

Asymmetry Disentanglement Network for Interpretable Acute Ischemic Stroke Infarct Segmentation in Non-Contrast CT Scans

Haomiao Ni^{1*}, Yuan Xue^{2*}, Kelvin Wong³, John Volpi⁴,
Stephen T.C. Wong³, James Z. Wang¹, and Xiaolei Huang¹

¹ The Pennsylvania State University, University Park, Pennsylvania, USA

² Johns Hopkins University, Baltimore, Maryland, USA

³ TT and WF Chao Center for BRAIN & Houston Methodist Cancer Center,
Houston Methodist Hospital, Houston, Texas, USA

⁴ Eddy Scurlock Comprehensive Stroke Center, Department of Neurology,
Houston Methodist Hospital, Houston, Texas, USA

Abstract. Accurate infarct segmentation in non-contrast CT (NCCT) images is a crucial step toward computer-aided acute ischemic stroke (AIS) assessment. In clinical practice, bilateral symmetric comparison of brain hemispheres is usually used to locate pathological abnormalities. Recent research has explored asymmetries to assist with AIS segmentation. However, most previous symmetry-based work mixed different types of asymmetries when evaluating their contribution to AIS. In this paper, we propose a novel Asymmetry Disentanglement Network (ADN) to automatically separate pathological asymmetries and intrinsic anatomical asymmetries in NCCTs for more effective and interpretable AIS segmentation. ADN first performs asymmetry disentanglement based on input NCCTs, which produces different types of 3D asymmetry maps. Then a synthetic, intrinsic-asymmetry-compensated and pathology-asymmetry-salient NCCT volume is generated and later used as input to a segmentation network. The training of ADN incorporates domain knowledge and adopts a tissue-type aware regularization loss function to encourage clinically-meaningful pathological asymmetry extraction. Coupled with an unsupervised 3D transformation network, ADN achieves state-of-the-art AIS segmentation performance on a public NCCT dataset. In addition to the superior performance, we believe the learned clinically-interpretable asymmetry maps can also provide insights towards a better understanding of AIS assessment. Our code is available at https://github.com/nihaomiao/MICCAI22_ADN.

1 Introduction

Stroke is one of the leading causes of death and disability worldwide [10]. In the United States, about 795,000 people experience a new or recurrent stroke every year, and 87% of all are acute ischemic strokes (AIS) [27]. Non-contrast CT

* These authors contributed equally to this work.




Fig. 1: Illustration of (a) different types of asymmetries in the human brain, and (b) our proposed asymmetry disentanglement framework. Asymmetry maps A , P , and Q in (b) are normalized for better visualization. (a) and (b) are from the same testing CT scan.

(NCCT) images are routinely used to assess the extent of infarction in AIS patients [23]. For computer-aided AIS estimation in NCCT scans, accurate infarct segmentation is a crucial step. However, it is challenging to segment AIS infarct in NCCT scans. First, NCCT is more difficult to process than other medical image modalities such as MRI due to the low signal-to-noise and contrast-to-noise ratios of brain tissues [17]. Second, infarct regions can be confounded by normal physiologic changes, and the density and texture variations in involved brain areas may be subtle [19]. In clinical practice, to differentiate subtle abnormalities, clinicians often locate suspicious regions by comparing bilateral differences along the mid-sagittal axis (Fig. 1a).

Leveraging such prior knowledge that bilateral asymmetries can indicate potential lesions, recent symmetry-based AIS segmentation approaches [5, 16, 19, 22, 23, 28] have shown impressive progress. Clèrigues *et al.* [9] exploited brain hemisphere symmetry by inputting CT, CT perfusion images, and their horizontally flipped versions to a 2D-patch-based CNN to segment AIS. Kuang *et al.* [15] proposed a dense multi-path contextual GAN (MPC-GAN) to integrate bilateral intensity difference, infarct location probability, and distance to cerebrospinal fluid (CSF) for AIS segmentation. Instead of exploring symmetry at the image level, Liang *et al.* [19] introduced a symmetry-enhanced attention network (SEAN) to segment AIS in NCCT images. The authors first ran a 2D alignment network to transform input images to be bilaterally quasi-symmetric in axial view. Then a symmetry-enhanced attention module was employed to capture both in-axial and cross-axial symmetry at the feature level. Though achieving promising results, most existing work simply fused all asymmetries, ignoring the fact that specific asymmetries caused by non-pathological factors cannot reveal clinical findings. For instance, in Fig. 1a, the asymmetries be-

tween the two yellow boxes indicate left-sided infarction. However, differences between the green circles originate from the natural variation of the brain. Several recent approaches [4, 6] explored how to highlight semantically-meaningful pathological asymmetries at the feature level to help identify abnormalities in X-ray or MR images. However, such feature-level pathological asymmetry has poor interpretability, making them less useful in clinical practice. Also, [4, 6] both focused on 2D asymmetry analysis, which may fail to make full use of 3D spatial contextual information when applied to CT scans. Moreover, clinicians usually only look for asymmetry in voxels of the same tissue type, *e.g.*, within the grey matter or within the white matter. Such critical domain knowledge has not been investigated or utilized in the literature.

In this paper, we aim to address the aforementioned weaknesses in existing methods by proposing a novel asymmetry disentanglement network (ADN) to automatically separate pathological and intrinsic anatomical asymmetries in NCCT scans for AIS segmentation with tissue-type awareness. ADN performs asymmetry disentanglement at the image level and achieves high interpretability by directly outputting learned pathological and anatomical 3D asymmetry maps for clinical examination. Furthermore, when asymmetry maps are applied to remove different types of asymmetries from the original NCCT, pathology-salient or pathology-compensated images can be obtained (Fig. 1b). Validated by our experiments, performing segmentation on pathology-salient NCCTs shows a noticeable improvement over using original images. A key novelty of our framework is inspired by the observation that non-pathological asymmetries can be due to intensity differences between different types of tissues, uneven distribution of CSF, or inherent brain anatomical asymmetry. Thus we design a tissue-aware regularization loss function to incorporate tissue-type information into the training of ADN, which further encourages ADN to extract pathologically meaningful asymmetries. Coupled with ADN, an unsupervised 3D transformation network is also proposed to align NCCT scans to obtain the mid-sagittal plane for asymmetry estimation. We conduct extensive experiments on a public NCCT dataset AISD [19], and the results show that our proposed ADN can achieve state-of-the-art performance while successfully disentangling different types of asymmetries in a clinically interpretable way.

2 Methodology

Fig. 2 shows the overview of our proposed framework ADN. In general, ADN includes three 3D volume-based modules: transformation network \mathcal{T} for input alignment and the subsequent total asymmetry estimation, asymmetry extraction network \mathcal{D} to detect pathological asymmetries and further separate different types of asymmetries, and segmentation network \mathcal{F} for the final AIS segmentation. Network \mathcal{T} is mainly based on a convolutional encoder while both \mathcal{D} and \mathcal{F} adopt an encoder-decoder architecture. For the input deskulled NCCT scan I , we first employ an unsupervised 3D transformation network \mathcal{T} to align I and generate a bilaterally quasi-symmetric scan X in which the mid-sagittal plane




Fig. 2: Overview of the asymmetry disentanglement network (ADN). Tissue-type information is generated by SPM12 [2] and is only required during training.

S is in the vertical center of the 3D volume space. We then feed X into the network \mathcal{D} to extract pathological asymmetry map P . P is later used to calculate the intrinsic anatomical asymmetry map and then synthesize an intrinsic-asymmetry-compensated scan \hat{X} . We subsequently input \hat{X} to network \mathcal{F} to obtain the segmentation result \hat{Y} . We finally apply inverse transformation \mathcal{T}^{-1} to \hat{Y} for getting the final output O corresponding to the input I . During the training, we also leverage tissue segmentation results generated from an existing tool to provide extra regularization for network \mathcal{D} . The voxel intensities of NCCT images are normalized to be in the range of $[0, 1]$ for network processing and asymmetry computation.

Transformation Network \mathcal{T} . Since the input deskulled NCCT scan I is usually not perfectly centered in the volume space, to make better use of bilateral symmetry, we propose an unsupervised 3D transformation network \mathcal{T} to align the mid-sagittal plane S to be in the vertical center of volume space for subsequent total asymmetry estimation. Different from the 2D alignment network in [19], which may be sensitive to texture variations in some slices, \mathcal{T} utilizes 3D spatial contextual information and outputs parameters $\alpha \in \mathbb{R}^6$, which represents an affine transformation $(\mathbf{R}, \mathbf{t}) \in \text{SE}(3)$, where $\alpha_{1:3}$ are rotation angles and $\alpha_{4:6}$ are translations along x , y , and z axes. We then apply the output parameters α to generate bilaterally quasi-symmetric CT image $X = \mathcal{T}(I)$ using parameterized sampling grid [12]. Intuitively, when S is transformed to align with the vertical center plane, X should have the minimum difference from its horizontally flipped version X' . Thus we adopt the following loss function to train the 3D transformation network \mathcal{T} :

$$L_{\mathcal{T}} = \|X - X'\|_1, \quad (1)$$

where $\|\cdot\|_1$ is L1 loss. Note that \mathcal{T} computes the 3D alignment in contrast to 2D, and no annotation is required for the training of \mathcal{T} .

Asymmetry Disentanglement. Based on X and its self-mirrored version X' , we can compute a total asymmetry map A and further employ an asymmetry extraction network \mathcal{D} to extract pathological asymmetry map P from A . Generally, the ischemic stroke areas appear darker than their contralateral regions in NCCT images (see Fig. 1a). Thus voxels that are darker (*i.e.*, having lower

intensity) in X than X' are suspicious voxels that could correspond to stroke lesions. To obtain these asymmetric darker voxels in X , we first compute the total asymmetry map A by:

$$A_i = \max(X'_i - X_i, 0), \quad (2)$$

where i is any voxel inside X . Note that voxels that are darker on one side of the midsagittal plane have positive values in A , whereas the value of other voxels is set to zero. We obtain the pathological asymmetry map P through the trained network \mathcal{D} , *i.e.*, $P = \mathcal{D}(X)$. Then the map of asymmetry due to intrinsic anatomical asymmetry is $Q = A - P$. Again, in both P and Q asymmetric voxels have positive values and symmetric voxels have zero values. Next, we generate a new image \hat{X} in which the intrinsic anatomical asymmetry is compensated or neutralized so that there is only pathological asymmetry in \hat{X} . \hat{X} is defined as:

$$\hat{X} = X + Q = X + A - P. \quad (3)$$

Note that although \hat{X} does not exist in reality since actual brain CTs always contain some normal anatomical asymmetry, using this synthetic \hat{X} can make pathological asymmetries more salient and thus make accurate infarct segmentation easier as we will demonstrate in our experiments. Similarly, we can compensate for pathological asymmetry and generate a synthetic image $\tilde{X} = X + P$ that contains only normal anatomical asymmetry. Our proposed ADN encourages learned anatomical asymmetry and pathological asymmetry to be meaningful via implicit supervision from the final segmentation so that they should not affect pathology information presented in \tilde{X} or lead to false positives. Example asymmetry disentanglement results are shown in Figs. 1b and 4. We then input \hat{X} to network \mathcal{F} to obtain segmentation map \hat{Y} . The final segmentation O corresponding to the original I can be calculated by applying inverse transformation \mathcal{T}^{-1} to \hat{Y} .

Training of Pathological Asymmetry Extraction Network \mathcal{D} and Segmentation Network \mathcal{F} . Intuitively, one can utilize the infarct ground truth map G to train both \mathcal{D} and \mathcal{F} with binary cross-entropy loss:

$$L_{\mathcal{F}} = \mathcal{L}_{\text{bce}}(O, G) = \mathcal{L}_{\text{bce}}(\mathcal{T}^{-1}(\hat{Y}), G), \quad (4)$$

Since $\hat{Y} = \mathcal{F}(\hat{X}) = \mathcal{F}(X + A - \mathcal{D}(X))$, which is related to both \mathcal{D} and \mathcal{F} , using the loss $L_{\mathcal{F}}$ can update both \mathcal{D} and \mathcal{F} . However, this can lead to a trivial solution $P = \mathcal{D}(X) = A$, such that \hat{X} is equivalent to X , downgrading the proposed model to a regular segmentation network. To ensure pathological asymmetry map P is clinically meaningful, we design a novel tissue-aware regularization loss function to provide extra constraints for the training of \mathcal{D} . This loss function is motivated by the fact that it is only meaningful to examine pathological asymmetry when the pair of mirroring voxels belong to the same tissue type. Thus P resulting from the network \mathcal{D} should exclude asymmetry in voxels whose mirrored counterparts belong to a different tissue type. To utilize such same-tissue constraint, during training, we first employ an off-the-shelf

tool, statistical parametric mapping SPM12 [2, 3] and a CT brain atlas [24] to segment GM, WM, and CSF regions from X ; these regions are represented as binary maps R_{GM} , R_{WM} and R_{CSF} . Then we improve these tissue segmentation results by removing the known stroke regions (based on ground truth G) from the GM, WM and CSF regions. The tissue-aware regularization loss for training \mathcal{D} is designed by considering the following loss terms:

- Tissue-type constraint loss $L_{\text{tissue}} = \|P \cdot (R_{GM} + R_{CSF})\|_1$, which aims to make P exclude asymmetric voxels that belong to different tissue types and regions containing CSF. In particular, R_{GM} indicates GM voxels whose mirrored counterparts are WM voxels or WM voxels whose mirrored counterparts are GM voxels; R_{GM} can be obtained by computing the intersection of horizontally flipped R_{GM} and the original R_{WM} .
- Size constraint loss $L_{\text{size}} = \|\text{mean}(P) - \text{mean}(\mathcal{T}(G))\|_1$, which aims to keep P be of similar average size as ground truth.
- Asymmetry loss $L_{\text{asymmetry}} = \|P \cdot (\mathbf{1} - A)\|_1$, which is to constrain all symmetric regions (non-zero in $(\mathbf{1} - A)$) to have zero value in P so that P has non-zero only for asymmetric voxels.
- Intrinsic asymmetry loss $L_{\text{intrinsic}} = -\|\text{mean}(A - P)\|_1$, which is to encourage intrinsic anatomical asymmetries $Q = A - P$ to be as large as possible to contain all non-pathological asymmetries.

The final regularization loss function to train \mathcal{D} is calculated by:

$$L_{\mathcal{D}} = L_{\text{tissue}} + L_{\text{size}} + L_{\text{asymmetry}} + L_{\text{intrinsic}} . \quad (5)$$

So the total loss function to jointly train \mathcal{D} and \mathcal{F} is:

$$L_{\mathcal{DF}} = L_{\mathcal{F}} + \lambda L_{\mathcal{D}} , \quad (6)$$

where λ is a scaling factor. For $L_{\mathcal{F}}$, we also add an extra generalized Dice loss function [25] to alleviate the class imbalance issue. Note that we only use tissue-type information during training. Using $L_{\mathcal{DF}}$, network \mathcal{D} is trained to help maximize the segmentation accuracy of network \mathcal{F} under the clinical constraints encoded in $L_{\mathcal{D}}$. Thus ADN can learn how to automatically separate different kinds of asymmetries in a clinically interpretable way.

3 Experiments

We validate our proposed method using AISD [19], a public non-contrast CT (NCCT) dataset of acute ischemic stroke (AIS), which includes 345 training scans and 52 testing scans. Ischemic lesions are manually contoured on NCCT by a doctor using MRI scans as the reference standard. The xyz spacing values are various in this dataset, where the x - and y -spacing values are $0.40 \sim 2.04$ mm, and z -spacing (slice thickness) varies from 3 to 10 mm. To ensure that each voxel represents a consistent volume throughout the whole dataset, after




Fig. 3: Qualitative comparison of different methods on AISD. Segmented regions are marked by red contours.

Table 1: Quantitative comparison of different methods on AISD. “Aligned” means whether the transformation is applied to the input CT to make it bilaterally quasi-symmetric.

Method	Aligned	Dice \uparrow	HD95 (mm) \downarrow
DeepMedic [13]	N	0.4412	58.20
3D U-Net [8]	N	0.4281	42.18
3D ResUNet [18]	N	0.4763	42.55
Clèrigues <i>et al.</i> [9]	Y	0.5051	43.17
SEAN [19]	Y	0.5047	40.07
3D ResUNet [18]	Y	0.4850	39.87
ADN w/o L_{tissue}	Y	0.5090	39.66
ADN (Ours)	Y	0.5245	39.18

skull stripping, we resample all NCCTs to be $1.2 \times 1.2 \times 5 \text{ mm}^3$ and reshape them to be $256 \times 256 \times 40$ using the Python package `nilearn` [1]. To evaluate 3D segmentation performance, we employ two volume-based metrics Dice coefficient and Hausdorff distance with 95% quantile (HD95). The definitions of the metrics can be found in [26].

Implementation Details. Our implementation is based on `PyTorch` [21] framework. We implement transformation network \mathcal{T} with four 3D residual blocks [11] followed by one fully-connected layer with `Tanh` activation to predict parameters α . The range of rotation degrees in xy plane is restricted to be no more than 60° and the translation distance is limited to be no more than half of the image size. For simplicity, we do not consider rotation and translation of input CT in z -axis. As a general framework, ADN can employ various networks as its backbone. Here we adopt the same architecture as 3D ResUNet [18] to implement asymmetry extraction network \mathcal{D} and segmentation network \mathcal{F} . Due to the GPU memory constraint, we first train unsupervised network \mathcal{T} using Eq. (1) and then fix \mathcal{T} and jointly train \mathcal{D} and \mathcal{F} using Eq. (6). The factor λ in Eq. (6) is set to be 10. All networks are trained with `AdamW` optimizer [20] with $(\beta_1, \beta_2) = (0.9, 0.999)$ and 5×10^{-4} weight decay. The initial learning rates are




Fig. 4: Some asymmetry disentanglement results of our proposed ADN. X : aligned NCCT; GT: ground truth; Prediction: AIS segmentation by our method; $\tilde{X} = X + Q$: pathology-salient NCCT; $\tilde{X} = X + P$: pathology-compensated NCCT; A : total asymmetry map; P : pathological asymmetry map; Q : intrinsic asymmetry map. Displayed asymmetry maps are normalized for better visualization.

1×10^{-5} for network \mathcal{T} and 1×10^{-3} for network \mathcal{D} and \mathcal{F} and we adopt the poly learning rate policy [7] with a power of 0.9. We take the whole CT scan of size $256 \times 256 \times 40$ as input and the training batch is set to be 6. We only train models using the training set until convergence. More specifically, we terminate the training once the Dice of the training set remains mostly unchanged for 2 epochs. To help network \mathcal{D} achieve good initialization and facilitate the joint training, at the first 2 epochs, we employ a warm-start strategy by using ground truth G to provide extra supervised cross-entropy loss for training \mathcal{D} .

Result Analysis. We compare our proposed ADN with current state-of-the-art methods, including 3D patch-based network DeepMedic [13], volume-based 3D U-Net [8] and 3D ResUNet [18], and symmetry-aware models Clèrigues *et al.* [9] and SEAN [19]. The comparison results are shown in Table 1 and Fig. 3. The original Clèrigues *et al.* is 2D-patch-based and it utilizes symmetry information by concatenating aligned CT patches with their mirrored versions. Here we adapt it to be 3D-volume-based by inputting both aligned CT and its horizontally flipped version to 3D ResUNet. We reimplement SEAN [19] according to their paper due to the lack of publicly available implementation. For the other models, we follow publicly available implementations. Note that both Clèrigues *et al.* and SEAN are based on aligned CT to better extract bilateral symmetry. All the aligned segmentation results will be inversely transformed to correspond to the original NCCT for final comparison. As shown in Table 1, ADN has achieved the best Dice and HD95, outperforming all other methods. We also show some qualitative results in Fig 3.

To further verify the effectiveness of the proposed asymmetry extraction and compensation module, we conduct an ablation study by training a 3D ResUNet

using aligned CT scans (see the third to the last line in Table 1). The only difference between this baseline and proposed ADN is the input to the segmentation network: 3D ResUNet uses the aligned CT scan (X in Fig. 2) as input while ADN uses the aligned CT scan with intrinsic asymmetry compensated (\hat{X}). Results in Table 1 demonstrate that using pathology-salient input can achieve a noticeably better Dice coefficient. We also demonstrate the effectiveness of tissue-type constraint loss by ablation study. Compared to the ADN model without using L_{tissue} , the full ADN model achieves both better Dice and HD95 scores. In addition to the state-of-the-art AIS segmentation performance, ADN also provides clinically interpretable asymmetry maps. We visualize some asymmetry disentanglement results of ADN in Fig. 4.

4 Discussion and Conclusion

We proposed a novel asymmetry disentanglement network, ADN, to separate different kinds of asymmetries in non-contrast CT (NCCT) images for effective acute ischemic stroke (AIS) infarct segmentation. Equipped with a clinically-inspired tissue-aware regularization loss function, ADN not only learns pathological and intrinsic anatomical asymmetry maps for clinical interpretation but also generates pathology-salient (and intrinsic asymmetry compensated) NCCT images for better AIS detection. We currently focus on NCCT because it is noisier and more challenging than other modalities when dealing with soft tissues. Besides CT, our ADN can be extended to other tasks/modalities that leverage bilateral asymmetries to identify abnormalities. Such tasks include but are not limited to stroke or multiple sclerosis (MS) in brain MRI [4], fractures in pelvic X-Rays [6], and infiltration in chest X-Rays [14]. One limitation of ADN that we observe is that it appears to ignore those bright bleeding spots inside stroke regions. With available annotations, bleeding spots could be detected by a network with inverse intensity change using the same architecture. We plan to explore bleeding spot detection and mitigate their effects in our future work.

References

1. Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., Varoquaux, G.: Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics* **8**, 14 (2014)
2. Ashburner, J., Barnes, G., Chen, C.C., Daunizeau, J., Flandin, G., Friston, K., Kiebel, S., Kilner, J., Litvak, V., Moran, R., et al.: SPM12 manual. Wellcome Trust Centre for Neuroimaging, London, UK **2464**, 4 (2014)
3. Ashburner, J., Friston, K.: Multimodal image coregistration and partitioning—a unified framework. *Neuroimage* **6**(3), 209–217 (1997)
4. Bao, Q., Mi, S., Gang, B., Yang, W., Chen, J., Liao, Q.: MDAN: Mirror difference aware network for brain stroke lesion segmentation. *IEEE Journal of Biomedical and Health Informatics* **26**(4), 1628–1639 (2021)

5. 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: Proceedings of the IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). pp. 1873–1877 (2019)
6. Chen, H., Wang, Y., Zheng, K., Li, W., Chang, C.T., Harrison, A.P., Xiao, J., Hager, G.D., Lu, L., Liao, C.H., Miao, S.: Anatomy-aware Siamese network: Exploiting semantic asymmetry for accurate pelvic fracture detection in X-ray images. In: Proceedings of the European Conference on Computer Vision. pp. 239–255. Springer (2020)
7. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **40**(4), 834–848 (2017)
8. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In: Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 424–432. Springer (2016)
9. Clerigues, A., Valverde, S., Bernal, J., Freixenet, J., Oliver, A., Lladó, X.: Acute ischemic stroke lesion core segmentation in CT perfusion images using fully convolutional neural networks. *Computers in Biology and Medicine* **115**, 103487 (2019)
10. Feigin, V.L., Stark, B.A., Johnson, C.O., Roth, G.A., Bisignano, C., Abady, G.G., Abbasifard, M., Abbasi-Kangevari, M., Abd-Allah, F., Abedi, V., et al.: Global, regional, and national burden of stroke and its risk factors, 1990–2019: A systematic analysis for the Global Burden of Disease Study 2019. *The Lancet Neurology* **20**(10), 795–820 (2021)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
12. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. *Advances in Neural Information Processing Systems* **28** (2015)
13. 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. *Medical Image Analysis* **36**, 61–78 (2017)
14. Kim, M., Park, J., Na, S., Park, C.M., Yoo, D.: Learning visual context by comparison. In: European Conference on Computer Vision. pp. 576–592. Springer (2020)
15. Kuang, H., Menon, B.K., Qiu, W.: Automated infarct segmentation from follow-up non-contrast CT scans in patients with acute ischemic stroke using dense multi-path contextual generative adversarial network. In: Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 856–863. Springer (2019)
16. Kuang, H., Menon, B.K., Sohn, S.I., Qiu, W.: EIS-Net: Segmenting early infarct and scoring ASPECTS simultaneously on non-contrast CT of patients with acute ischemic stroke. *Medical Image Analysis* **70**, 101984 (2021)
17. Kuang, H., Najm, M., Menon, B.K., Qiu, W.: Joint segmentation of intracerebral hemorrhage and infarct from non-contrast CT images of post-treatment acute ischemic stroke patients. In: Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 681–688. Springer (2018)
18. Lee, K., Zung, J., Li, P., Jain, V., Seung, H.S.: Superhuman accuracy on the SNEMI3D connectomics challenge. arXiv preprint arXiv:1706.00120 (2017)

19. Liang, K., Han, K., Li, X., Cheng, X., Li, Y., Wang, Y., Yu, Y.: Symmetry-enhanced attention network for acute ischemic infarct segmentation with non-contrast CT images. In: Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 432–441. Springer (2021)
20. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
21. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems* **32** (2019)
22. Peter, R., Korfiatis, P., Blezek, D., Oscar Beitia, A., Stepan-Buksakowska, I., Horinek, D., Flemming, K.D., Erickson, B.J.: A quantitative symmetry-based analysis of hyperacute ischemic stroke lesions in noncontrast computed tomography. *Medical Physics* **44**(1), 192–199 (2017)
23. Qiu, W., Kuang, H., Teleg, E., Ospel, J.M., Sohn, S.I., Almekhlafi, M., Goyal, M., Hill, M.D., Demchuk, A.M., Menon, B.K.: Machine learning for detecting early infarction in acute stroke with non-contrast-enhanced CT. *Radiology* **294**(3), 638–644 (2020)
24. Rorden, C., Bonilha, L., Fridriksson, J., Bender, B., Karnath, H.O.: Age-specific CT and MRI templates for spatial normalization. *Neuroimage* **61**(4), 957–965 (2012)
25. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Jorge Cardoso, M.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pp. 240–248. Springer (2017)
26. Taha, A.A., Hanbury, A.: Metrics for evaluating 3D medical image segmentation: Analysis, selection, and tool. *BMC Medical Imaging* **15**(1), 1–28 (2015)
27. Virani, S.S., Alonso, A., Aparicio, H.J., Benjamin, E.J., Bittencourt, M.S., Callaway, C.W., Carson, A.P., Chamberlain, A.M., Cheng, S., Delling, F.N., et al.: Heart disease and stroke statistics—2021 update: A report from the American Heart Association. *Circulation* **143**(8), e254–e743 (2021)
28. Wang, Y., Katsaggelos, A.K., Wang, X., Parrish, T.B.: A deep symmetry convnet for stroke lesion segmentation. In: *Proceedings of the IEEE International Conference on Image Processing (ICIP)*. pp. 111–115 (2016)