Context-Aware Convolutional Neural Networks for Stroke Sign Detection in Non-contrast CT Scans

Aneta Lisowska[1,2(✉)], Alison O'Neil[1], Vismantas Dilys[1], Matthew Daykin[1,2], Erin Beveridge[1], Keith Muir[3], Stephen Mclaughlin[2], and Ian Poole[1]

[1] Toshiba Medical Visualization Systems Europe Ltd., Edinburgh, UK
alisowska@tmvse.com
[2] School of Engineering and Physical Sciences,
Heriot-Watt University, Edinburgh, UK
[3] Queen Elizabeth University Hospital, Glasgow, UK

Abstract. Detection of acute stroke signs in non-contrast CT images is a challenging task. The intensity and texture variations in pathological regions are subtle and can be confounded by normal physiological changes or by old lesions. In this paper we investigate the use of contextual information for stroke sign detection. In particular, the appearance of the contralateral anatomy and the atlas-encoded spatial location are incorporated into a Convolutional Neural Network (CNN) architecture. CNNs are trained separately for the detection of dense vessels and of ischaemia. The network performance is evaluated on 170 datasets by cross-validation. We find that atlas location is important for dense vessel detection, but is less useful for ischaemia, whereas bilateral comparison is crucial for detection of ischaemia.

1 Introduction

In acute ischaemic stroke the ultimate goal of treatment is tissue reperfusion via recanalisation of an occluded vessel. This can be successfully achieved, with subsequent improved clinical outcome, by reperfusion therapies or mechanical thrombectomy. However, the offsetting of risks versus benefit in the selection of patients for treatment options is currently suboptimal, and thus there is a recognised need for imaging biomarkers [1]. Whilst the debate of a gold standard imaging profile and optimal modalities continues, the role of non-contrast computed tomography (NCCT) in the emergency setting is undisputed. Beyond its primary role of intracerebral haemorrhage detection, signs of early ischaemic change can be identified. These subtle imaging features include the parenchymal changes as well as the dense vessel sign, which represents causative vessel occlusion. Objective, automated methods to detect and quantify these signs of early ischaemic change hold value in a time critical clinical setting for treatment triage, especially out-of-hours in the absence of a specialist.

Automatic detection of acute stroke signs in NCCT images is challenging. Dense vessels and areas of ischaemia are difficult to detect due to the proximity
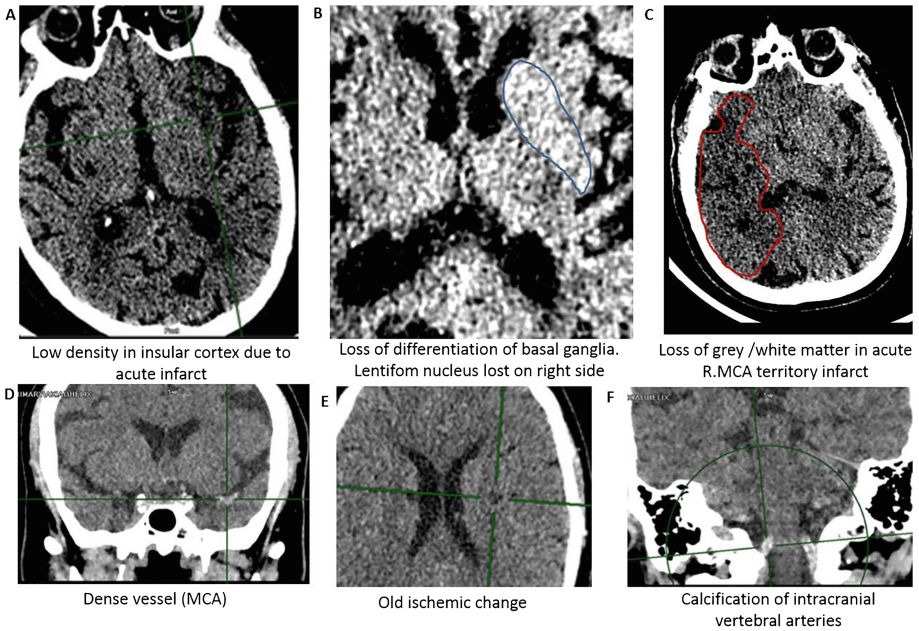
Low density in insular cortex due to acute infarct

Loss of differentiation of basal ganglia. Lentifom nucleus lost on right side

Loss of grey /white matter in acute R.MCA territory infarct

Dense vessel (MCA)

Old ischemic change

Calcification of intracranial vertebral arteries

Fig. 1. A-E show examples of various subtle stroke signs. F shows normal calcification of arteries, which could be confused with dense vessel signs by a naïve classifier.

of bone in the former case, and to the subtlety of intensity and texture changes in the latter case (see Fig. 1).

For the easier problem of haemorrhagic tissue detection in NCCT, common solutions include histogram-based thresholding [2] or clustering [3], followed by morphological operations. More sophisticated methods of stroke sign detection use classifiers with handcrafted features [4]. Chawala *et al.* developed a system for the classification of both ischaemic and haemorrhagic stroke signs [5]. Their system computes slicewise histograms for each hemisphere and classification of the image slice is based on a comparison between the left and right side histograms (see Sect. 5). This bilateral comparison is an example of where the provision of contextual information might be helpful to a system for stroke sign identification.

There have been CNN-based methods suggested for ischaemia lesion segmentation in MR images [6,7] (see Sect. 5), nevertheless the architecture for CT may require different design choices. CT brain images show less structural detail in regions of soft tissue and take on a textured appearance, therefore there are fewer higher-level concepts for a CNN to learn than for MR brain images.

To inform the network design, we observed an experienced neuroradiologist during reading of the NCCT scan. They routinely compared the appearance and Hounsfield Unit intensities of the left and right hemispheres when searching for stroke signs and used their knowledge of brain anatomy to navigate straight to the regions most commonly affected in stroke episodes. Therefore, we hypothesise that incorporation of the bilateral comparison and atlas information into the

CNN architecture might be helpful for the detection of the dense vessels and ischaemia (see Fig. 2). These insights have been previously incorporated in a CNN applied to a thrombus detection problem [8]. In this study we examine the impact of each of the architectural choices on the detection performance of different type of stroke signs, specifically we:

– Demonstrate a CNN-based solution applied to the detection of dense vessel and ischaemic regions in non-contrast CT scans.
– Evaluate the impact of including contralateral features in a CNN architecture for dense vessel and ischaemia detection.
– Evaluate the impact of including atlas location features in a CNN architecture for dense vessel and ischaemia detection.

2   Proposed Solution: A CNN for Stroke Sign Detection

Bilateral Comparison: The exploitation of anatomical symmetry has previously been incorporated in unsupervised approaches to pathology detection.
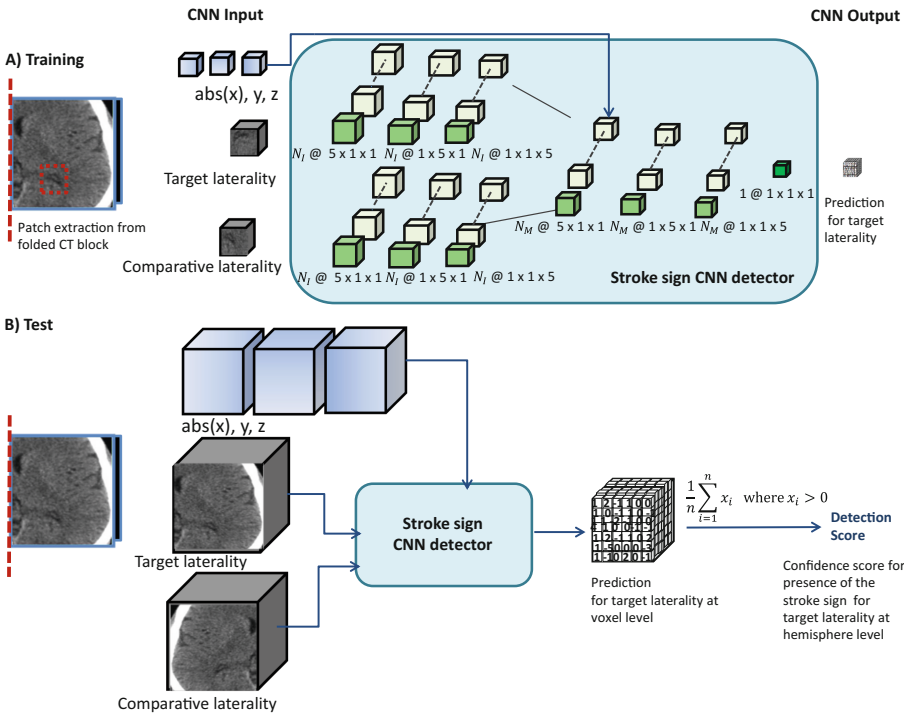

Fig. 2. (A) Schematic of the CNN, including filter sizes and number of layers. Pairs of contralateral 3D image intensity patches are input to the network at training time. Atlas coordinate inputs are fused at the merge point of the intensity channels. (B) Application of the detector at test time. The whole folded block is input to the network, and predictions are generated separately for each side.

Researchers utilised within-organ symmetry for the detection of brain tumours [9] and examined symmetry between a pair of organs for the detection of breast tumours [10]. In these approaches, the pathology is found by searching for the most dissimilar regions between the left and right sides of the organ [9], or by identification of asymmetry between paired organs [10]. We incorporate right and left hemisphere comparison in the CNN architecture by inputting image patches extracted from both hemispheres to parallel CNN channels and allow left/right comparison functionality to emerge as a part of supervised training. To ensure extraction of *corresponding* patches from contralateral regions, we first align the CT volume to the reference dataset, then we extract symmetrical blocks of interest (Fig. 4) about the sagittal midline, and finally we fold the CT block along the midline (Fig. 2A).
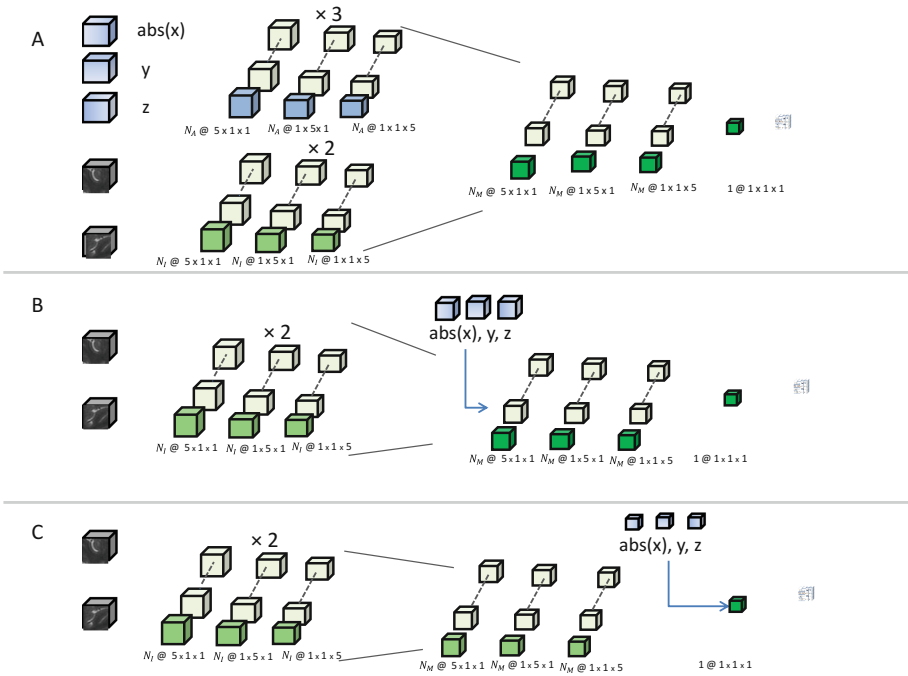

Fig. 3. Adding atlas information at different network stages: (A) Adding the atlas location by creating an additional 3 input channels (×3) alongside two intensity channels (×2). (B) Adding the atlas location midway, at the merge point of the bilateral intensity channels. (C) Adding the atlas location to the pre-classification layer.

Anatomical Context: One way of discovering abnormalities in the image is to compare the patient image with a normative atlas created from healthy examples of anatomy. In this approach, the patient image is registered to the normative atlas and the pathology is identified by the differences between the reference atlas

and the examined image [11]. Anatomical atlases have also been used in supervised medical imaging applications as they provide useful anatomical context information. Researchers have previously employed explicit anatomical context mechanisms when training random forest classifiers [12]. We propose to supply the CNN with this information by adding three channels encoding the x, y and z atlas locations. Furthermore, we investigate the level at which this information should be provided to the network (see Fig. 3).

3D Architecture: NCCT scans are three-dimensional (3D), therefore architectures designed for 2D images cannot be directly applied. When moving to 3D data, 3D convolutions [6,13] may be used in place of 2D convolutions. A 3D CNN could potentially lead to better results than a 2D CNN applied slicewise [13] since information within a 3D neighbourhood may be leveraged at a deep level. In this paper we adopt a 3D CNN, but we apply spatial decomposition of the kernels such that convolutions are applied one dimension at the time. This allows us to reduce the risk of overfitting as the network has smaller number of parameters.

3   Data Sets

We use data from the following studies: ATTEST [14], POSH and WYETH [15]. Ground truth was collected on the acute NCCT scans for 170 patients with suspected acute ischaemic stroke within 6 hours of onset.

Capturing ground-truth for subtle stroke signs is only achievable by manual segmentation, given the diversity of shape, size and location of these signs. 3D Slicer 4.5.0 was used for this task which generated a label map for each of the acute NCCT datasets. Manual segmentations of ischaemia and dense vessels were generated by a clinical researcher under the supervision of an experienced neuroradiologist. Annotations were blind to additional scans (e.g. CT angiography, CT perfusion, follow-up scans) and clinical information except for the radiology report which included laterality of symptoms.

Thromboembolism is most frequently seen in the middle cerebral artery, thus we currently focus our detection on dense vessels within the anterior circulation and on ischaemic regions which lie downstream in the associated vascular territory (see Fig. 4).

4   Experiments

Data Preparation: Volumetric datasets are pre-aligned to a reference dataset, designated as an atlas. The transformation between a given dataset and the reference atlas is discovered via landmarks, which are detected in the novel volume by a random forest as proposed in [16]. The datasets are isotropically resampled to 1 mm per voxel for dense vessels and 2 mm per voxel for ischaemia. Blocks of interest are extracted symmetrically about the sagittal midline and folded. Volume intensities are clamped in the range $\{-125, 225\}$ HU for dense vessels
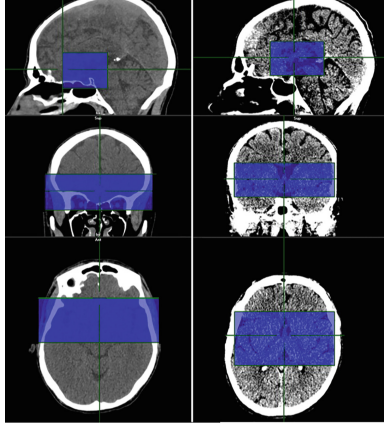
Fig. 4. Left: anterior circulation block used in dense vessel detection experiments. Right: subcortex block used in ischaemia detection experiments.

and $\{0, 80\}$ HU for ischaemia to imitate the typical window levels chosen by a radiologist when searching for these signs.

CNN Model: There are two data input channels for 3D patches selected symmetrically from the left and right hemispheres. Weights are shared between intensity channels for ischaemia, but not for dense vessels. We also insert x, y and z atlas coordinates into the architecture. Each full convolution operation comprises a series of orthogonal one-dimensional convolutions, with kernels of $5 \times 1 \times 1$, $1 \times 5 \times 1$ and $1 \times 1 \times 5$. $N_I = 32$ kernels are used for the data channels and $N_A = 4$ kernels are used for the atlas channels. Channels are then merged and another convolution operation is applied, with $N_M = 32$ kernels in the case of dense vessels and $N_M = 2$ in the case of ischaemia. The number of kernels and the filter size were chosen empirically. ReLU activation functions are used. The output is fully convolutional allowing for the efficient prediction of all voxels of the dataset in a single pass (see Fig. 2). The models were implemented in Python using the Keras [17] library built on top of Theano [18].

CNN Training: We use 71 datasets for training, 48 for validation and 51 for testing. To compensate for the large imbalance between abnormal voxels and normal voxels, we adopt a biased patch selection process and use a weighting factor w, which is defined as the ratio of normal to abnormal voxels in the training set. The patches for training of ischemia are of size $18 \times 18 \times 18$ voxels and for dense vessels are $24 \times 24 \times 24$ voxels. There is inevitably some uncertainty around precise ground truth segmentation of stroke signs in NCCT, due to their diffuse appearance, especially in case of ischaemia. Therefore we are not interested in refining the segmentation boundary, but in the detection of the presence or absence of a stroke sign. We mark the regions around the pathology as "do not care". We adjust the loss function to ignore the voxels with this label. For implementation, we use labels of -1 and $+1$ for the normal and abnormal classes

respectively and use the label 0 to represent the "do not care" voxels for which loss is not computed. The border is created by a dilation operation and it is 1 mm thick for dense vessel and 6 mm for ischaemia. Training is performed at the voxel level, meaning that each voxel in the patch is a training example. Training is performed using the Adam optimiser [19] on normalised data samples, to optimise the squared *hinge loss* function, with a learning rate of 0.001, a momentum of 0.9 and L2 regularisation of 0.002. For training and testing of the CNN classifier we used a Titan X GPU.

Post-processing: We train the dense vessel and ischaemia detectors at the voxel level as this gives us a larger number of training samples. Nevertheless, we are interested in determining whether the stroke sign is present or absent at the level of the brain hemisphere. To arrive at the detection confidence score at half block level we compute the mean of all voxel level predictions with a score above 0. This threshold was necessarily as the majority of voxels are negative (normal tissue) and had more influence over the mean than the confident positive detection of smaller stroke signs.

Evaluation: We evaluate the performance of the ischaemia (Sect. 4.1) and dense vessel (Sect. 4.2) detectors at the brain hemisphere level (half-patient) in terms of the area under the curve for the Receiver Operating Characteristic (ROC AUC) and the Precision-Recall (PR AUC) curve. It is suggested that PR curves should be used when the positive class samples are rare compared to the negative class samples [20], because precision is more sensitive to any change in the number of false positives, while specificity is not due to the large number of negative samples. To determine whether the inclusion of bilateral comparison or the atlas are helpful in stroke sign detection we compare the detection performance of four different CNN architectures.
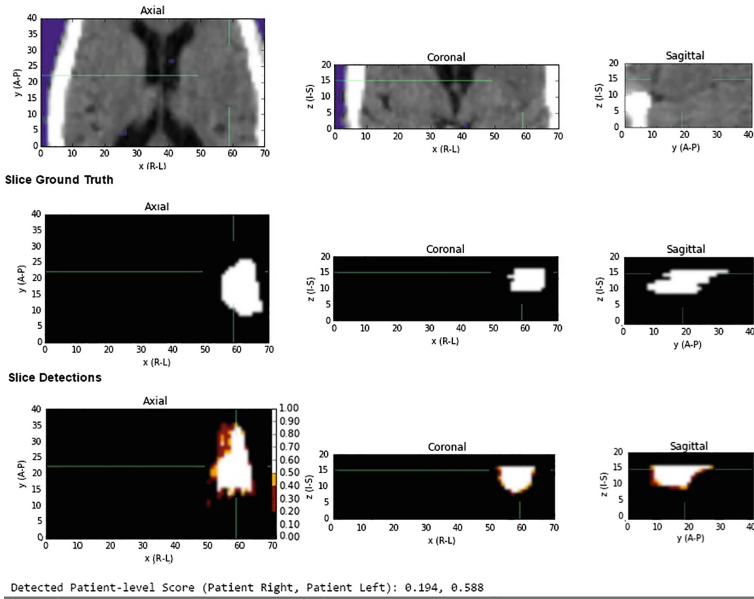
4.1   Ischaemia Detection

Table 1 presents the detection results obtained for ischaemia detection. Incorporation of bilateral features in the CNN architecture improves detection performance compared to a single intensity channel CNN, which is not trained with pairs of contralateral patches, but follows standard single intensity patch training. Incorporation of atlas information is not helpful for the ischaemia detection task as there is no significant different in performance between CNN architectures with and without atlas coordinates provided to the network.

4.2   Dense Vessel Detection

Table 2 presents the dense vessel detection results. The addition of atlas coordinates to the network has a large impact on the performance of the dense vessel detector. It gives rise to larger improvements than inclusion of bilateral channels, which has little impact on dense vessel detection. We incorporated atlas information in the network at three different levels (see Fig. 3) to investigate if the point at which this information is injected to the network affects the detection results. The earliest incorporation of the atlas led to the best performance.
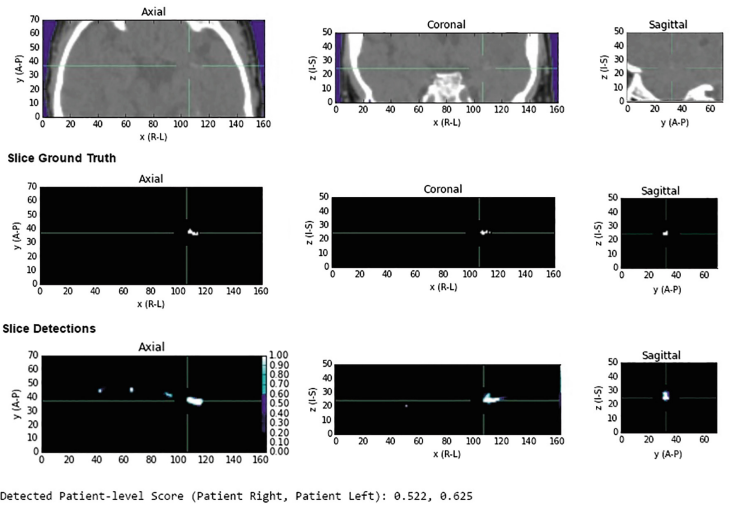
Fig. 5. Examples of ischaemia detection (top) and dense vessel detection (bottom). For each example, we display the volume slice with the highest number of abnormal voxels. Nevertheless the predictions are computed for all voxels in the volume. Image brightness corresponds to confidence level. For dense vessels we selected an example with both true and false positive detections. The false positives are detected with lower confidence, as indicated by the brightness. The *hemisphere*-level scores for the right and left hemisphere are printed below each example.

Table 1. Ischaemia detection results. Detection of the presence/absence of a sign was performed at the hemisphere level. Each experiment was run 3 times with different random seeds, and we report the mean and standard deviation

Model	ROC AUC [std]	PR AUC [std]
Bilateral CNN + atlas	0.915 [0.006]	0.783 [0.014]
Bilateral CNN	0.912 [0.007]	0.782 [0.006]
Single intensity channel CNN + atlas	0.738 [0.003]	0.483 [0.021]
Single intensity channel CNN	0.743 [0.012]	0.461 [0.022]

Table 2. Dense vessel detection results. Detection of the presence/absence of a sign was performed at the hemisphere level. Each experiment was run 3 times with different random seeds, and we report the mean and standard deviation.

Model	ROC AUC [std]	PR AUC [std]
Bilateral CNN + atlas (A)	0.964 [0.005]	0.898 [0.029]
Bilateral CNN + atlas (B)	0.950 [0.011]	0.817 [0.062]
Single intensity channel CNN + atlas (A)	0.936 [0.026]	0.790 [0.063]
Bilateral CNN + atlas (C)	0.927 [0.019]	0.718 [0.072]
Bilateral CNN	0.891 [0.011]	0.691 [0.036]
Single intensity channel CNN	0.876 [0.013]	0.514 [0.060]

5 Related Work and Discussion

In this paper we were inspired by the workflow of a radiologist, to investigate the role of contextual information when interpreting NCCT scans to identify acute stroke signs. Frequently in medical images, local and global spatial context is informative. Pathology might have a characteristic distribution relative to different anatomical structures and tissues. This is the case with dense vessel signs which are expected to appear along the vasculature. Precise anatomical location is less useful for ischaemic changes, which are diffuse in appearance and more texture-like than dense vessels, therefore the type of context required to detect those signs may differ.

Some researchers have tried the use of foveation [21] or the similar method of non-uniform sampling [22] in order to capture anatomical contextual information. Foveation refers to under-sampling of pixels closed to a window boundary, whilst keeping the central pixels at the original sampling level. Ciresan *et al.* [21] propose that this method is well suited to the training of networks whose task is to classify the central pixel of the window, as the network is then forced to ignore fine details at the periphery of the window, whilst still having access to general context information. The challenge with this approach is that it assumes classification of one voxel at a time using sliding window approach, which leads to long detection times. Although these methods may offer efficient training, it

is not obvious how to incorporate them in a fully convolutional network, and application to stroke sign detection might not be feasible due to constraints on detection time in the clinical setting.

In tackling ischaemia lesion detection in MR, a couple of authors have pursued a dual-pathway approach, in which two pathways are devoted to local and global context respectively, before combining at the pre-classification stage. Kamnitsas *et al.* [6] achieves this through the use of patches at two different image scales, whereas Dutil *et al.* [7] uses kernels of different sizes. As the pathways are disjoint the network can process them in parallel, which is convenient for run time speed up. We adopted a similar architectural design, but designated the two input pathways for bilateral inputs, rather than two scales.

We are not the first authors to notice that comparison of the appearance of the left and right hemisphere might be helpful in ischaemia detection. Chawla *et al.* proposed a two stage system for differentiation between chronic infarcts, haemorrhagic and ischaemic stroke [5]. For each image slice, histograms of intensity values are computed for each hemisphere and hemisphere similarity is assessed using the correlation coefficient as a measure. In the first stage they classify the datasets containing chronic infarcts or haemorrhage by histogram thresholding. In the second stage they employ a wavelet decomposition of the histograms for slices assigned to neither of the categories to further discriminate between acute infarcts and normal slices. The authors report average recall of 90% for acute stoke categorisation at slice level [5]. By contrast, our system first produces a prediction for all voxel in the block so we are able to highlight the region suspected of ischaemia alongside the hemisphere detection score computed at half patient level (see Fig. 5).

Since we have a strong prior belief about the spatial location of dense vessels, we also injected explicit atlas location information. In this we follow [23] who, for coronary calcium scoring in cardiac CT angiography, registered each image to an atlas image and fed the resulting atlas coordinates $\{x, y, z\}$ of each voxel as three additional inputs to the network at the dense fully connected layer. We build upon their work by showing that atlas coordinates may be effectively utilised even if the fully connected layers are implemented as convolutions. Furthermore, we investigated at which level to inject atlas information to the architecture and we found that providing it at the input layer alongside the intensity inputs led to best performance. This agrees with the finding of Havaei et al., who investigated at which level the information from the output of an initial CNN should be incorporated to a second CNN in a cascade of classifiers. They also tried providing this additional information at three different levels and found that concatenation of the prediction of the first classifier with intensity patches at the input to the second classifier gave best brain tumour segmentation in MR [24].

6   Conclusion and Future Work

The design of the CNN classifier depends on the task at hand. We have investigated the type of contextual information required for two different types of

stroke sign and suggested how this information might be incorporated in the CNN architecture. We found that providing atlas information is helpful for dense vessel detection, where the signs appear in typical locations. Furthermore, the earlier this information is provided to the network, the sooner the detector is able to focus in on the critical spatial region. Atlas coordinates are less useful for detecting ischaemic regions since they vary in location and size. However, the incorporation of contralateral features in our CNN design enables bilateral comparison, and leads to more successful ischaemia detection.

In future work it would be interesting to compare the CNN-based solutions designed for ischaemia lesion segmentation in MR, such as [6], with the proposed CNN architecture for stoke sign detection in NCCT. This compares solutions designed for similar tasks but for different imaging modalities and so would enable evaluation of the extent to which the CNN architectures designs should be modality specific.

References

1. Wintermark, M., Albers, G.W., Broderick, J.P., Demchuk, A.M., Fiebach, J.B., Fiehler, J., Grotta, J.C., Houser, G., Jovin, T.G., Lees, K.R., et al.: Acute stroke imaging research roadmap II. Stroke **44**(9), 2628–2639 (2013)
2. Chan, T.: Computer aided detection of small acute intracranial hemorrhage on computer tomography of brain. Comput. Med. Imaging Graph. **31**(4), 285–298 (2007)
3. Dhawan, A.P., Loncaric, S., Hitt, K., Broderick, J., Brott, T.: Image analysis and 3-D visualization of intracerebral brain hemorrhage. In: Proceedings of Sixth Annual IEEE Symposium on Computer-Based Medical Systems, pp. 140–145. IEEE (1993)
4. Usinskas, A., Dobrovolskis, R.A., Tomandl, B.F.: Ischemic stroke segmentation on CT images using joint features. Informatica **15**(2), 283–290 (2004)
5. Chawla, M., Sharma, S., Sivaswamy, J., Kishore, L.: A method for automatic detection and classification of stroke from brain CT images. Eng. Med. Biol. Soc. **2009**, 3581–3584 (2009)
6. Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B.: Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Med. Image Anal. **36**, 61–78 (2017)
7. Dutil, F., Havaei, M., Pal, C., Larochelle, H., Jodoin, P.-M.: A convolutional neural network approach to brain segmentation. In: Ischemic Stroke Lesion Segmentation, p. 53 (2015)
8. Lisowska, A., Bereridge, E., Muir, K., Poole, I.: Thrombus detection in ct brain scans using a convolutional neural network. In: Proceedings of the 10th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2017), Bioimaging, vol. 2, pp. 24–33. SCITEPRESS (2017)
9. Hasan, A., Meziane, F., Khadim, M.: Automated segmentation of tumours in MRI brain scans. In: Proceedings of the 9th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2016), pp. 55–62. SCITEPRESS (2016)
10. Erihov, M., Alpert, S., Kisilev, P., Hashoul, S.: A cross saliency approach to asymmetry-based tumor detection. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 636–643. Springer, Cham (2015). doi:10.1007/978-3-319-24574-4_76

11. Doyle, S., Vasseur, F., Dojat, M., Forbes, F.: Fully automatic brain tumor segmentation from multiple MR sequences using hidden Markov fields and variational EM. In: Proceedings of NCI-MICCAI BraTS, pp. 18–22 (2013)
12. O'Neil, A., Murphy, S., Poole, I.: Anatomical landmark detection in CT data by learned atlas location autocontext. In: Medical Image Understanding and Analysis (MIUA), pp. 189–194 (2015)
13. Payan, A., Montana, G.: Predicting Alzheimers disease: a neuroimaging study with 3D convolutional neural networks. arXiv:1502.02506 (2015)
14. Huang, X., Cheripelli, B.K., Lloyd, S.M., Kalladka, D., Moreton, F.C., Siddiqui, A., Ford, I., Muir, K.W.: Alteplase versus tenecteplase for thrombolysis after ischaemic stroke (ATTEST): a phase 2, randomised, open-label, blinded endpoint study. Lancet Neurol. **14**(4), 368–376 (2015)
15. Wardlaw, J.M., Muir, K.W., Macleod, M.J., Weir, C., McVerry, F., Carpenter, T., Shuler, K., Thomas, R., Acheampong, P., Dani, K., Murray, A.: Clinical relevance and practical implications of trials of perfusion and angiographic imaging in patients with acute ischaemic stroke: a multicentre cohort imaging study. J. Neurol. Neurosurg. Psychiatry **84**(9), 1001–1007 (2013). http://jnnp.bmj.com/content/84/9/1001
16. Dabbah, M.A., Murphy, S., Pello, H., Courbon, R., Beveridge, E., Wiseman, S., Wyeth, D., Poole, I.: Detection and location of 127 anatomical landmarks in diverse CT datasets. In: SPIE Medical Imaging, pp. 903415–903415. International Society for Optics and Photonics (2014)
17. Chollet, F.: Keras (2015). https://github.com/fchollet/keras
18. Theano Development Team. Theano: a Python framework for fast computation of mathematical expressions. arXiv e-prints, vol. abs/1605.02688, May 2016
19. Kingma, D., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
20. Davis, J., Goadrich, M.: The relationship between precision-recall and ROC curves. In: Proceedings of the 23rd International Conference on Machine Learning, ICML 2006, pp. 233–240 (2006)
21. Ciresan, D., Giusti, A., Gambardella, L.M., Schmidhuber, J.: Deep neural networks segment neuronal membranes in electron microscopy images. In: Advances in Neural Information Processing Systems, pp. 2843–2851 (2012)
22. Ghafoorian, M., Karssemeijer, N., Heskes, T., van Uder, I., de Leeuw, F., Marchiori, E., van Ginneken, B., Platel, B.: Non-uniform patch sampling with deep convolutional neural networks for white matter hyperintensity segmentation. In: 2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI), pp. 1414–1417. IEEE (2016)
23. Wolterink, J.M., Leiner, T., Viergever, M.A., Išgum, I.: Automatic coronary calcium scoring in cardiac CT angiography using convolutional neural networks. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9349, pp. 589–596. Springer, Cham (2015). doi:10.1007/978-3-319-24553-9_72
24. Havaei, M., Davy, A., Warde-Farley, D., Biard, A., Courville, A., Bengio, Y., Pal, C., Jodoin, P.-M., Larochelle, H.: Brain tumor segmentation with deep neural networks. Med. Image Anal. **35**, 18–31 (2017)