Automatic Cerebral Perfusion Imaging using Deep Learning for Digital Subtraction Angiography

Flavius-Gabriel Marc (MSc Student), Ruisheng Su (Daily Supervisor), Matthijs van der Sluijs (Daily Supervisor), prof. dr. Josien Pluim (Examiner) and dr. ir. Theo van Walsum (Supervisor host institute)

Minor Research Project, September 2023
Erasmus MC, Rotterdam, The Netherlands

ARTICLE INFO

Keywords:
deep learning
computer vision
internal carotid artery segmentation
deconvolution
X-ray
DSA
perfusion

ABSTRACT

Visualization of blood flow in brain vessels is crucial for neurovascular disease patients, including the evaluation of ischemic stroke treatments. X-ray Digital Subtraction Angiography (DSA) is the standard imaging modality used for this purpose. So far, visual inspection is the primary way to assess DSA series. Due to the high temporal resolution of DSA, it holds great potential in facilitating quantitative assessment of cerebral hemodynamics. Various parametric perfusion images have been generated from DSA based on temporal blood flow characteristics. This technique is commonly referred to as perfusion DSA. Such parameters include cerebral blood volume (CBV), cerebral blood flow (CBF), time to maximum (Tmax), and mean transit time (MTT) and can be generated using deconvolution techniques. To obtain these deconvolution-based images, an arterial input function (AIF) is extracted from the internal carotid artery (ICA) region of interest. However, current perfusion DSA methods require manual annotation of the ICA. In this work, a supervised deep-learning model for semantic ICA segmentation was trained. Subsequently, an automated application was developed to generate perfusion DSA images. The results of quantitative statistics indicate that there is no significant difference between using the automated application that utilizes the AI segmentation model and the manual annotation method. These findings suggest that this application framework may offer important benefits in clinical practice and future research.

1. Introduction

Cerebrovascular diseases, such as ischemic stroke caused by blood vessel blockage, stenosis or aneurysm, are major contributors to global mortality and long-term disability (Roth et al., 2020). Ischemic stroke is the most prevalent type of stroke and develops when there is an obstruction in an artery of the brain. This occurs when a blood clot blocks blood flow to the brain, causing tissue damage and neurological issues (Venema et al., 2017). Endovascular thrombectomy (EVT) improves outcomes in patients with acute ischemic stroke by restoring blood flow to the affected area of the brain (Goyal et al., 2016). This involves injecting a contrast agent into the brain vessels to enable accurate imaging and identification of the blockage location (Berkhemer et al., 2015).

Digital Subtraction Angiography (DSA) is the primary imaging modality used for evaluation during endovascular treatment. In this minimally invasive treatment, a catheter is guided towards the occlusion after which the occlusion is removed (Su et al., 2023). DSA involves capturing a series of actual X-ray images. Subsequently, DSA images are generated by taking one image before the injection of the contrast agent and then subtracting this image from a series of images taken after the contrast agent has been injected (Ovitt et al., 1979). This process helps to visualize the blood flow within the brain and exclude any unrelated background elements (i.e., skull) from the images. DSA images are captured in quick succession, enabling the visualization of the contrast agent's movement through the blood vessels. By tracking the intensity values, across the series of images, time-intensity curves (TICs) can be generated. This is used to provide information about the temporal changes in contrast agent concentration within the DSA sequence. The arterial input function (AIF) refers to the average TIC of a single artery that contains contrast intensity values. Parameters such as the AUC (area under the curve), peak concentration, and TTP (time-to-peak) can be automatically obtained from TICs. On the other hand, to provide quantitative measurements of other hemodynamic parameters, AIF can be used to compute additional perfusion parameters.

Perfusion DSA is therefore a concept that involves constructing perfusion images from DSA data (Su et al., 2023). Studies have shown that parametric images can provide quantitative information about cerebral hemodynamics and can aid the visual assessment of treatment (Strother et al., 2010; Scalzo et al., 2016). In this work, the internal carotid artery (ICA) is selected as the input artery in order to extract the AIF curve, as it is a vessel used frequently for contrast injection. Subsequently, by employing deconvolution techniques, perfusion parameters such as cerebral blood flow (CBF), cerebral blood volume (CBV), mean transit time (MTT), and time to maximum (Tmax) are computed for visual interpretation and quantitative analysis.

Current perfusion DSA methods require manual annotation of the ICA to extract the AIF (Scalzo et al., 2016). This is time-consuming. In this work, a supervised deep learning model is developed for the automatic semantic segmentation of ICA, from which the AIF is extracted. Subsequently, deconvolution-based perfusion DSA images are obtained. The outcome is an automated framework that takes a DICOM DSA series as input and generates perfusion

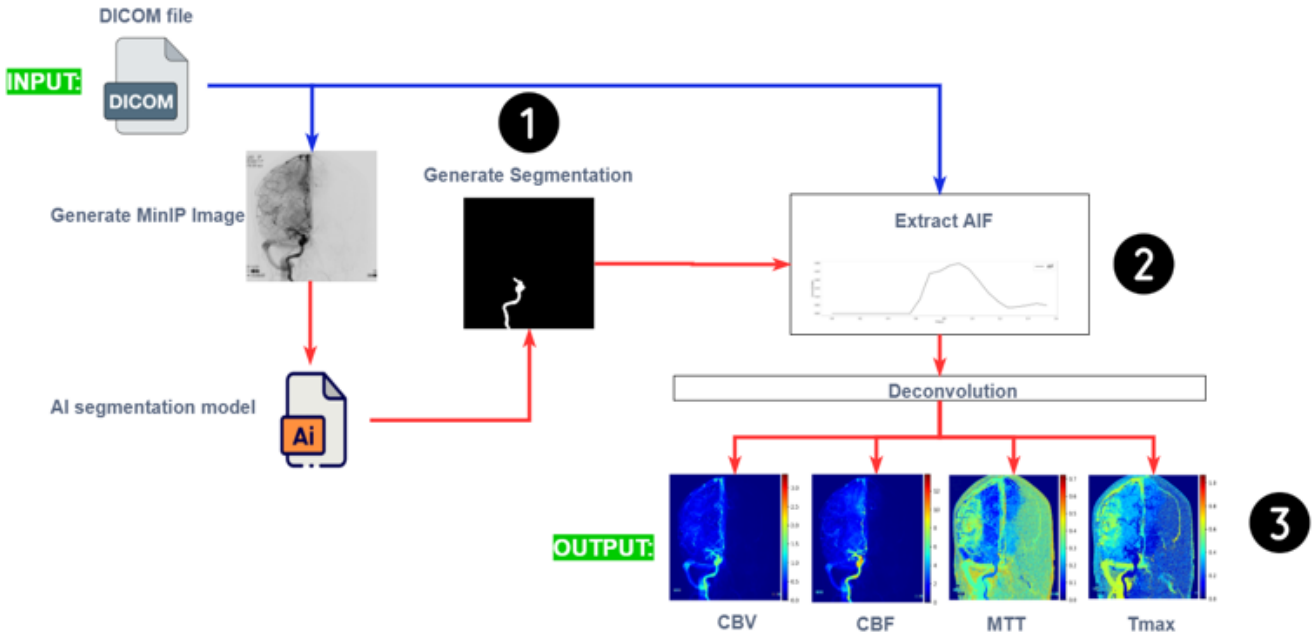

Figure 1: Diagram illustrating the execution steps: 1. Generating the segmentation using MinIP and AI segmentation model file; 2. Extracting the AIF; 3. Using deconvolution algorithms, compute perfusion parameter maps.

images. The framework is evaluated using a national multi-center DSA image dataset. Furthermore, experiments are carried out to improve the segmentation model and generate comparative statistics. These statistics are used to investigate the performance of the automated segmentation model in comparison to manual annotation, as well as to identify future enhancements for the application framework. To the best of our knowledge, this is the first automated end-to-end framework for perfusion DSA generation. Therefore, a comparison with another method is not available in the current literature.

The subsequent sections of this paper are structured as follows: Firstly, the methods used in this work are described in Section 2. Then, in Section 3, the experiments conducted, the dataset used, the annotation process, and the evaluation metrics are presented. The results obtained from the experiments are presented in Section 4, followed by a discussion in Section 5. Finally, the paper concludes with a summary of the main findings in Section 6.

2. Methods

In this work, an end-to-end framework application was developed. The diagram in Figure 1 illustrates the overall process of the automated application, including the main operations involved in generating the perfusion parameter maps. In order to generate the segmentation, it is necessary to possess both the DSA DICOM files and an AI segmentation model (weights) file. After generating a MinIP image, the AI segmentation model generates the ICA segmentation, which is then combined with the DICOM series to extract the AIF. From one DICOM series of DSA images, a MinIP image is generated. MinIP stands for Minimum Intensity Projection, which is a technique used to create an image that displays the minimum pixel intensity value from a stack of images. The pixel values of MinIP are the minimum pixel values along the time direction of the DSA image series. Having the extracted AIF values from ICA, using deconvolution-based algorithms, the color-coded parameter images are obtained. This application allows for the generation of parameter maps for multiple patients (multiple DICOM files as input).

The subsections include Section 2.1, which describes the AI segmentation model, and Section 2.2, which details the process of obtaining perfusion DSA images.

2.1. AI segmentation model

The segmentation model was trained using a standard deep learning network, U-Net (Ronneberger et al., 2015). The network (Figure 2), takes DSA MinIP Images as input. Rather than using DSA frame images, MinIP is used to enhance the visibility of blood vessels. Hence, MinIP images are employed for spatial encoding and decoding in a U-shaped design, comprising five downward layers (yellow boxes) and five upward layers (green boxes). The architecture of each layer is made up of double-convolution blocks that apply instance normalization and ReLU activation. The contracting path uses max pooling while the expanding path uses bilinear upsampling. The number of feature channels starts at 64, doubles after each max pooling operation, and halves after each bilinear upsampling operation. The loss function is defined as a combination of cross-entropy loss and a Dice loss (Kervadec and de Bruijne, 2023).

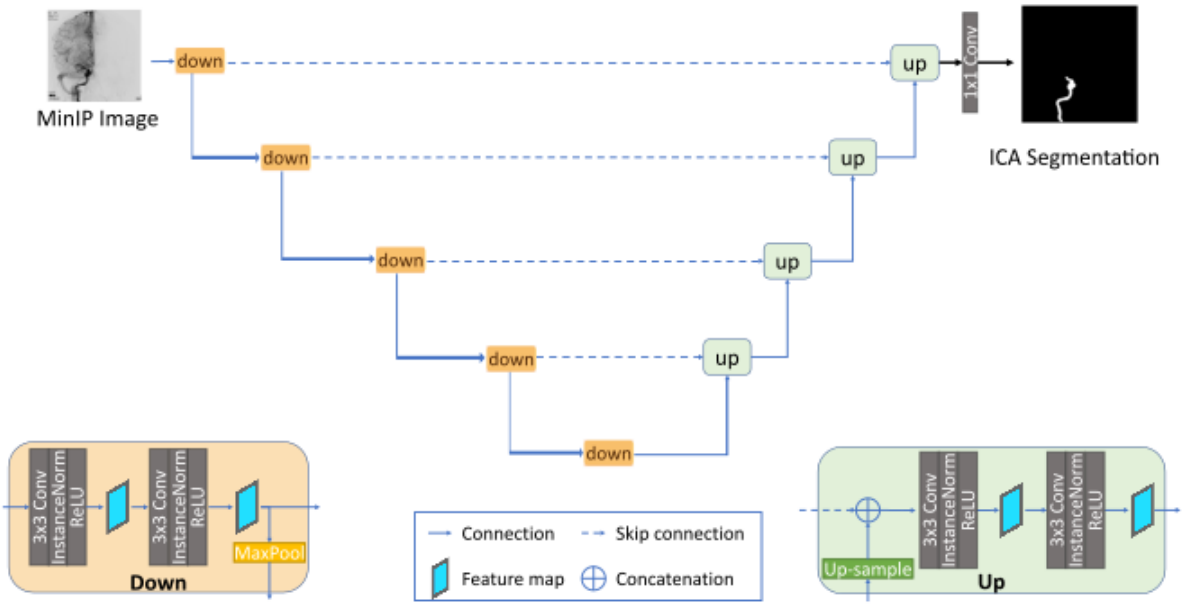

Figure 2: Network architecture for ICA segmentation in DSA: Input is a MinIP image and the output is a two-channel segmentation image with black representing the background and white representing the mask. Adapted from Su et al. (2022).

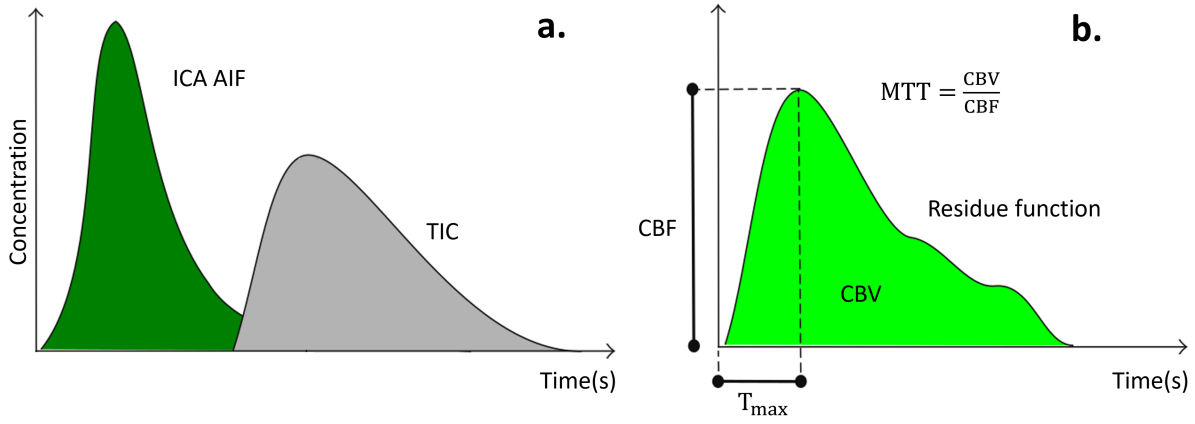

Figure 3: a. Illustration of a time-intensity curve (TIC) (gray) with respect to the arterial input function extracted from ICA (ICA AIF) (dark green). The deconvolution of the TIC curve with ICA AIF produces the residue function (b.). CBF is extracted at the maximum value reached at T_{max}, while MTT is calculated as CBV/CBF, where CBV is determined as the AUC (light green).

2.2. Computing deconvolution-based images

Commonly used perfusion DSA parameters, such as CBF, CBV, MTT, and Tmax, offer valuable insights into microvascular hemodynamics in brain tissue. Therefore, perfusion parameters can be used to assess the severity of ischemia (Gaillard, 2023).

CBF, short for Cerebral Blood Flow, measures the volume of blood that passes through a specific brain region per unit of time. CBF is crucial because it provides valuable insights into brain health and function (ScienceDirect, 2023). However, CBF in a specific region of brain tissue cannot be estimated based on the contrast agent concentration curve alone. Instead, the concentration curve should be considered in conjunction with the time-dependent curve of the contrast agent's entry or exit from the tissue, known as the AIF (Liebeskind et al., 2019).

CBV, or Cerebral Blood Volume, complements CBF by providing information about the amount of blood present in a specific area of the brain. It is essentially a measure of blood storage within the brain's vascular system. Understanding CBV helps in diagnosing conditions like brain tumors and vascular malformations (Wintermark et al., 2005).

MTT, which stands for Mean Transit Time, is a parameter that quantifies the average time taken by blood to pass through a given region. It is an important metric when assessing the efficiency of blood flow within a tissue. Deviations in MTT can indicate problems such as ischemia or impaired blood circulation.

Tmax, or Time to Maximum Enhancement, focuses on pinpointing when a particular tissue or region receives its maximum blood supply. This parameter is particularly useful in identifying areas of delayed perfusion, which could signify vascular obstructions or other medical concerns (Wintermark et al., 2005).

An illustration of how these parameters are computed is shown in Figure 3. You may refer to Fieselmann et al. (2011) and Brix et al. (2010) for mathematical derivations.

3. Experiments

3.1. Dataset

The data used is from the MR CLEAN Registry (Jansen et al., 2018), a multi-center (16 centers) observational cohort study of acute ischemic stroke patients in the Netherlands between March 2014 and December 2018. The DSA frame images are 1024 x 1024 pixels, acquired on various systems (Philips, GE, Siemens) at 0.5-4 frames per second over 10-50 frames.

The ICA segmentations were manually created in MeVisLab (Heckel et al., 2009) using an in-house tool that uses the region-growing technique. The annotation was done for 150 random patients including both AP and lateral views. An example of how the masks were created is shown in Figure 4. A radiologist supervised the creation of these masks, and they are used as a reference in this work.

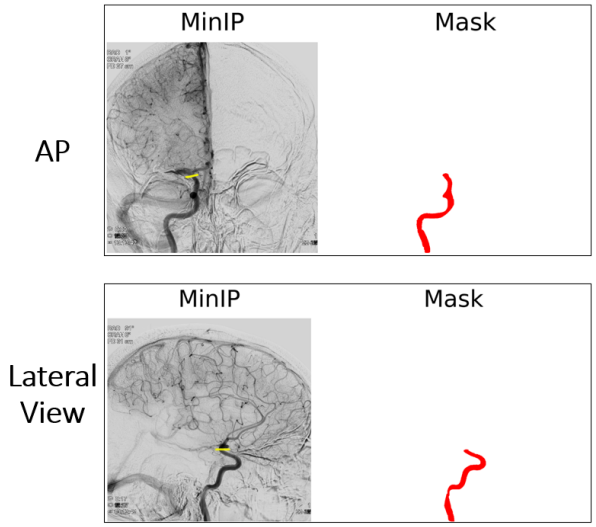

Figure 4: The figure shows masks created for one patient for the anterior-posterior (AP) and lateral view. Segmentation was performed on DSA MinIP images. The top boundary of the mask is highlighted in yellow on the MinIP image.

3.2. Implementation

The automatic application is implemented in Python and the segmentation model was implemented in PyTorch (Steiner et al., 2019).

As a pre-processing step for both, the automatic application and the training of the segmentation model, all DSA frames are left to their original size (1024 × 1024 pixels) and the intensity values are normalized to [0,255]. For the automatic application, there is additional pre-processing made. The temporal resolution of the DSA frames is linearly resampled to 1 fps to enable uniform analysis. Furthermore, the code checks whether there is a frame time vector in the DICOM header and if there is not, the patient is discarded from the processing.

The segmentation model was trained on an NVIDIA RTX A40 GPU. The dataset used is randomly split into training (50%), validation (30%) and testing (20%) on the patient level. Both orientations, AP and lateral views are included in the dataset and randomly distributed into the splits. The used optimizer is RMSpropo (Hinton et al., 2012) with 0.9 momentum and weight decay set to 10^{-8}. The chosen scheduler is ReduceLROnPlateau (Steiner et al., 2019), with patience set to 10 epochs, 0.5 decay factor, and 10^{-8} minimum learning rate. Early stopping is also implemented which configures how long to wait (epochs) after the last time validation loss improved. It monitors the validation loss and stops training if the validation loss stops increasing for 50 epochs. The best-performing model (see Section 3.4) weights file is included in the automatic application for predicting the masks.

3.3. Evaluation metrics

The best-performing ICA segmentation model is selected by utilizing the Dice coefficient (Dice score), comparing the predicted masks generated by the model with the reference annotated masks.

Comparative statistics are conducted to quantify and assess the differences between using generated masks and manually segmented masks for the computation of the perfusion parameter maps. The Root Mean Square Error (RMSE) between the extracted AIFs is calculated. Additionally, the absolute differences between the pixel values of the perfusion parameter maps are calculated, along with the absolute percentage difference for each pixel value in a reference parameter map (LLC, 2023). The percentage difference is calculated as follows:

$$\frac{V_2 - V_1}{|V_1|} \times 100 \tag{1}$$

where V_2 represents the pixel value generated from processing the generated masks, and V_1 represents the pixel value obtained from processing the reference masks. Finally, differences in perfusion parameters are tested for statistical significance using the paired t-test.

3.4. Experiment 1 - Data Augmentation

In the initial experiment, data augmentation was introduced to the training set. This included applying horizontal flips and "ShiftScaleRotate" transformations randomly during training, each with a 0.5 probability. The rotation was confined to a range of -10° to 10°, while the shift and scale remained at their default settings of -6.25% to 6.25% and -10% to 10%, respectively (Albumentations, 2023). The

Table 1
Data Split experiment details.

Model	Reduced Training Data Percentage	Training Set Size (patients)	Init Learning Rates
Best-performing	N/A	75	N/A
1	10%	67	$10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}$
2	20%	60	
3	30%	53	
4	40%	45	
5	50%	37	
6	60%	30	
7	70%	23	

rotation parameter was adjusted to accommodate the variability observed in DSA images, but care was taken not to over-rotate but to maintain realism in the training data. The initial learning rate was set to 10^{-4} for this experiment. The goal of this experiment was to establish the best-performing model for the AI segmentation model, using the validation Dice score. Furthermore, the overall performance of the test dataset is reported, taking into account the best-performing model.

3.5. Experiment 2 - Data Split Analysis

In a second experiment, a Data Split Analysis (see Table 1) was performed to assess the need for acquiring more data by investigating the impact of reducing the training data. In this experiment, a certain percentage of data is randomly subtracted from the initial set of training data images, simulating scenarios with reduced training data. While the training data was reduced for each set, the validation dataset was kept the same with the best-performing model (see Section 3.4). New models were trained using the reduced training datasets while retaining the same architecture as the best-performing model. Most of the hyperparameters were also kept the same, but the initial learning rate was also varied along with the reduced training sets. The impact of reduced training data models on performance was analyzed by visualizing them in relation to different initial learning rates. The results are reported using the validation Dice score and the standard deviation of the Dice score.

3.6. Experiment 3 - Post-processing

In some cases, the AI segmentation model can produce a low Dice score because it incorrectly predicts multiple connected regions. When visually inspecting these connected regions, it becomes clear that the model has the potential to provide a well-segmented ICA result. However, it can still yield a low Dice score due to inaccuracies in predicting additional regions. To address this, a post-processing script was implemented in this experiment to consider only the largest connected component among all generated masks. The impact of this approach was quantitatively assessed by examining the average mean Dice score for the test dataset.

3.7. Experiment 4 - Quantitative analysis

For the quantitative analysis, 63 DSA sequences from the test dataset, with a valid frame time vector in the DICOM header, were selected. ICA segmentations were generated by the best-performing model chosen in Section 3.4. A quantitative approach is employed to investigate the differences in generated perfusion parameters when utilizing the generated masks in comparison to the reference masks. The comparative statistics are first made at the level of AIF curves, calculating the RMSE between these curves. Next, for the perfusion parameters, the absolute differences between the pixel values in the generated and reference parameter maps are computed, as well as the absolute percentage change for each individual pixel in the reference parameter maps (Equation 1). Percentage change maps yield extremely large values for extremely low pixel values. Thus, when reporting percentage change, pixel values lower than 0.01 are adjusted to 0. For statistics, the mean absolute difference and mean absolute percentage are reported for each DSA sequence based on Dice coefficients. Achieving a low RMSE, mean absolute difference, and mean absolute percentage change is crucial to accurately quantify the perfusion images. Finally, a paired t-test was performed to determine the statistical significance of perfusion parameters. The expectation was that there would be no significant difference between using automated segmentation masks and manual segmentation.

To visually compare the generated and reference parameter maps, you can find the absolute difference maps for three patients in the Appendix (refer to Figure 11).

4. Results

4.1. Experiment 1 - Data Augmentation

Figure 5 displays the results for the ICA segmentation predicted by the U-Net on the test dataset. The mean Dice achieved for this dataset was 0.875 with a standard deviation of 0.126.

Applying data augmentation increases efficiency. The model that applies horizontal flip and shiftscalerotate is the best-performing (Table 2).

4.2. Experiment 2 - Data Split Analysis

The results of the data split experiment are presented in Figure 6. Based on the findings, it cannot be concluded that a significant increase in performance can be achieved by adding more patients to the training dataset in steps of 10%. In relation to the initial learning rates, fluctuations in performance are observed more often when the reduction

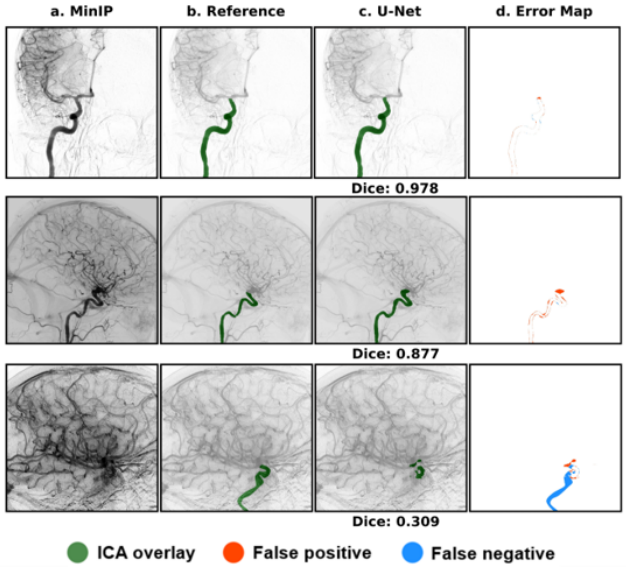

Figure 5: Figure showing ICA segmentation results of U-Net. a: the MinIP image; column b: the reference annotation of ICA; column c: segmentation output of the U-Net model; column d: the error map of U-Net showing the false positives with orange and false negatives with blue. In the top row, the maximum Dice score performance is shown. A Dice score closer to the mean is shown in the middle row, and the minimum Dice score performance is displayed in the third row.

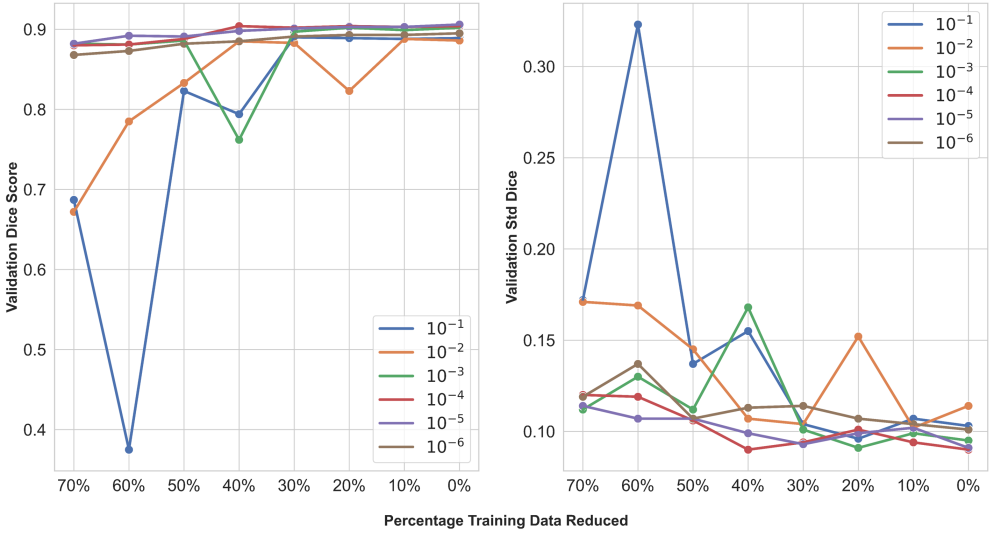

Figure 6: Data Split Experiment results. Performance is reported on the mean validation Dice score and mean validation standard deviation Dice score.

Table 2
Data Augmentation results reported on the validation Dice mean (column 2) and validation Dice standard deviation (column 3).

Model	Val Dice Mean	Val Dice Std
horizontal flip + shiftscalerotate	0.905	0.089
horizontal flip	0.891	0.104
no augmentation	0.889	0.099

in training data exceeds 30%, particularly for learning rates higher than 10^{-3}.

4.3. Experiment 3 - Post-processing

Before implementing the post-processing step, the average Dice coefficient for the test dataset was 0.875. However, after applying the post-processing step, which involves selecting the largest connected region in the generated segmentation masks, the mean Dice coefficient decreased below 0.875 due to a worst-case scenario. Figure 7 illustrates both a favourable and an unfavourable scenario. It was noted that such worst-case scenarios can occur when using this post-processing step. Nevertheless, if we exclude the specific patient from the processing, the mean Dice coefficient for the test dataset increases to 0.882 after post-processing.

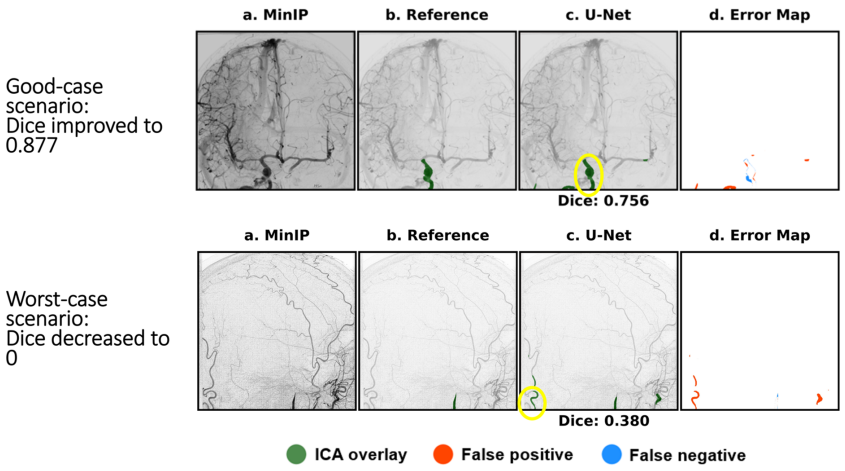

Figure 7: Good-case scenario and worst-case scenario in the post-processing experiment. With yellow is encircled the largest connected region considered after post-processing. Columns a. to d. are explained in Figure 5.

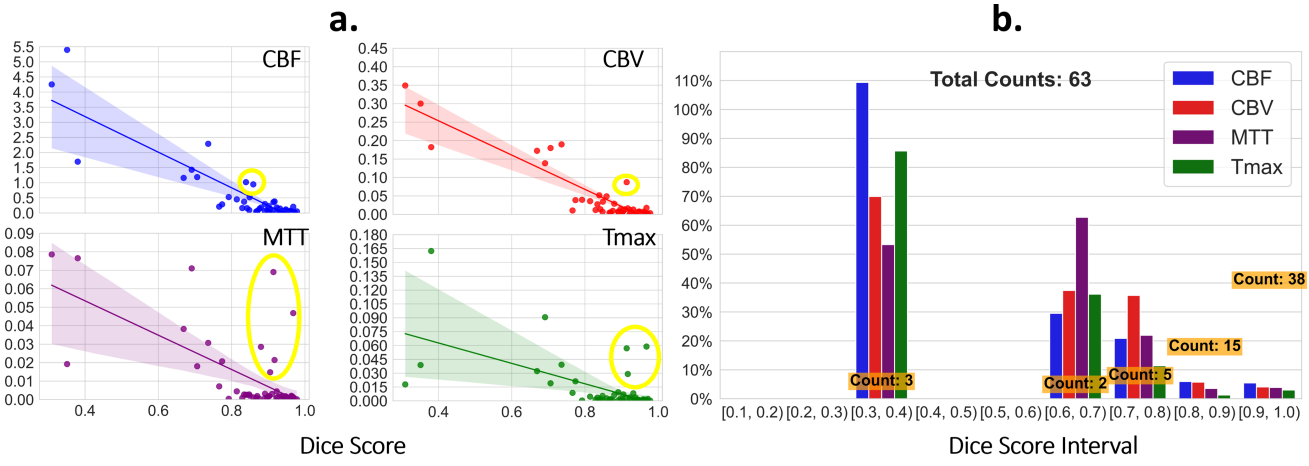

Figure 8: Figure showing differences between using generated masks and manual masks: a. The mean of absolute differences (y-axis) in perfusion parameters related to the Dice Score on the x-axis. Outliers encircled with yellow for a Dice greater than 0.8; b. The mean absolute percentage change (y-axis) for each perfusion parameter, related to Dice intervals (x-axis).

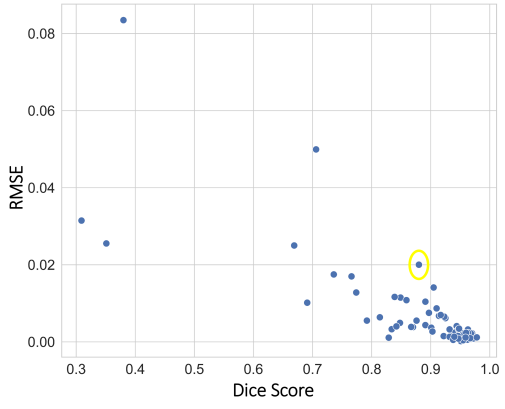

Figure 9: RMSE between the extracted AIFs related to the Dice Score. Outlier encircled with yellow.

4.4. Experiment 4 - Quantitative analysis

In Figure 8a, it can be observed that predictions with a Dice score above 0.8 exhibit notably low mean absolute differences across all perfusion parameters. This observation is reaffirmed in the barplot (8b), which also shows that the predictions with a Dice score exceeding 0.8 have mean absolute percentage changes below 10%. Furthermore, this pattern is evident in Figure 9 as well, where overall low RMSE values are obtained for a Dice higher than 0.8. However, it's worth noting some outlier values encircled with yellow in Figure 8a and Figure 9. Upon closer examination of these outliers, a specific pattern was identified. Figure 10 shows a low-high-low (LHL) intensity pattern. This suggests that even if the Dice score is higher than 0.8, the AI segmentation model can especially be improved for these examples.

Ultimately, Table 3 shows that the difference between using generated and reference masks is not significant ($p>0.05$) for the generation of all perfusion parameter images.

Table 3
Paired t-test. The analysis was performed on the 63 DSA patients considered for the quantitative analysis. The predetermined significance level is set to 0.05. AI Mean is the mean computed for each perfusion parameter obtained using the generated ICA masks. Manual Mean is the mean computed for each perfusion parameter obtained using manually created ICA annotations.

Perfusion Parameters	AI Mean	Manual Mean	Difference	T-Statistic	P-Value	Result
CBV	0.337	0.342	-0.005	-0.700	0.487	No Significant Difference
CBF	3.740	3.869	-0.129	-1.139	0.259	No Significant Difference
MTT	0.110	0.110	0.000	-0.082	0.935	No Significant Difference
Tmax	0.180	0.178	0.002	0.782	0.437	No Significant Difference

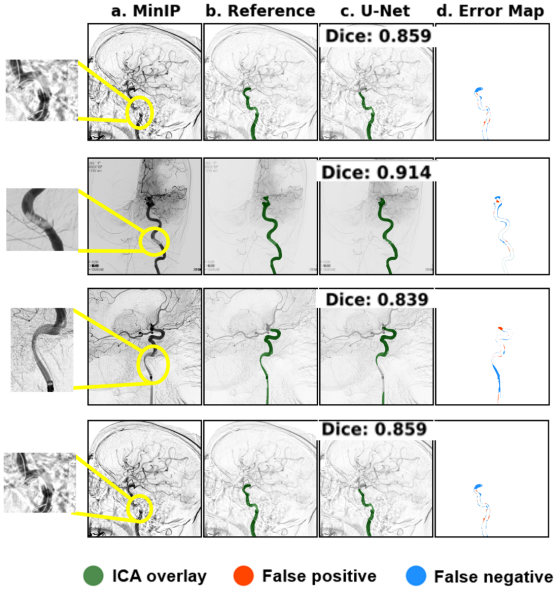

Figure 10: In the left side of the MinIP column, The Low-High-Low (LHL) intensity pattern is zoomed in for proper visualization. Columns a. to d. are explained in Figure 5.

5. Discussion

In this work, a fully automated deep learning-based method for generating perfusion parameter images from DSA images is proposed. First, an AI segmentation model for ICA segmentation is trained and next, the best-performing model weights file is included in the automated application. The AI segmentation model is capable of taking multiple MinIP images and output ICA segmentations.

The experiments performed for tuning the AI segmentation model proved that data augmentation is an added benefit. The figure showing ICA segmentation results (Figure 5) shows in the last row, the lowest segmentation Dice of being 0.309. However, this MinIP presents significant patient motion. In the work of Su et al. (2021), a motion correction algorithm was implemented with visible results. Therefore, this algorithm can be integrated as a pre-processing step for all/motion-related DSA frames before training. This could lead to enhancements in the lowest Dice prediction. The Data Split Analysis did not show that increasing the quantity of the current training dataset in steps of 10% would result in a notable better performance.

The quantitative analysis demonstrates that using the AI segmentation model for generating ICA masks does not significantly differ (paired t-test) from using manual reference which is time-consuming to create. It can be noted that the majority of randomly selected patients included in the processing with an ICA segmentation Dice score higher than 0.8 report a low mean absolute difference, mean absolute percentage change, and RMSE. This indicates that having segmentations with a Dice of at least 0.8 can be good enough for generating promising perfusion parameters. Furthermore, the quantitative analysis also suggests that further enhancement is achievable by addressing the accurate segmentation of LHL intensity patterns. These patterns were identified as outliers in both the mean absolute difference and RMSE graphs. One way to achieve this is by incorporating additional data augmentation during/before training, such as utilizing a Generative Adversarial Network (GAN) (Zhao and Bilen, 2021). This refinement is expected to result in higher Dice scores for the current segmentations that already exceed 0.8.

Applying the post-processing step of taking the largest connected component shows some potential but it has a good and worst-case scenario. The worst-case scenario would make this approach unfeasible for automated processing. It is noted that, in the good case scenario, a Dice score of 0.877 (from 0.756) can be obtained. This means that the potential is there. If we also take into consideration the results obtained from the quantitative analysis and discard from the processing all the patients with a Dice score lower than 0.7, there is a chance that patients who have an ICA segmentation score between 0.7 and 0.8 can also be considered for processing. However, a more thorough analysis is further required to completely assess this.

This work employs the foundational U-Net as the AI segmentation model's architecture. It is worth noting that the U-Net model may not be considered one of the most recent and advanced models in the current literature. Newer state-of-the-art networks, such as nnUNet (Isensee et al., 2020), UNet++ (Zhou et al., 2018), and TransUNet (Jieneng Chen et al., 2021), could potentially offer better performance than U-Net which is based on the spatial features from MinIP images.

Expanding on this research, from a clinical perspective, the AI segmentation model demonstrates its potential to adapt to different patients. Its implementation in clinical

practice is highly valuable as it reduces the laborious manual annotation of the ICA artery. This application could be integrated with a user-friendly graphical interface (GUI). Another viable approach worth considering is a semi-automatic processing method. In this approach, radiologists can use the automated tool to achieve accurate segmentation for most cases, reserving manual intervention for more complex scenarios. Although this would necessitate a more extensive GUI, it could prove highly advantageous if the goal is to implement this work in practice as soon as possible.

6. Conclusion

In conclusion, the automated deep learning-based application for generating perfusion parameter images has been found to produce results comparable to the manual-based application, which is a time-consuming process due to the need for manual segmentation of the internal carotid artery. For most patients in the random test dataset, the automated application shows minimal differences compared to its manual counterpart. However, there is place for improvement in the AI segmentation model. Additionally, pre-processing or post-processing techniques can be added to further enhance the performance of the automated application. Overall, the automated application offers important advantages that may prove valuable both in future research and in clinical practice.

References

G. A. Roth, G. A. Mensah, C. O. Johnson, G. Addolorato, E. Ammirati, L. M. Baddour, N. C. Barengo, A. Z. Beaton, et al., Global Burden of Cardiovascular Diseases and Risk Factors, 1990–2019: Update From the GBD 2019 Study, Journal of the American College of Cardiology 76 (2020) 2982–3021.

E. Venema, M. J. Mulder, B. Roozenbeek, J. P. Broderick, S. D. Yeatts, P. Khatri, O. A. Berkhemer, B. J. Emmer, Y. B. Roos, C. B. Majoie, et al., Selection of patients for intra-arterial treatment for acute ischaemic stroke: development and validation of a clinical decision tool in two randomised trials, bmj 357 (2017).

M. Goyal, B. K. Menon, W. H. Van Zwam, D. W. Dippel, P. J. Mitchell, A. M. Demchuk, A. Dávalos, C. B. Majoie, A. van Der Lugt, M. A. De Miquel, et al., Endovascular thrombectomy after large-vessel ischaemic stroke: a meta-analysis of individual patient data from five randomised trials, The Lancet 387 (2016) 1723–1731.

O. A. Berkhemer, P. S. Fransen, D. Beumer, L. A. Van Den Berg, H. F. Lingsma, A. J. Yoo, W. J. Schonewille, J. A. Vos, P. J. Nederkoorn, M. J. Wermer, et al., A randomized trial of intraarterial treatment for acute ischemic stroke, New England Journal of Medicine 372 (2015) 11–20.

R. Su, P. M. van der Sluijs, J. Bobi, A. Taha, H. M. van Beusekom, A. van der Lugt, W. J. Niessen, D. Ruijters, T. van Walsum, Towards quantitative digital subtraction perfusion angiography: An animal study, Medical Physics (2023).

T. Ovitt, M. P. Capp, P. Christenson, H. Fisher, M. Frost, S. Nudelman, H. Roehrig, G. Seeley, Development of a digital video subtraction system for intravenous angiography, in: Recent and Future Developments in Medical Imaging II, volume 206, SPIE, 1979, pp. 73–76.

C. Strother, F. Bender, Y. Deuerling-Zheng, K. Royalty, K. Pulfer, J. Baumgart, M. Zellerhoff, B. Aagaard-Kienitz, D. Niemann, M. Lindstrom, Parametric color coding of digital subtraction angiography, American Journal of Neuroradiology 31 (2010) 919–924.

F. Scalzo, D. S. Liebeskind, et al., Perfusion angiography in acute ischemic stroke, Computational and Mathematical Methods in Medicine 2016 (2016).

O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18, Springer, 2015, pp. 234–241.

H. Kervadec, M. de Bruijne, On the dice loss variants and sub-patching, in: Medical Imaging with Deep Learning, short paper track, 2023.

R. Su, M. van der Sluijs, S. Cornelissen, P. J. van Doormaal, R. v. d. Broek, W. van Zwam, J. Hofmeijer, D. Ruijters, W. Niessen, A. van der Lugt, et al., Spatio-temporal u-net for cerebral artery and vein segmentation in digital subtraction angiography, arXiv preprint arXiv:2208.02355 (2022).

F. Gaillard, CT perfusion in ischemic stroke | Radiology Reference Article | Radiopaedia.org, 2023. URL: https://radiopaedia.org/articles/ct-perfusion-in-ischaemic-stroke?lang=us. doi:10.53347/rID-24526.

T. ScienceDirect, Cerebral Blood Flow - an overview | ScienceDirect Topics, 2023. URL: https://www-sciencedirect-com.proxy.library.uu.nl/topics/agricultural-and-biological-sciences/cerebral-blood-flow.

A. Liebeskind, A. Deshpande, J. Murakami, F. Scalzo, Automatic estimation of arterial input function in digital subtraction angiography, in: International Symposium on Visual Computing, Springer, 2019, pp. 393–402.

M. Wintermark, M. Sesay, E. Barbier, K. Borbély, W. P. Dillon, J. D. Eastwood, T. C. Glenn, C. B. Grandin, S. Pedraza, J.-F. Soustiel, T. Nariai, G. Zaharchuk, J.-M. Caillé, V. Dousset, H. Yonas, Comparative Overview of Brain Perfusion Imaging Techniques, Stroke 36 (2005) e83–e99. Publisher: American Heart Association.

A. Fieselmann, M. Kowarschik, A. Ganguly, J. Hornegger, R. Fahrig, Deconvolution-based CT and MR brain perfusion measurement: theoretical model revisited and practical implementation details, International Journal of Biomedical Imaging 2011 (2011) 467563–467563. MAG ID: 2049152314 S2ID: 14b8ebec39b6738720eb391e137fb7313145d8ea.

G. Brix, J. Griebel, F. Kiessling, F. Wenz, Tracer kinetic modelling of tumour angiogenesis based on dynamic contrast-enhanced CT and MRI measurements., European Journal of Nuclear Medicine and Molecular Imaging 37 (2010) 30–51. MAG ID: 2037631380 S2ID: 9b041bdea02b5af8aee36b5a7f38eb7c50b4bab5.

I. G. Jansen, M. J. Mulder, R.-J. B. Goldhoorn, Endovascular treatment for acute ischaemic stroke in routine clinical practice: prospective, observational cohort study (mr clean registry), bmj 360 (2018).

F. Heckel, M. Schwier, H.-O. Peitgen, H.-O. Peitgen, Object-oriented application development with MeVisLab and Python, GI Jahrestagung (2009) 1338–1351. MAG ID: 2675101 S2ID: ac2e84709387035118a96e5e38923727def6d765.

B. Steiner, Z. DeVito, S. Chintala, S. Gross, A. Paszke, F. Massa, A. Lerer, T. Killeen, G. Chanan, Z. Lin, E. Yang, A. Desmaison, Andreas Kopf, A. Tejani, A. Kopf, Z. DeVito, J. Bradbury, M. Raison, L. Antiga, Martin Raison, N. Gimelshein, S. Chilamkurthy, L. Fang, Trevor Killeen, Lu Fang, J. Bai, PyTorch: An Imperative Style, High-Performance Deep Learning Library, Neural Information Processing Systems 32 (2019) 8026–8037. ARXIV_ID: 1912.01703 MAG ID: 2970971581 S2ID: 3c8a456509e6c0805354bd40a35e3f2dbf8069b1.

G. Hinton, N. Srivastava, K. Swersky, Neural networks for machine learning—lecture 6e—rmsprop: Divide the gradient by a running average of its recent magnitude., 2012. URL: :https://www.cs.toronto.edu/tijmen/csc321/slides/lectureslideslec6.pd.

C. LLC, Percentage Change Calculator, 2023. URL: https://www.calculatorsoup.com/calculators/algebra/percent-change-calculator.php.

Albumentations, Albumentations Documentation - Geometric transforms (augmentations.geometric.transforms), 2023. URL: https://albumentations.ai/docs/api_reference/augmentations/geometric/transforms/.

R. Su, S. A. Cornelissen, M. Van Der Sluijs, A. C. Van Es, W. H. Van Zwam, D. W. Dippel, G. Lycklama, P. J. Van Doormaal, W. J. Niessen, A. Van Der Lugt, et al., autotici: Automatic brain tissue reperfusion scoring on 2d dsa images of acute ischemic stroke patients, IEEE transactions on medical imaging 40 (2021) 2380–2391.

B. Zhao, H. Bilen, Dataset Condensation with Differentiable Siamese Augmentation, 2021. URL: http://arxiv.org/abs/2102.08259. doi:10.48550/arXiv.2102.08259, arXiv:2102.08259 [cs].

F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, K. H. Maier-Hein, nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation, Nature Methods 18 (2020) 203–211. MAG ID: 3112701542 S2ID: f28e387d4229c5f690ce4570a391c0f47e7155c7.

Z. Zhou, M. R. Siddiquee, N. Tajbakhsh, J. Liang, Unet++: A nested u-net architecture for medical image segmentation, DLMIA/ML-CDS@MICCAI 11045 (2018) 3–11. ARXIV_ID: 1807.10165 MAG ID: 2884436604 S2ID: a6876ea89e677a7cc42dd43f27165ff6fd414de5.

Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, E. Adeli, Yan Wang, Le Lu, A. Yuille, Yuyin Zhou, TransUNet: Transformers Make Strong Encoders for Medical Image Segmentation, arXiv.org (2021). ARXIV_ID: 2102.04306 S2ID: 24b8a0b02bcb7934967757fc59d273a71ba67e30.

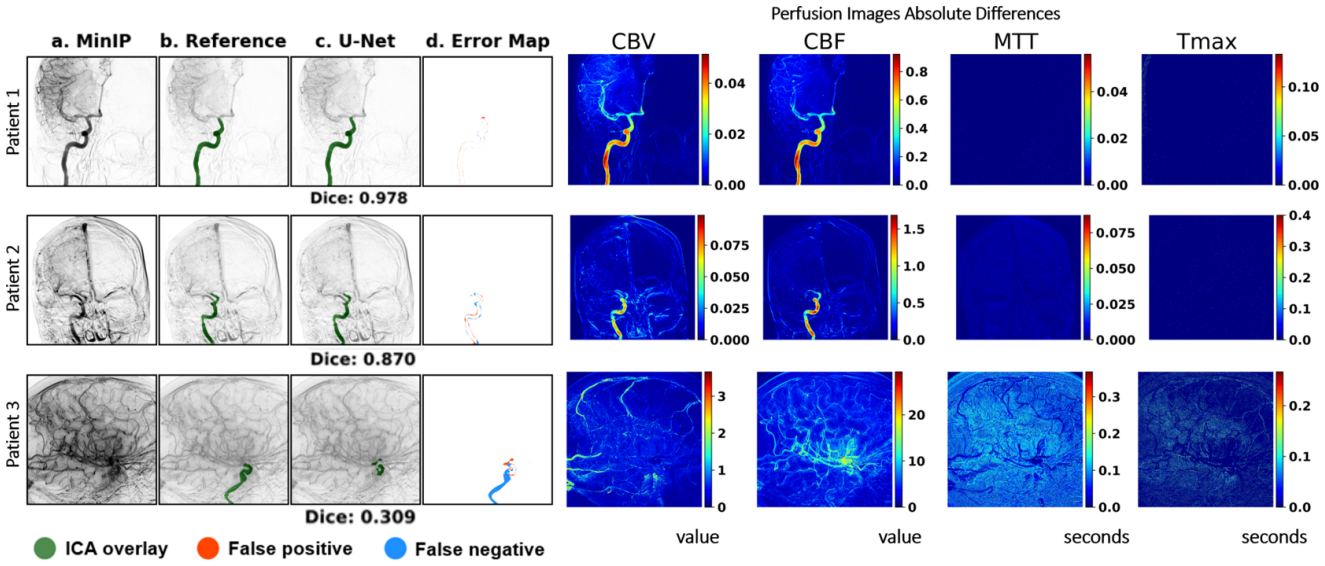

Figure 11: In this figure, the absolute difference on all perfusion images is reported for three patients: Patient 1 with the highest Dice, Patient 2 with a Dice close to the mean and Patient 3 with the lowest Dice. Columns a. to d. are explained in Figure 5.

A. Appendix