the brain (bone) should be removed.

Another possible improvement is the removal of the CSF area since the intensity of the pixels is the same as the ischemic zone and can mislead the model.

Finally, the limited size of the dataset is one of the causes for the limited performance of the model. In total, slices from only 60 patients were used, with 35 patients associated with class 1 and 25 patients assigned to class 0. One avenue to increase the performance of the model would be to label more data and select the same number of slices from each patient, avoiding overfitting, data leakage, and selection bias.

Acknowledgments. This research was supported by the Portuguese Science Foundation, through the Projects PRELUNA - PTDC/CCI-INF/4703/2021 and UIDB/50021/2020.

References

1. Barman, A., Inam, M.E., Lee, S., Savitz, S., Sheth, S., Giancardo, L.: Determining ischemic stroke from ct-angiography imaging using symmetry-sensitive convolutional networks. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), pp. 1873–1877. IEEE (2019)
2. Barzegar, Z., Jamzad, M.: Fully automated glioma tumour segmentation using anatomical symmetry plane detection in multimodal brain MRI. IET Comput. Vis. **15**(7), 463–473 (2021)
3. Chen, L., Bentley, P., Rueckert, D.: Fully automatic acute ischemic lesion segmentation in DWI using convolutional neural networks. NeuroImage Clin. **15**, 633–643 (2017)
4. Chin, C.L., et al.: An automated early ischemic stroke detection system using CNN deep learning algorithm. In: 2017 IEEE 8th International Conference on Awareness Science and Technology (iCAST), pp. 368–372. IEEE (2017)
5. Dey, S., Dutta, A., Toledo, J.I., Ghosh, S.K., Lladós, J., Pal, U.: SigNet: convolutional siamese network for writer independent offline signature verification. CoRR abs/1707.0 (2017)
6. Gibicar, A., Moody, A.R., Khademi, A.: Automated midline estimation for symmetry analysis of cerebral hemispheres in FLAIR MRI. Front. Aging Neurosci. **13**, 644137 (2021)
7. Goldman, L.W.: Principles of CT and CT technology. J. Nucl. Med. Technol. **35**(3), 115–128 (2007)
8. Herzog, N.J., Magoulas, G.D.: Brain asymmetry detection and machine learning classification for diagnosis of early dementia. Sensors **21**(3), 778 (2021)
9. Lindsay, M.P., et al.: World Stroke Organization (WSO): global stroke fact sheet 2019. Int. J. Stroke **14**(8), 806–817 (2019)

10. Liu, S.X.: Symmetry and asymmetry analysis and its implications to computer-aided diagnosis: a review of the literature. J. Biomed. Inf. **42**(6), 1056–1064 (2009)
11. Liu, X., Zhou, Y., Zhao, J., Yao, R., Liu, B., Zheng, Y.: Siamese convolutional neural networks for remote sensing scene classification. IEEE Geosci. Remote Sens. Lett. **16**(8), 1200–1204 (2019)
12. Melekhov, I., Kannala, J., Rahtu, E.: Siamese network features for image matching. In: 2016 23rd International Conference on Pattern Recognition (ICPR), pp. 378–383. IEEE (2016)
13. Miletto Petrazzini, M.E., Sovrano, V.A., Vallortigara, G., Messina, A.: Brain and behavioral asymmetry: a lesson from fish. Front. Neuroanatomy **14**, 11 (2020)
14. Öman, O., Mäkelä, T., Salli, E., Savolainen, S., Kangasniemi, M.: 3D convolutional neural networks applied to CT angiography in the detection of acute ischemic stroke. Eur. Radiol. Exp. **3**(1), 8 (2019)
15. Phipps, M.S., Cronin, C.A.: Management of acute ischemic stroke. BMJ **368**, l6983 (2020)
16. Pohl, M., et al.: Ischemic stroke mimics: a comprehensive review. J. Clin. Neurosci. **93**, 174–182 (2021)
17. Qiu, K., Ai, Y., Tian, B., Wang, B., Cao, D.: Siamese-ResNet: implementing loop closure detection based on siamese network. In: 2018 IEEE Intelligent Vehicles Symposium (IV), pp. 716–721. IEEE (2018)
18. Ribolsi, M., Daskalakis, Z.J., Siracusano, A., Koch, G.: Abnormal asymmetry of brain connectivity in schizophrenia. Front. Human Neurosci. **8**, 1010 (2014)
19. Ruppert, G.C.S., Teverovskiy, L., Yu, C.P., Falcao, A.X., Liu, Y.: A new symmetry-based method for mid-sagittal plane extraction in neuroimages. In: 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 285–288. IEEE (2011)
20. Vupputuri, A., Dighade, S., Prasanth, P.S., Ghosh, N.: Symmetry determined superpixels for efficient lesion segmentation of ischemic stroke from MRI. In: 2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 742–745. IEEE (2018)
21. Wu, H., Chen, X., Li, P., Wen, Z.: Automatic symmetry detection from brain MRI based on a 2-channel convolutional neural network. IEEE Trans. Cybern. **51**(9), 4464–4475 (2021)
22. Xue, Z., Antani, S., Long, L.R., Demner-Fushman, D., Thoma, G.R.: Window classification of brain CT images in biomedical articles. In: AMIA Annual Symposium proceedings 2012, pp. 1023–1029 (2012)
23. Yu, C.P., Ruppert, G., Nguyen, D., Falcão, A., Liu, Y.: Statistical asymmetry-based brain tumor segmentation from 3D MR images. In: Biosignals (2012)