Deep Learning-Based Segmentation of the Atherosclerotic Carotid Plaque in Ultrasonic Images

Conference Paper · June 2022

DOI: 10.1007/978-3-031-08341-9_16

6 authors, including:

Georgia Liapi
Cyprus University of Technology
2 PUBLICATIONS **0** CITATIONS
SEE PROFILE

Efthyvoulos C Kyriacou
Cyprus University of Technology
199 PUBLICATIONS **3,645** CITATIONS
SEE PROFILE

Christos P Loizou
Cyprus University of Technology
198 PUBLICATIONS **2,826** CITATIONS
SEE PROFILE

Andreas S. Panayides
CYENS CoE
107 PUBLICATIONS **1,159** CITATIONS
SEE PROFILE

Deep Learning-Based Segmentation of the Atherosclerotic Carotid Plaque in Ultrasonic Images

Georgia D. Liapi[1], Efthyvoulos Kyriacou[1(✉)], Christos P. Loizou[1], Andreas S. Panayides[2], Constantinos S. Pattichis[2,3], and Andrew N. Nicolaides[4]

[1] Department of Electrical Engineering, Computer Engineering and Informatics, Cyprus University of Technology, 3036 Limassol, Cyprus
gd.liapi@edu.cut.ac.cy, efthyvoulos.kyriacou@cut.ac.cy
[2] Centre of Excellence, CYENS, Nicosia, Cyprus
[3] Department of Computer Science, University of Cyprus, Nicosia, Cyprus
[4] Vascular Screening and Diagnostic Center, Nicosia, Cyprus

Abstract. Early stroke risk stratification in individuals with carotid atherosclerosis is of great importance, especially in high-risk asymptomatic (AS) cases. In this study, we present a new computer-aided diagnostic (CAD) system for the automated segmentation of the atherosclerotic plaque in carotid ultrasound (US) images and the extraction of a refined set of ultrasonic features to robustly characterize plaques in carotid US images and videos (AS vs symptomatic (SY)). So far, we trained a UNet model (16 to 256 neurons in the contracting path; the reverse, for the expanding path), starting from a dataset of 201 (AS = 109 and SY = 92) carotid US videos of atherosclerotic plaques, from which their first frames were extracted to prepare three subsets, a training, an internal validation, and final evaluation set, with 150, 30 and 15 images, respectively. The automated segmentations were evaluated based on manual segmentations, performed by a vascular surgeon. To assess our model's capacity to segment plaques in previously unseen images, we calculated 4 evaluation metrics (mean ± std). The evaluation of the proposed model yielded a 0.736 ± 0.10 Dice similarity score (DSC), a 0.583 ± 0.12 intersection of union (IoU), a 0.728 ± 0.10 Cohen's Kappa coefficient (KI) and a 0.65 ± 0.19 Hausdorff distance. The proposed segmentation workflow will be further optimized and evaluated, using a larger dataset and more neurons in each UNet layer, as in the original model architecture. Our results are close to others published in relevant studies.

Keywords: Carotid ultrasound video · Atherosclerotic carotid plaques · Automated segmentation · Deep learning-based segmentation · Computer-aided diagnosis

Published by Springer Nature Switzerland AG 2022
I. Maglogiannis et al. (Eds.): AIAI 2022 Workshops, IFIP AICT 652, pp. 187–198, 2022.
https://doi.org/10.1007/978-3-031-08341-9_16

1 Introduction

Carotid atherosclerosis, which may cause stroke [1] is a dangerous condition, with most studies assessing its risk of occurrence by detecting and analyzing developed plaques in carotid ultrasound (US) images. A recent meta-analysis study [2] estimated that, in 2020, approximately 28% of the global population, aged 30 to 79 years, would have abnormal carotid intima-media thickness (IMT), while 21% of the people would have carotid plaques detected. Multiple studies have associated presence, size and ultrasonic appearance of carotid arteries plaques with the risk of ischemic stroke [3–5], while 16% of all ischemic strokes occur due to carotid atherosclerosis [6]. Although carotid-IMT (c-IMT) has been a valuable marker of subclinical atherosclerosis [7], the focus is extended towards the automated detection and analysis of carotid atherosclerotic plaques and their association with future stroke events. In earlier studies, our group has identified carotid asymptomatic (AS) individuals at high risk [8] by analyzing US image features of plaques. Additionally, we had investigated how texture features of segmented plaques in carotid US videos vary among AS and symptomatic (SY) individuals, during the cardiac cycles [9].

In order to automate the process of localization and visualization of atherosclerotic plaques in carotid US images, a plethora of machine learning (ML) and deep learning (DL)-based studies have been presented, during the last 10 years, with convolutional neural networks (CNNs) frequently used to segment carotid plaques effectively, such as in [10, 11] and [12]. Lately, DL models, either trained from scratch or used under transfer learning (TL), have been widely used to classify atherosclerotic plaques in carotid US images and videos, into AS or SY [13–15]. Aside from DL-based carotid US plaque image classification studies, there is notable research on plaque motion characteristics in the current imaging modality, as well as on whether these characteristics are related to stable or unstable plaques [16, 17].

In most of the above-mentioned studies, there were multiple attempts to create workflows for the accurate detection and characterization of atherosclerotic plaques in US images, a process crucial for early stroke risk assessment. Even though research in this area has been extensive over the years, the addition of extra features for better stroke risk estimation, such as video analysis parameters, is still an open research area.

In this study, we present the first steps towards the development of a new computer aided design (CAD) system for stroke risk assessment based on the analysis of US carotid images and videos. The aim of the currently presented work is to establish DL-based segmentation of US carotid plaques (see also Fig. 1, A, B and C), which is going to be part of the overall CAD system. The system presented in Fig. 1 will perform carotid US plaque analysis to support experts in decision making and stroke risk stratification. This is going to be the evolution of systems presented by our group [8, 9, 13, 17] and will be based on video analysis, extraction of new additional features and DL-based methods. More specifically, the proposed system will include: (a) a data preprocessing module (see Fig. 1, A), (b) a UNet-based model to automatically segment plaques in carotid US images (see Fig. 1, B and C), (b) extraction of ultrasonic image textural, morphological and motion features from the segmented plaque regions (see Fig. 1, D), (c) image feature selection (see Fig. 1, E), and (d) classification of selected mixed features into AS or SY, using a pre-trained ResNet50 [18] with TL (see Fig. 1, F). The ResNe50 (under Tl) has

demonstrated a promising capacity to classify AS and SY plaques in US images, based on previous work of our group [13].

The rest of the paper is organized as follows. In Sect. 2, our methodology and materials are described. Section 3 summarizes the results, which are discussed in Sect. 4.

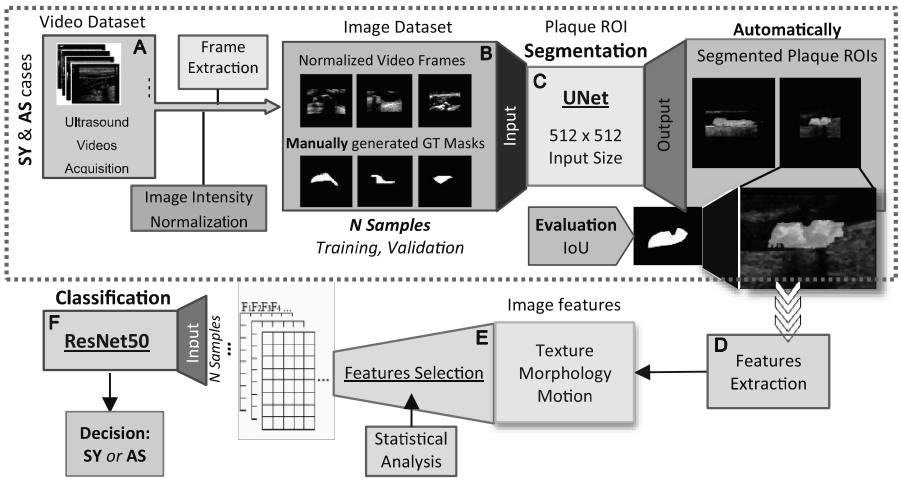

Fig. 1. General flow diagram illustrating the different modules in the CAD system proposed in this work for stroke risk stratification. **A**: Acquisition of carotid US B-mode videos, **B**: Video normalization and frame extraction, **C**: Automatic plaque ROI segmentation using a UNet and evaluation of the segmented area, **D**: Image-based textural, morphological and motion features extraction per segmented plaque ROI, **E**: Features selection based on statistical analysis, **F**: CNN-based image features classification to derive the plaque type. GT: Ground Truth, IoU: Intersection of Union; SY: Symptomatic, AS: Asymptomatic.

2 Materials and Methods

2.1 Video Dataset

A dataset of 201 carotid B-mode longitudinal US videos, from 196 subjects (123 males/73 females, AS = 108, SY = 88)), was used in this study. Five subjects had either 2 videos available or were presented with 2 plaques per video. The degree of stenosis in the subjects ranged between 50–95%. Overall, different carotid areas were included (Bifurcation, common carotid artery, external carotid artery, and internal carotid artery; both right and left sides). Plaque regions of interest (ROI) were manually selected on the first frame of each video, by an expert ultrasonographer, using a Matlab® [19]-based software created by our group. Each video was intensity-normalized, using the Video Despeckle Filtering (VDF) tool for carotid US video processing [20] (see also Fig. 1, A). Then, from each video the first frame was extracted (jpeg format, with the default OpenCv settings for the image quality), using OpenCv in Python [21–23] and resulting in a dataset of 201 different US images, with a total of 201 different plaque ROIs. Each

of these images was cropped to include the plaque ROI at the image center, but also a part of the surrounding area (maximum number of included surrounding pixels was 70).

Then, black borders were added to all frames in order to reach 512 × 512 overall image size. The image input dimensions were set to 512 × 512, in order to accommodate plaque ROIs having width higher than 500 pixels (x-axis). For each of these new images, a corresponding mask was generated (type: uint8, black and white), following given ROIs coordinates (see Fig. 1, B) and using OpenCv to derive mask contours. It is important to mention that no resizing took place for the plaque ROIs. The images were separated into training, validation and evaluation subsets, with 156 (76 SY and 80 AS), 30 (14 SY and 16 AS) and 15 (6 SY and 9 AS) images, respectively.

2.2 Model for the Automatic Plaque Segmentation

An edited version of UNet architecture [24] (see Fig. 2), for binary segmentation (plaque ROI in white, background in black), was used in Keras-Tensorflow [22, 25], and Anaconda [26]. Compared to the original model, we trained in this work, less filters per convolutional layer (CNVL, from 16 to 256 filters), due to current GPU memory constraints. Binary cross-entropy was the loss function and 'Adam' was the optimizer, used with its default learning rate value, 0.001. Data augmentation was applied, namely rotation range (80), horizontal flip, and width and height shift range (both at 0.1), with the fill mode set to 'nearest'. The UNet was trained for 220 epochs, with a batch size of 4, an input size of 512 × 512 and an early-stopping callback of patience 30. We used a computer with an NVIDIA RTX 2060 GPU (6 GB), an i7 Intel Core processor and a 16 GB RAM. The best model version was saved when the validation loss had the lowest value.

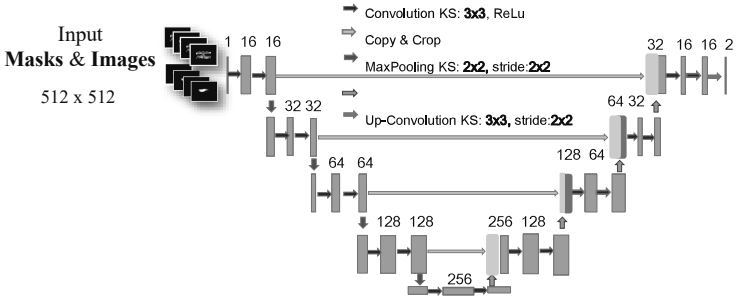

Fig. 2. The UNet architecture proposed in this work. The padding, in each convolutional layer, was set to 'same'. Kernel Size, KS.

2.3 Evaluation Metrics

In order to evaluate how similar the automatically segmented plaque ROIs were to their ground truth counterparts, we used the following evaluation metrics: 1. the Dice Similarity Coefficient (DSC), 2. the Intersection of Union (IoU), 3. the Cohen's Kappa Coefficient (KI) [27]. We also calculated the area under the receiver operating characteristic curve (AUC) and the directed Hausdorff distance (HD) [28] (see formula 4).

When evaluating model's capacity to segment, in order to obtain the segmented areas, the output of the model was thresholded (thresholding resulted probabilities), meaning that not all predictions were kept, as not all of them truly belonged to the given ROIs. To find the optimum probability threshold for the identification of the segmented ROIs, we performed predictions using all the training data, with 100 different threshold values (0.00 to 1.00) applied on the predictions and kept the threshold that yielded the highest DSC and IoU, as a guide. The metrics used to evaluate the automatic segmentations are:

$$\mathrm{DSC} = \frac{2 * \mathrm{TP}}{2 * \mathrm{TP} + \mathrm{FN} + \mathrm{FP}} \tag{1}$$

$$\mathrm{IoU} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP} + \mathrm{FN}} \tag{2}$$

$$\mathrm{KI} = \frac{(2 * (\mathrm{TP} * \mathrm{TN} - \mathrm{FN} * \mathrm{FP}))}{((\mathrm{TP} + \mathrm{FP}) * (\mathrm{FP} + \mathrm{TN}) + (\mathrm{TP} + \mathrm{FN}) * (\mathrm{FN} + \mathrm{TN}))} \tag{3}$$

$$\mathrm{HD} = \max_{a \in A} \left[\min_{b \in B} \{ d(a, b) \} \right] \tag{4}$$

where TP are the true positives, TN are the true negatives, FN are the false negatives and FP are the false positives. In formula 4, a and b are points in the sets A and B, where A is the set of points forming a ground truth plaque ROI and B is the set of points of its predicted ROI counterpart, while d(a,b) is considered as the Euclidean distance between a and b.

3 Results

After training the model, 0.34 was the threshold that yielded the highest mean DSC and IoU values. In Fig. 3 the Interquartile Ranges (IQRs) for all our evaluation metrics for the 15 automatic plaque ROI segmentations are presented for the selected threshold. Also, the resulted AUC for all the examined probability thresholds on the evaluation images reached 0.80. Table 1 also shows the mean ± std values for all evaluation metrics. We reached a 0.736 ± 0.10 mean DSC, a 0.583 ± 0.12 mean IoU and a 0.728 ± 0.10 mean KI. Measurement of the Hausdorff distance showed an 8.18 ± 2.36 mean (normalized as 0.65 ± 0.19), suggesting that our segmentation methodology needs further refinement.

In Fig. 4, we present four out of the 15 cases of the automatically segmented plaque ROIs. We see that for the first two cases (see Fig. 4 A and B), there is an agreement (overlapping), higher than 70%, when comparing to the GT with the automatically segmented area, with the corresponding DSC values both being higher than 80%. In Fig. 4 B, we notice that the model falsely detects a small area (not part of the plaque), an overall segmentation that can be possibly refined with further optimization in the model's training process.

Table 1. Evaluation metrics of the proposed segmentation system based on four model performance metrics (N = 15).

Evaluation metrics	Average (±std)
DSC	0.736 ± 0.10
IoU	0.583 ± 0.12
KI	0.728 ± 0.10
HD	0.65 ± 0.19

DSC, Dice Similarity Coefficient; IoU, Intersection of Union; KI, Cohen's Kappa Coefficient; HD, Hausdorff Distance.

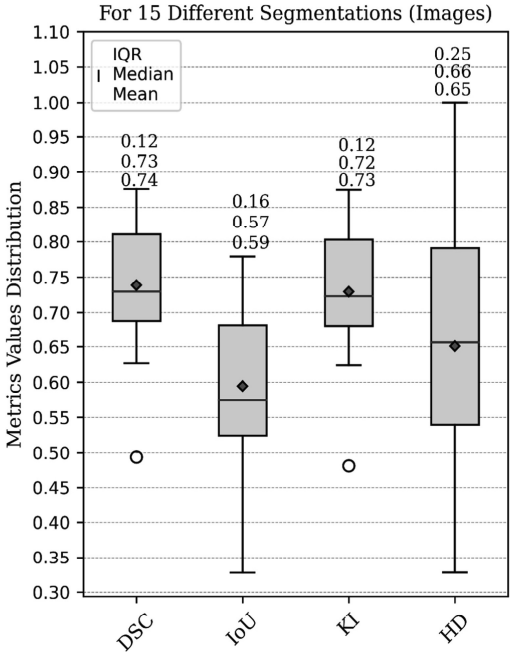

Fig. 3. Distribution of the segmentation performance evaluation metrics for all images (N = 15) investigated in this study. The ± IQR, the median and the mean values, per metric, are given at the top of each boxplot, while mean values are in red and median values are in dark blue. The HD values were normalized for visualization purposes. DSC: Dice similarity coefficient; HD, Hausdorff Distancce; IoU: Intersection of union; KI: Cohen's Kappa Coefficient. (Color figure online)

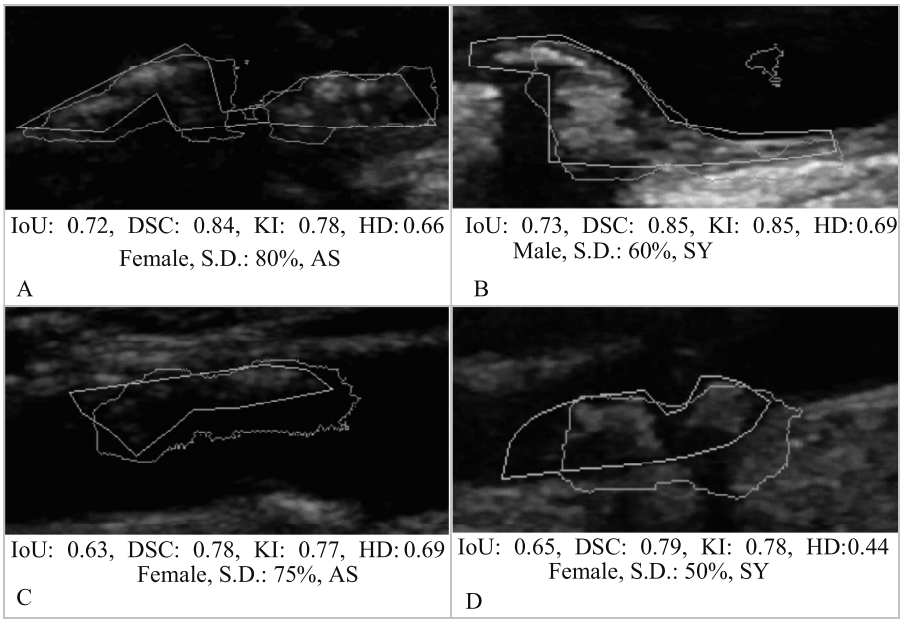

Fig. 4. Demonstration of the automatically segmented plaque versus the ground truth (GT) plaque ROIs in two different cases, AS and SY. The light-blue line is the GT plaque ROI, while the orange is the automatically segmented area, after thresholding. The DSC and the IoU values are given to depict the agreement of the GT with the automatically segmented areas, along with other patient information. For visualization purposes, the black borders were removed from each image. AS: Asymptomatic; DSC: Dice similarity coefficient; HD, Hausdorff Distancce; IoU: Intersection of union; KI: Cohen's Kappa Coefficient; S.D.: Stenosis Degree; SY: Symptomatic.

4 Discussion

Stroke risk stratification in individuals with atherosclerotic carotid plaques is of great importance, especially in cases considered as high risk, whether stenosis exists or there is a low risk. During the past ten years, there have been multiple attempts to identify valuable atherosclerotic plaque-derived features in carotid US images and videos in order to characterize the plaques based on DL methods. The two main processes, in which DL can facilitate stroke risk stratification in individuals with carotid atherosclerosis are: carotid plaque automatic segmentation and plaque classification. Studies have demonstrated the capacity of DL models to automatically segment plaque ROIs in US images or to classify plaques, in US images, into AS or SY.

In this study, we present the first step towards the development of a new CAD system for stroke risk stratification, using US carotid plaque images, along with the preliminary results we have obtained from automatic segmentation of atherosclerotic plaques in carotid US images, by training a UNet-based model. Our model is a UNet version that hosts fewer neurons per CNVL, compared to the original architecture. Initially, we applied image intensity normalization and prepared the image dataset such that no plaque ROI was resized; no plaque ROI content loss existed. With 220 epochs of training and

156 training samples, UNet segmented 15 previously unseen images, yielding DSC and IoU mean values at 0.736 ± 0.10 and 0.583 ± 0.12, respectively, with some of our DSC values being above 0.80, not very different from other relevant published results [10, 30]. Our IoU and HD resulting values suggest that our current approach needs improvement, although they might also be partially attributed to the level of image quality in our utilized dataset. Our current approach, although still premature, indeed depicts the capacity of the utilized model to segment plaques in carotid US images.

Related Work

In 2021, Jain *et al.* [12] presented a series of solo DL models, namely UNet, UNet +, SegNet, and hybrid DL models, namely SegNet-UNet and SegNet-UNet+. All models were trained with either DSC or binary cross-entropy loss to automatically segment plaques in carotid US images. They further quantified the segmented plaque ROIs (mm^2) to compare them with the GT ROIs. All models had the exact number of filters per CNVL as the standard UNet. The SegNet-UNet was their best-performing model, yielding a mean ± std IoU and DSC of 80.44 ± 1.59 88.98 ± 1.04, respectively, when trained for 100 epochs, under the K10 protocol (90% data for training and 10% data for testing), with data augmentation and with an input size of 128x128 pixels (no reported image cropping or up- or downsizing). To obtain their image dataset, they extracted 10 frames from 97 different carotid US videos (overall 970 images, with an interval of 10 per video). They have also stated that the 10 frames in each video were actually treated as different plaque ROI examples, although no metric for comparison of all GT ROIs per video was provided, in order to support their claim and despite their resulted IoU and DSC values from all their models. Also, no image intensity normalization was applied prior processing.

Zhou *et al.* [29], developed an automated plaque segmentation method, based on DL, in order to obtain the total plaque area automatically in carotid US images. They modified the original UNet architecture in terms of number of filters per CNVL, in a way similar to our UNet version, although they set 32 neurons in their first and last CNVL. They trained their model two times (two experiments, with GT masks from two different observers), for 500 epochs, using a dataset of 510 total images (2/3 used for training and 10% of the training data used for internal validation), with a batch size of 64 and used different input sizes among batches, as UNet can accept different input sizes, with the condition that the image size within each batch will be uniform. Where needed, they padded training examples to reach the mean required batch image size. They used data augmentations similar to those in our approach and they monitored the training process using the DSC loss. They used 1/3 of the primary data for testing. Their optimized model was finally trained on the whole first dataset and tested on another image US dataset. They reported a 0.05 ± 7.13 mm^2 and a 0.8 ± 8.7 mm^2 mean ± std TPA difference, when comparing each UNet with the manual segmentations. No image intensity normalization was reported.

In 2020, under a semantic segmentation approach, Xie *et al.* [30], introduced a method for plaque and vessel automatic segmentation in carotid US images based on UNet. They developed two types of UNet-based models, a two-stage model, where first a UNet architecture segments the lumen in carotid US images and its output is the input of a second UNet, which segments the plaque, and a dual-decoder, where instead of one

main decoder unit in the UNet, they set a pair of decoders, one for the lumen ROIs and one for the plaque ROIs. In the two above-mentioned models, both the plaque and lumen areas were segmented on the same output image. They trained the models for 20 epochs, with carotid US images, cropped to remove specifications surrounding the US area and resized to 224 × 224 pixel size. They applied 10-fold cross validation and had a batch size of 4. Their images were not intensity-normalized. Their highest reported DSC was 0.69 for the dual-decoder model.

Finally, Meshram *et al.* [10] developed a fully- and a semi-automatic approach for plaque segmentation in carotid US images. They compared a standard UNet with a dilated UNet, hosting dilated CNVLs in its bottleneck. They overall performed 4 experiments. In their UNet models, they also started and ended with 32 filters in the CNVLs. Their semi-automatic approach allowed a sonographer to provide a bounding box as a guide to focus on plaque ROI, followed by an exterior 12.5% buffer added around the given box, on all sides, such that each input image had 75% plaque and 25% background included. They used 862 carotid US images, separated into 90% for training (from which 5% were used for internal validation) and 10% for testing and resampled to 512 × 512 pixels for the automatic approaches, while in the semi-automatic experiments, plaque ROI bounding boxes were resampled to 256 × 256 pixels. They reached a DSC of 0.84, when comparing semi-automatic plaque segmentations from the dilated UNet to their ground truth counterparts.

Compared to all studies described above, it is clear that our model's plaque segmentation performance was highly dependent on the size of the utilized dataset and the number of neurons trained per UNet CNVL, as well as on the batch size our hardware can currently support. For these reasons, at the moment, our model's full potential to segment plaque ROIs in carotid US images is not fully examined. Also, there is a possibility that in some images, the carotid anatomy is more complex than in others, which, accompanied also by speckle noise, might hinder model's ability to generate segmented areas with clear and smooth borders.

As in all DL-based classification tasks, preparation of the image dataset is of great importance, our primary focus in this study was to avoid content loss or resampling in the carotid US images we used as input to our UNet model. In contrast, in [12, 30] and [10], plaque ROIs were resized. More reliably, in [29], cropping and padding was applied in the input images, only where needed. Secondary, the preparation of our image dataset, where not all of the plaque ROI surrounding was included is similar to the data preparation approach in [29], in order to acquire the desirable input size, and to that in [10], regarding the area covered by the plaque in each image's carotid content.

As shown from the above-explained studies, it is still not clear if plaque automatic segmentation in carotid US images is improved when the input includes mainly the plaque ROI (with some carotid surrounding included) and when no intervention in the content occurs (resampling).

In our future steps, we will repeat the current segmentation experiment using carotid US whole images, the portable graphics format (png) for lossless data compression and a 5-fold cross validation workflow. We will have a larger dataset and a GPU of higher capacity, in order to optimize more parameters simultaneously, with a larger batch size. We will also extract segmentation metrics from a larger evaluation dataset,

and we will investigate if speckle removal facilitates UNet-based automatic carotid plaque segmentation in US images. Finally, different comparison metrics between the automatically segmented and the GT total plaque areas will be included.

By following the above-described steps, we expect to have a robust UNet-based segmentation model that will be the heart of our developing CAD program, which will be used by experts for automated plaque characterization in carotid US images and videos, in clinical routine.

5 Limitations

Our current preliminary results for the automated plaque segmentation in carotid US images is accompanied by some limitations. Regarding the preparation of the dataset, the image intensity normalization (normalization applied on the video level), required a considerable amount of time and user interaction. Additionally, we had to prepare an automated workflow to acquire an appropriate dataset of a uniform size from plaques with quite different area sizes. For the training process of the segmentation model, we could obtain improved segmentation results by following the exact UNet architecture. It should be however noted that our current hardware is not designed for such an experiment, yet. In the automatically segmented plaque ROIs, after the thresholding process, the borders were noisy. This possibly implies that speckle noise or primary video quality play a role, both of which should be further investigated.

Funding – Acknowledgements. This study is funded by the project 'Atherorisk' "Identification of unstable carotid plaques associated with symptoms using ultrasonic image analysis and plaque motion analysis", code: Excellence/0421/0292, funded by the Research and In-novation Foundation, the Republic of Cyprus.

References

1. Flaherty, M.L., Kissela, B., Khoury, J.C., Alwell, K., et al.: Carotid artery stenosis as a cause of stroke. Neuroepidemiology **40**(1), 36–41 (2013)
2. Song, P., Fang, Z., Wang, H., Cai, Y., et al.: Global and regional prevalence, burden, and risk factors for carotid atherosclerosis: a systematic review, meta-analysis, and modelling study. Lancet Glob. Health **8**(2), 721–729 (2020)
3. Hollander, M., Bots M.L, Iglesias del Sol A., Koudstaal P.J et al.: Carotid plaques increase the risk of stroke and subtypes of cerebral infarction in asymptomatic elderly: The Rotterdam Study. Circulation 105(24), 2872–2877 (2002)
4. Paraskevas, K.I., Nicolaides, A.N., Kakkos, S.K.: Asymptomatic Carotid Stenosis and Risk of Stroke (ACSRS) study: what have we learned from it? Annals Translational Med. **8**(19), 1271 (2020)
5. Howard, D.P.J., Gaziano, L., Rothwell, P.M.: Risk of stroke in relation to degree of asymptomatic carotid stenosis: a population-based cohort study, systematic review, and meta-analysis. Lancet Neurol. **20**(3), 193–202 (2021)
6. Petty, G.W., Brown, R.D., Whisnant, J.P., Sicks, J.D., et al.: Ischemic stroke subtypes: a population-based study of incidence and risk factors. Stroke **30**(12), 2513–2516 (1999)

7. Stein, J.H., Korcarz, C.E., Hurst, R.T., Lonnet, E., et al.: Use of carotid ultrasound to identify subclinical vascular disease and evaluate cardiovascular disease risk: a consensus statement from the American Society of Echocardiography carotid intima-media thickness task force endorsed by the Society for Vascular Medicine. J. Am. Soc. Echocardiogr. **21**(2), 93–111 (2008)
8. Kyriacou, E.C., Petroudi, S., Pattichis, C.S., Pattichis, M.S., et al.: Prediction of high-risk asymptomatic carotid plaques based on ultrasonic image features. IEEE Trans. Inform. Technol. Biomed. **16**(5), 966–973 (2012)
9. Loizou, C.P., Pattichis, C.S., Pantziaris, M., Kyriacou, E.C., Nicolaides, A.: Texture feature variability in ultrasound video of the atherosclerotic carotid plaque. IEEE J. Transl. Eng. Health Med. **5**, 1–9 (2017)
10. Meshram, N.H., Mitchell, C.C., Wilbrand, S., Dempsey, R.J., Varghese, R.J.: Deep learning for carotid plaque segmentation using a dilated U-Net architecture. Ultrason. Imaging **42**(4–5), 221–230 (2020)
11. Del Mar Vila, M., Remeseiro, B., Grau, M., Elosua, R., et al.: Semantic segmentation with DenseNets for carotid artery ultrasound plaque segmentation and CIMT estimation. Artif. Intell. Med. **103**, 101784 (2020)
12. Jain, P.K., Sharma, N., Giannopoulos, A.A., Saba, L., et al.: Hybrid deep learning segmentation models for atherosclerotic plaque in internal carotid artery B-mode ultrasound. Comput. Biol. Med. **136**, 104721 (2021)
13. Panayides, A., Kyriacou, E.C., Nicolaides, A., Pattichis, C.S.: Stroke risk stratification using transfer learning. In: 41st Engineering in Medicine and Biology Conference (EMBC), Berlin, Germany (2019)
14. Guang, Y., He, W., Ning, B., Zhang, H., et al.: Deep learning-based carotid plaque vulnerability classification with multicentre contrast-enhanced ultrasound video: a comparative diagnostic study. BMJ Open **11**(8), 047528 (2021)
15. Sanagala, S.S., Nicolaides, A., Gupta, S.K., Koppula, V.K., et al.: Ten fast transfer learning models for carotid ultrasound plaque tissue characterization in augmentation framework embedded with heatmaps for stroke risk stratification. Diagnostics **11**(11), 2109 (2021)
16. Golemati, S., Patelaki, E., Gastounioti, A., Andreadis, I., et al.: Motion synchronisation patterns of the carotid atheromatous plaque from B-mode ultrasound. Sci. Rep. **10**(1), 11221 (2020)
17. Giannopoulos, A.A., Kyriacou, E., Griffin, M., Pattichis, C.S., et al.: Dynamic carotid plaque imaging using ultrasonography. J. Vasc. Surg. **73**(5), 1630–1638 (2021)
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. arXiv:1512.03385 [cs] (2015)
19. MATLAB and Statistics Toolbox Release 2012b, The MathWorks, Inc., Natick, Massachusetts, United States (2012)
20. Loizou, C.P., Kasparis, T., Christodoulides, P., Theofanous, C., et al.: Despeckle filtering toolbox for medical ultrasound video. Int. J. Monitoring Surveillance Technol. Res. **1**(4), 61–79 (2013)
21. Bradski, G., Kaehler, A.: Learning OpenCV: Computer vision with the OpenCV library. Reilly Media, Inc. (2008)
22. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., et al.: TensorFlow: large-scale machine learning on heterogeneous systems, Software available from tensorflow.org. (2015)
23. Raybaut, P.: Spyder-documentation. Available online at: pythonhosted org. (2009)
24. Ronneberger O., Fischer P., Brox T. U-Net: Convolutional networks for biomedical image segmentation. arXiv:1505.04597 [cs] (2015)
25. Chollet, F., et al.: Keras. Available at: https://github.com/fchollet/keras. (2015)
26. Anaconda Software Distribution. Computer software Version 2–2.4.0. Anaconda (2016)

27. Cohen, J.: A coefficient of agreement for nominal scales. Educ. Psychol. Measur. **20**(1), 37–46 (1960)
28. Taha, A.A., Hanbury, A.: An efficient algorithm for calculating the exact Hausdorff distance. IEEE Trans. Pattern Anal. Mach. Intell. **37**, 2153–2163 (2015)
29. Zhou, R., Azarpazhooh, M.R., Spence, J.D., Hashemi, S., et al.: Deep learning-based carotid plaque segmentation from B-Mode ultrasound images. Ultrasound Med. Biol. **47**(9), 2723–2733 (2021)
30. Xie, M., Li, Y., Xue, Y., Huntress, L., et al.: Two-stage and dual-decoder convolutional U-Net ensembles for reliable vessel and plaque segmentation in carotid ultrasound images. In: 19th IEEE International Conference on Machine Learning and Applications (ICMLA), pp. 1376–1381, Miami, FL, USA (2020)