Detecting Large Vessel Occlusions using Graph Deep Learning

Jad Kassam JAD.KASSAM@FAU.DE
Pattern Regocnition Lab, Friedrich-Alexander University, Erlangen, Germany

Florian Thamm FLORIAN.THAMM@FAU.DE
Pattern Regocnition Lab, Friedrich-Alexander University, Erlangen, Germany
Siemens Healthcare GmbH, Forchheim, Germany

Leonhard Rist LEONHARD.RIST@FAU.DE
Pattern Regocnition Lab, Friedrich-Alexander University, Erlangen, Germany
Siemens Healthcare GmbH, Forchheim, Germany

Oliver Taubmann OLIVER.TAUBMANN@SIEMENS-HEALTHINEERS.COM
Siemens Healthcare GmbH, Forchheim, Germany

Andreas Maier ANDREAS.MAIER@FAU.DE
Pattern Regocnition Lab, Friedrich-Alexander University, Erlangen, Germany

Abstract

Large vessel occlusions (LVO) typically lead to severe ischemia of brain parenchyma. Identifying such LVOs is thus a crucial objective in stroke diagnosis. As shortening the time to treatment is essential for a good outcome, fast automated detection can be a valuable tool in clinical routine. This can be achieved using deep learning approaches. In a CTA scan, an LVO can be detected as an unexpected interruption in the contrast-enhanced vessel tree. These cerebrovascular trees can be represented as graphs and analyzed using graph deep learning (GDL) methods. Representing the vasculature as a graph instead of a (very sparsely populated) Euclidean volume massively reduces the model input dimensionality, which promotes time and memory efficiency. In this study, we investigate the use of graph deep learning methods for classifying the presence of a large vessel occlusion compared to state-of-the-art image-based methods. Furthermore, the influence of vascular attributes and different graph topologies is investigated. The proposed model achieves performance comparable to the baseline with an accuracy of 0.95 and an AUC of 0.89. Compared to the image-based approach, the graph-based approach is ten times faster and requires 80% less memory.

Keywords: Graph Deep Learning . Computed Tomography Angiography · Stroke

1. Introduction

Acute ischemic strokes (AIS) are the second leading cause of death and the leading cause of physical disability worldwide and have increased over the last decade (Murphy and Werring, 2020). One of the leading causes of an AIS is large vessel occlusions (LVO), accounting for more than one-third of all AIS cases (Malhotra et al., 2017) and responsible for 90% of mortality after an AIS. Usually, these occlusions occur in the major arteries around the Circle of Willis (CoW) in the brain, the internal carotid artery (ICA) and/or the middle cerebral artery (MCA). Finding LVOs is the critical objective in the diagnostic process and treatment (Lakomkin et al., 2019). Reperfusion with intravenous thrombolytic medication

therapy or mechanical thrombectomy improves the probability of a disability-free recovery following an AIS. However, the time window for reperfusion is limited (Palaniswami and Yan, 2015; ElTawil and W Mui Keith, 2017); hence a fast diagnosis of LVO-positive patients is required. Computed Tomography Angiography (CTA) represents one of the essential modalities to visualize the cerebrovascular tree by injecting a contrast medium, which is beneficial in identifying LVOs. An automated classification of patients suffering from LVOs based on a CTA scan would benefit the clinical workflow by shortening the diagnostic time.

In the last decade, many image-based approaches have been proposed for automated classification of the presence/absence of an LVO in CTA scans. S. A. Amukotuwa et al. (2019) developed an image processing pipeline for this task by applying a rule-based classifier. Their approach achieved an area under the receiver operator characteristic curve (AUC) between 0.86 and 0.94, depending on the patient cohort. Stib et al. (2020) investigated LVO detection on multi-phase CTA, for which they trained 2D-DenseNets. Several experiments were performed by using combinations of three CTA phases, yielding AUCs between 0.74 and 0.85. Thamm et al. (2022b) applied 3D-DenseNets on segmented CTA scans and proposed the randomized recombination of the patient's hemispheres as a novel augmentation method to create new synthetic patients. In a 5-fold cross validation, they achieve an AUC of 0.91.

However, image-based approaches can be computationally expensive, and a more efficient representation is desirable. Graphs offer an efficient representation for images of sparse objects, where the pixels of interest are considered nodes, and the edges between nodes represent the pixel neighborhoods. In some cases, the imaged object itself also exhibits a graph-like structure. For instance, the cerebrovascular system can be considered a graph as well, with vessel bifurcations (branching points) as nodes and the vessel segments as edges.

Wolterink et al. (2019) applied graph neural networks for coronary artery segmentation in cardiac CTA by identifying the spatial location of nodes in a tube-shaped surface mesh. Ye et al. (2019) categorized strokes using graph attention networks by assigning each node to a stroke and forming temporal and spatial interactions between adjacent nodes as edges. Popp et al. (2022) utilized graph neural networks for thrombus detection in non-contrast head CT (NCCT), arranging multiple candidate regions of interest per patient in a graph and jointly classifying them.

This paper aims to investigate the benefits of graph deep learning (GDL) compared to conventional image-based deep learning methods in classifying the presence/absence of an LVO in cerebrovascular trees treated as graphs. We hypothesize that by using graphs, it is possible to achieve comparable performance more efficiently due to the reduction in input dimensionality. Classification is performed by applying GDL to vascular tree structures represented as graphs after a preceding segmentation in CTA data. Node features are used to describe several vascular properties. A task-specific augmentation technique, namely recombination of hemispheres from different cerebrovascular trees as proposed by Thamm et al. (2022b), is adapted to the graph-based approach to increase the number of training samples.

2. Method

2.1. Data

In total, 171 CTA scans of the head region with a cohort of AIS patients were available. Of these, 20 had to be excluded due to insufficient image quality preventing a proper segmentation of the vessels and/or registration of the head. Of the remaining 151 scans, 57% were LVO positive with an occlusion either in the MCA or the ICA region. In order to compute the graph of each vessel tree, the vasculature is first segmented using VirtualDSA++ (Thamm et al., 2020, 2022a) followed by skeletonization that determines both the vessel centerlines, resulting in a graph representation of the vasculature, and the vessel radii. The overall data is limited to the quality of the segmentation. However, the baseline method is based on the same segmentation approach. Hence, a fair comparison between both approaches is possible.

2.2. Graph Construction

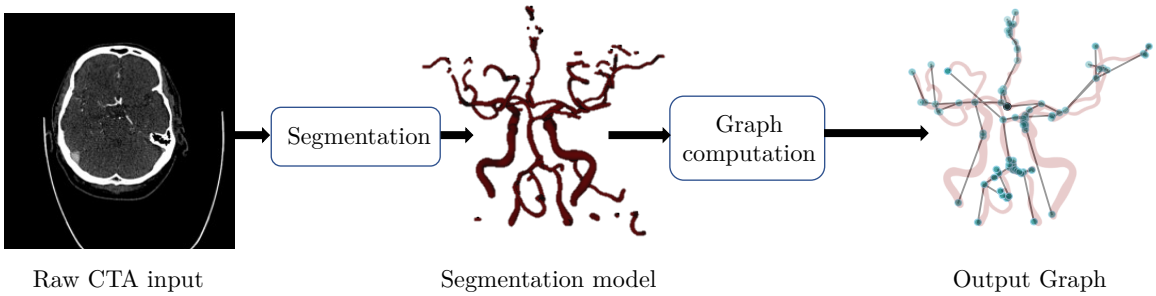

Figure 1: Pipeline for generating graph data using CTA scan as an input, where the graph's nodes are bifurcations in the vessel tree.

For each CTA data set, a graph $G = (\mathcal{V}, \mathcal{E})$ is generated following the pipeline shown in Figure 1. Each graph G consist of a set of bifurcations as nodes $v_i \in \mathcal{V}$ and set of undirected edges $\varepsilon_{ij} \in \mathcal{E}$, describing the vessel skeleton between two bifurcations v_i and v_j. Furthermore, each node v_i has a 3-dimensional feature vector containing vascular attributes about the radius (Figure 2(a)), geodesic distance to a root node on the Circle of Willis (CoW) (Figure 2(b)) and the sagittal position (Figure 2(c)).

For obtaining the node features, a landmark on the communicans anterior artery on the CoW is defined using the approach by Ghesu et al. (2017). Then a node on the CoW, referred to as the root node, is sought using the shortest Euclidean distance between nodes and the landmark. First, the average vessel radius at each bifurcation is obtained through the skeletonization process of the vascular structure by (Selle et al., 2002). Second, the geodesic distance to CoW is calculated by the shortest geodesic path between all nodes and the root node. Third, we compute each node's sagittal position (x coordinate) sorting each node to its belonging hemisphere by re-referencing the positions around the root node. The

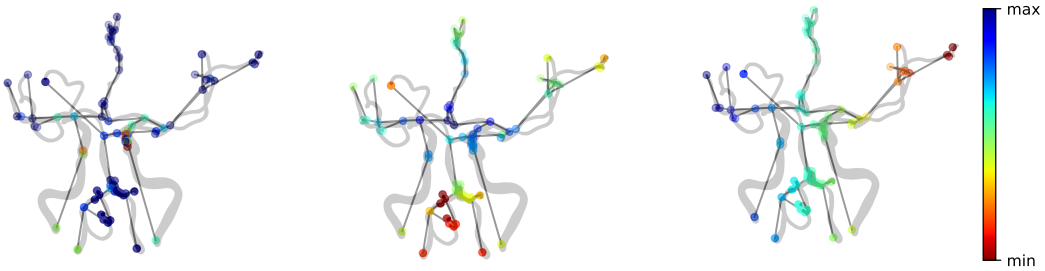

(*a*) Radius (*b*) Distance (*c*) Sagittal position

Figure 2: Heatmaps illustrating the features at each node. (a) Radius of each vessel at the corresponding node. (b) Distance of each node to the root node at the CoW. (c) Sagittal position of each node.

case-level labels are one-hot encoded 3-vectors for the three classes, LVO negative, LVO left and LVO right.

2.3. Augmentation

We apply two augmentation techniques. One operates on feature-level, the other on graph-level. For the feature-level augmentation, we jitter the value of our features by applying uniformly distributed noise with a variance of $\sigma = 10\%$.

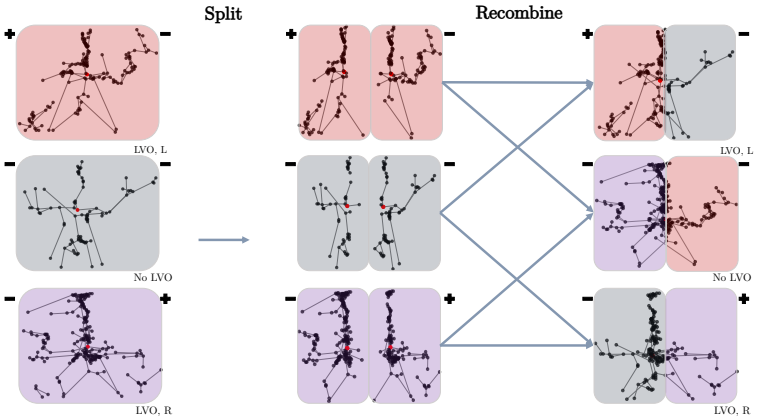

Figure 3: An example of the recombination process, with (+/− denoting LVO-pos/LVO-neg cases respectively). Three patients color-coded (red, grey and violet) are split at the root node sagittally, then recombining two hemispheres from different patients creating three new artificial patients with three different classes.

We adopt the recombination method by Thamm et al. (2022b) for the graph-level augmentation. To this end, we generate new artificial data (Figure 3) by splitting the brain

along the mid-plane into the two hemispheres, honoring the quasi-symmetrical property of the human brain and, therefore, the cerebral vasculature. Adapting this concept to graph structures, the graph is split during the generation process at the root node with a variance of $\sigma = 2\,\text{mm}$ around the sagittal axis. Each hemisphere graph is then labeled regarding the presence of an LVO along with the side of each hemisphere. During the recombination process, the overall label $y\prime_{i,j}$ is decided based on the labels of both hemisphere graphs y_i, y_j:

$$y\prime_{i,j} = \begin{cases} y_i & \text{if } y_i \neq 0 \ \wedge \ y_j = 0, \\ n.a. & \text{if } y_i = y_j \neq 0, \\ y_j & \text{else.} \end{cases} \tag{1}$$

Optionally, one hemisphere can be mirrored around the midpoint in the sagittal axis, in our case $x = 90\,\text{mm}$, of the nodes' positions depending on the label and side of each hemisphere, where the mirroring is described as

$$x\prime_{H_i} = 90 + (90 - x_{H_i}) \quad \text{if } x_{H_i} \neq \text{root node } \forall x_{H_i} \in \text{hemisphere } i. \tag{2}$$

In total, 81k recombinations are possible using the 151 patients. During training, new artificial patients can be drawn randomly with a uniform distribution of the labels to ensure the class balance, where each combination appears once in one epoch.

2.4. Architectures

We utilize GraphSAGE by Hamilton et al. (2017) as the graph convolutional layer of choice in the proposed graph networks. The first variant applies a node-level MaxPool layer after randomly clustering the nodes using the greedy clustering algorithm by Fagginger Auer and Bisseling (2012). The second approach relies on the ASAPooling layer by Ranjan et al. (2019). The last variant utilizes both pooling approaches by first applying the MaxPool followed by the ASAPooling layer later on in the network. All three architectures share the same basic structure (Figure 4).

The basic structure consists of two blocks of three GraphSAGE layers in combination with the ReLU activation function. Inspired by Corso et al. (2020), each block employs three different aggregation methods sequentially (max, sum, mean), introducing non-linearity in the message passing. The suggested combination was based on an empirical evaluation of different network combinations. However, other possible network topologies were only outperformed by slight margins.

The output of each layer is normalized, generalizing the node representation between graphs. The network starts with a width of 1024, capturing as much information as possible throughout each graph, then decreasing by half after each graph convolutional layer. Then, a vector representation is returned by a graph readout network consisting of a dropout layer followed by a global max pooling operation, which returns a batch-wise graph-level output by the channel-wise maximum across the node dimension. Finally, the representation vector is fed into the output prediction layer, consisting of a linear layer with three nodes at the end for the three classes.

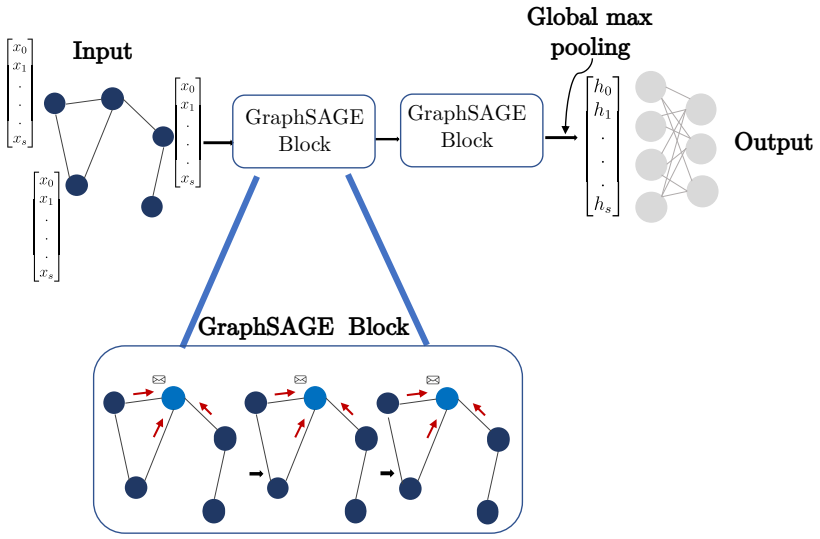

Figure 4: The basic structure of the proposed GNN consisting of two GraphSAGE blocks to update nodes representation followed by a dropout and global maxpool layer acting as a readout layer and a FC layer as a prediction network. Each block is composed out of three stacked GraphSAGE layers using thee different aggregation method in the sequence max, sum, mean.

2.5. Experiments

In this study, we use Adam by (Diederik P. Kingma and Jimmy Ba, 2015) with a learning rate of 10^{-5} and a weight decay of 10^{-4} alongside the BCE-loss function implemented in Pytorch (Paszke et al., 2019) with Python 3.8. We apply a batch size of 32 to compute a stable loss robustly. We utilize both augmentation techniques mentioned in section 2.3 separately and in combination with each method. For the recombination, we randomly draw a total of 20×32 patients from each training fold during each epoch. To evaluate our models and the impact of the recombination method, we apply a 5-fold micro-average cross-validation using a 3-1-1 split for training, validation and testing. Validation/testing is done on original, not recombined data. Additionally, early stopping is utilized with a patience of 120 epochs monitoring the validation loss.

As a baseline for comparison, we consider the approach by Thamm et al. (2022b) that utilizes 3D-DenseNets while also applying similar augmentation methods. They proposed two approaches utilizing the segmentations of the cerebrovascular system as a surrogate to the CTA scan and the recombination method to classify the existence of LVOs. In the first approach, H-Stack, the hemispheres are concatenated channel-wise. In the IM-Stack approach, the hemispheres are cropped to regions covering the ICA and MCA branches, respectively, left and right, to integrate more prior knowledge. Additional labels for the respective ICA and MCA sub-volumes are used to train 3D-DenseNets for the LVO classification task.

Table 1: Overview of all best performing models for augmentation the AUC of the respective class-wise prediction for the presence/absence of an LVO, and the accuracy for the affected side on LVO-pos. cases. The abbreviation "REC" stands for the recombination method, "D" for the deformation and "M" for mirroring.

Model		AUC	ACC
H-Stack, Thamm et al. (2022b)	H Stack	0.73	0.91
	H-Stack + D	0.82	0.86
	H-Stack + REC	0.87	0.93
	H-Stack + REC + D	0.89	0.92
IM-Stack, Thamm et al. (2022b)	IM Stack	0.84	0.92
	IM-Stack + D	0.86	**0.96**
	IM-Stack + REC	0.88	0.94
	IM-Stack + REC + D	**0.91**	**0.96**
MaxPool	original	0.68	0.76
	original + jitter	0.71	0.76
	REC	0.85	0.90
	REC + jitter	0.83	0.91
ASAPooling	original	0.83	0.64
	original + jitter	0.77	0.80
	REC	0.85	0.90
	REC + jitter	0.82	0.95
MaxPool/ASAPooling	original	0.64	0.82
	original + jitter	0.72	0.89
	REC	0.87	0,89
	REC + jitter	0.89	0.90

To evaluate the models' efficiency, we measure the time and GPU memory consumption per epoch for training and testing. For the measurement, we apply a batch size of 6 and the number of patients/epoch of 180, similar to the baseline. Therefore, both approaches process the same amount of patients in each epoch, ensuring a fair comparison.

3. Results

Table. 1 shows the AUC values achieved for the the class LVO-neg (no LVO) against LVO-pos (sum of LVO-left and LVO-right). The column ACC denotes the side accuracy by measuring the accuracy of taking the argmax of the left or right class prediction on the LVO-pos cases, thus evaluating the model's tendency between the two hemispheres.

The results show that the overall performance of the three proposed graph networks is boosted when the recombination method is introduced. When both augmentation methods are applied, each graph network achieves its best performance, where the Max/ASAPooling

model delivers the best AUC of 0.89, which is similar to H-Stack by Thamm et al. (2022b), and the ASAPoolng approach the best ACC of 0.95, outperforming the H-Stack model. Notably, only one graph network, the ASAPooling approach, performs well without any augmentation when predicting the presence/absence of an LVO, and the Max/ASAPooling approach performs well regarding the accuracy of the affected side. However, the IM-stack approach has the best overall performance when applying the subvolume recombination method.

Efficiency-wise, the proposed GDL models ($\sim 1.4M$ parameters) complete a training epoch in 2 seconds using $1\,GB$ of memory and testing/inference in 1.5 seconds allocating $0.8\,GB$ of memory. On the other hand, the IM-Stack ($\sim 6M$ parameters) approach completes in 22/5 seconds consuming $7.4/2.9GB$ of memory. Furthermore, utilizing a similar number of parameters to our models, the IM-Stack finishes in 16/4 seconds consuming $5.5/2.4GB$ of memory.

4. Conclusion

We investigated the use of GDL methods for classifying the presence of an LVO in a cerebral vessel tree. The proposed graph networks achieve comparable performance to a state-of-the-art image-based method Thamm et al. (2022b). While the GDL models do not outperform this baseline, they train up to ten times faster while allocating only a fifth of the memory. This demonstrates that representing cerebrovascular vasculature as graphs dramatically reduces the amount of redundant model input data while maintaining crucial structural information. We also adapted a previously described method for augmentation by recombination of single-hemisphere vessel trees to the proposed graph-based models and could show that it substantially benefits the learning process.

References

Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal neighbourhood aggregation for graph nets. In *Advances in Neural Information Processing Systems*, pages 13260–13271. Curran Associates, Inc., 2020.

Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *ICLR*, pages 1–15, 2015.

Salwa ElTawil and W Mui Keith. Thrombolysis and thrombectomy for acute ischaemic stroke. *Clinical Medicine*, pages 161–165, 2017.

Bas O. Fagginger Auer and Rob H. Bisseling. A GPU Algorithm for Greedy Graph Matching. In *Facing the Multicore - Challenge II*, pages 108–119. Springer Berlin Heidelberg, 2012.

Florin C. Ghesu, Bogdan Georgescu, Sasa Grbic, Andreas K. Maier, Joachim Hornegger, and Dorin Comaniciu. Robust Multi-scale Anatomical Landmark Detection in Incomplete 3D-CT Data. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2017*, pages 194–202. Springer International Publishing, 2017.

William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive Representation Learning on Large Graphs. *Conference on Neural Information Processing Systems*, 2017.

Nikita Lakomkin, Mandip Dhamoon, Kirsten Carroll, Inder Paul Singh, Stanley Tuhrim, Joyce Lee, Johanna T. Fifi, and J. Mocco. Prevalence of large vessel occlusion in patients presenting with acute ischemic stroke: a 10-year systematic review of the literature. *Journal of neurointerventional surgery*, pages 241–245, 2019.

Konark Malhotra, Jeffrey Gornbein, and Jeffrey L. Saver. Ischemic Strokes Due to Large-Vessel Occlusions Contribute Disproportionately to Stroke-Related Dependence and Death: A Review. *Frontiers in neurology*, pages 1–5, 2017.

Stephen Jx Murphy and David J. Werring. Stroke: causes and clinical features. *Medicine (Abingdon, England : UK ed.)*, pages 561–566, 2020.

Murugan Palaniswami and Bernard Yan. Mechanical Thrombectomy Is Now the Gold Standard for Acute Ischemic Stroke: Implications for Routine Clinical Practice. *Interventional neurology*, pages 18–29, 2015.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, and et al. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2019.

Antonia Popp, Oliver Taubmann, Florian Thamm, Hendrik Ditt, Andreas Maier, and Katharina Breininger. Thrombus Detection in Non-contrast Head CT Using Graph Deep Learning. In *Bildverarbeitung für die Medizin 2022*, pages 153–158. Springer Fachmedien Wiesbaden, 2022.

Ekagra Ranjan, Soumya Sanyal, and Partha Pratim Talukdar. ASAP: Adaptive Structure Aware Pooling for Learning Hierarchical Graph Representations. *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5470–5478, 2019.

S. A. Amukotuwa, Matus Straka, M. Smith, H. Chandra, R.V. Dehkharghani, S., Fischbein, and N.J. Bammer R. Automated Detection of Intracranial Large Vessel Occlusions on Computed Tomography Angiography. *American Heart Association, Inc.*, pages 2790–2799, 2019.

Dirk Selle, Bernhard Preim, Andrea Schenk, and Heinz-Otto Peitgen. Analysis of vasculature for liver surgical planning. *IEEE transactions on medical imaging*, pages 1344–1357, 2002.

Matthew T. Stib, Justin Vasquez, M.P. Dong, Y.H. Kim, S.S. Subzwari, H.J. Triedman, A. Wang, H.L.C. Wang, A.D. Yao, M. Jayaraman, and et al. Detecting Large Vessel Occlusion at Multiphase CT Angiography by Using a Deep Convolutional Neural Network. *Radiology*, pages 640–649, 2020.

Florian Thamm, Markus Jürgens, Hendrik Ditt, and Andreas Maier. VirtualDSA++: Automated Segmentation, Vessel Labeling, Occlusion Detection and Graph Search on CT-Angiography Data. pages 151–155. The Eurographics Association, 2020.

Florian Thamm, Markus Juergens, Oliver Taubmann, Aleksandra Thamm, Leonhard Rist, Hendrik Ditt, and Andreas Maier. VirtualDSA++: Automated Segmentation, Vessel Labeling, Occlusion Detection and Graph Search on CT-Angiography Data. pages 1–12, 2022a. doi: 10.1088/2057-1976/ac9415.

Florian Thamm, Oliver Taubmann, Markus Jürgens, Aleksandra Thamm, Felix Denzinger, Leonhard Rist, Hendrik Ditt, and Andreas Maier. Building brains: Subvolume recombination for data augmentation in large vessel occlusion detection. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, pages 634–643. Springer Nature Switzerland, 2022b. ISBN 978-3-031-16437-8.

Jelmer M. Wolterink, Tim Leiner, and Ivana Išgum. Graph convolutional networks for coronary artery segmentation in cardiac ct angiography. In *Graph Learning in Medical Imaging - 1st International Workshop, GLMI 2019, held in Conjunction with MICCAI 2019, Proceedings*, pages 62–69. Springer, 2019.

Jun-Yu Ye, Yan-Ming Zhang, Qing Yang, and Cheng-Lin Liu. Contextual stroke classification in online handwritten documents with graph attention networks. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 993–998, 2019.

Appendix A. Detailed results

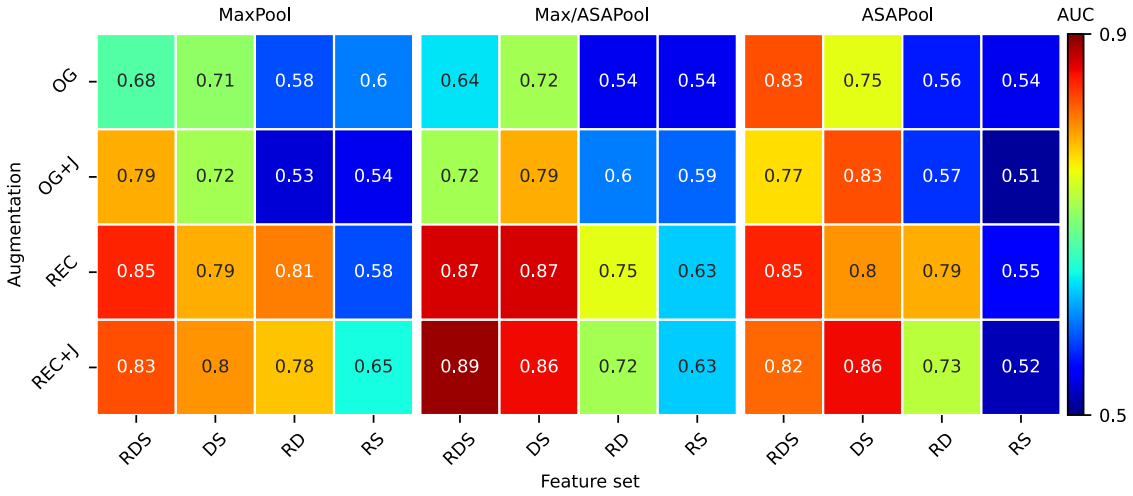

(*a*) AUC heatmap

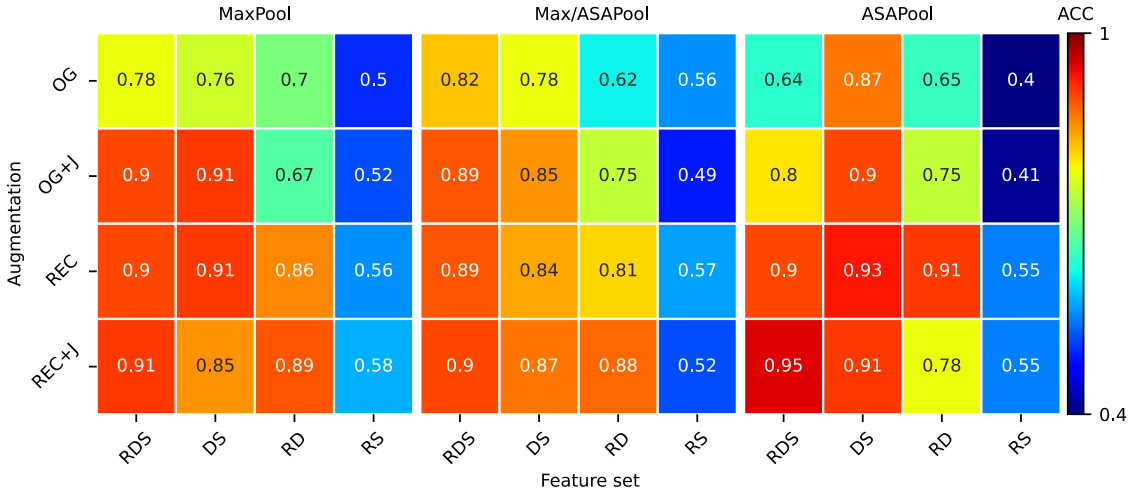

(*b*) ACC heatmap

Figure 5: Heatmaps showing the performance regarding (a) AUC and (b) ACC of all proposed models using different combinations of features and augmentation methods with the following legend; R: radius, D: distance, S: side, OG: original data, J: jittering, REC: recombination.