



HAL
open science

Carotid artery wall segmentation in ultrasound image sequences using a deep convolutional neural network

Nolann Lainé, Herve Liebgott, Guillaume Zahnd, Maciej Orkisz

► To cite this version:

Nolann Lainé, Herve Liebgott, Guillaume Zahnd, Maciej Orkisz. Carotid artery wall segmentation in ultrasound image sequences using a deep convolutional neural network. International Conference on Computer Vision and Graphics, Sep 2022, Varsovie, Poland. pp.1-12. hal-03898051

HAL Id: hal-03898051

<https://hal.science/hal-03898051v1>

Submitted on 14 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Carotid artery wall segmentation in ultrasound image sequences using a deep convolutional neural network

Nolann Lainé¹, Hervé Liebgott¹, Guillaume Zahnd², and Maciej Orkisz¹

¹ Univ Lyon, Université Claude Bernard Lyon 1, INSA-Lyon, CNRS, Inserm, CREATIS UMR 5220, U1294, F-69621, Lyon, France

{nolann.laine,herve.liebgott,maciej.orkisz}@creatis.insa-lyon.fr,
WWW home page: <https://www.creatis.insa-lyon.fr/site7/en>

² Institute of Biological and Medical Imaging, Helmholtz Zentrum München, Neuherberg, Germany

guillaume.zahnd@helmholtz-muenchen.de

Abstract. Intima-media thickness (IMT) of the common carotid artery is routinely measured in ultrasound images and its increase is a marker of pathology. Manual measurement being subject to substantial inter- and intra-observer variability, automated methods have been proposed to find the contours of the intima-media complex (IMC) and to deduce the IMT thereof. Most of them assume that these contours are smooth curves passing through points with strong intensity gradients expected between artery lumen and intima, and between media and adventitia layers. These assumptions may not hold depending on image quality and arterial wall morphology. We therefore relaxed them and developed a region-based segmentation method that learns the appearance of the IMC from data annotated by human experts. This deep-learning method uses the dilated U-net architecture and proceeds as follows. First, the shape and location of the arterial wall are identified in full-image-height patches using the original image resolution. Then, the actual segmentation of the IMC is performed at a finer spatial resolution, in patches distributed around around the location thus identified. Eventually, the predictions from these patches are combined by majority voting and the contours of the segmented region are extracted. On a public database of 2676 images the accuracy and robustness of the proposed method outperformed state-of-the-art algorithms. The first step was successful in 98.7% of images, and the overall mean absolute error of the estimated IMT was of $100 \pm 89 \mu\text{m}$.

Keywords: segmentation, ultrasound images, deep learning.

1 Introduction

Atherosclerosis, the “silent killer” that can progress to acute events such as stroke or heart attack, begins to develop in the intima-media complex (IMC), the innermost structure of the artery wall. IMC thickening is widely considered as a

marker of the atherosclerosis onset. Its screening is usually performed in ultrasound (US) images of the carotid artery. In the clinical routine, the measurement of the intima-media thickness (IMT) is performed manually, which may lead to substantial inter- and intra-observer variability.




Fig. 1. Longitudinal US images of the common carotid artery, examples. The curves represent the lumen-intima (red) and media-adventitia (green) boundaries of the far wall interpolated from expert’s annotations (dots). The annotations were restricted to the exploitable region (ROI) where the characteristic double-line pattern of the IMC is discernible. Yellow arrows point examples of inconsistent annotations.

Numerous methods, compared in several surveys [10, 2, 3], have been developed to automate the measurement of the carotid artery IMT. Most of them infer the IMT as the distance between the lumen-intima (LI) and media-adventitia (MA) contours extracted within a region of interest (ROI) cropped manually or automatically by seeking the window where the characteristic dual-line pattern is discernible (Fig. 1). Contour extraction is performed using various “conventional” algorithms, such as snakes or dynamic programming. Virtually all assume that the contours are smooth curves passing through points with strong intensity gradients expected between artery lumen and intima, and between media and adventitia layers. Lately, deep-learning-based (DL) methods have also been proposed for this task and obtained promising results, but the comparisons [1] were performed on different datasets. A recent study [6] compared five state-of-the-art conventional methods on an open access dataset CUBS1 [7] containing 2176 US images from two centers, and concluded that their accuracy is comparable to skilled human experts, while presenting substantially smaller variability.

In that study, the best results were achieved by a method based on dynamic programming [15]. A subsequent study [4] confronted the same methods with two DL approaches on another dataset CUBS2 [5] containing 500 US images from five different centers. This comparison has demonstrated the sensitivity of several methods to the arterial wall morphology, namely a decreased accuracy in curved or inclined arteries. One of the DL-based methods, previously unpublished, outperformed the others in almost all comparisons except for processing time. While the selected network architecture was the well-known U-net [12], the original contribution of the method resides in a specific patch-based strategy devised to cope with inclined and curved arteries. In this article, we describe for the first time the details of this method, and then report main results in terms of accuracy and robustness on both open-access datasets CUBS1 [7] and CUBS2 [5], compared to available expert annotations and the other methods. We also assess how much the processing time can be improved by changing the parameter setting, without degrading the accuracy.

2 Datasets

In this work, for training and evaluation purposes, we have used both previously mentioned publicly-available datasets CUBS1³ and CUBS2⁴, *i.e.*, a total of 2676 static US images acquired with different clinical equipment and including a subset ($n = 100$) of simulated images. These images have pixel sizes ranging from 29 to 267 μm (mean size 64 μm in CUBS1 and 60 μm in CUBS2) and come along with LI and MA contours independently traced by two or more experts, one of them having traced them twice. The experts independently selected a ROI, where they considered the LI and MA interfaces as sufficiently perceptible, and traced control points within it. To obtain smooth contours, piecewise cubic Hermite interpolating polynomial (*PCHIP*) was applied using *MATLAB, Version 2020b* (The Math Works, Inc.). The same annotations as in the previous studies [6, 4] were used as reference. These will be referred to as A1. According to A1, the mean IMT was of 725 μm in CUBS1 and 857 μm in CUBS2. The same expert’s second set of annotations (A1’) was used to calculate the intra-observer variability, while the inter-observer variability was calculated between A1 and the second expert’s annotations (A2).

3 Segmentation method

Similarly to many existing methods, *e.g.* [15], the first step is a very simple user interaction (two mouse clicks) defining the exploitable ROI, where the IMC is perceptible. The remainder is fully automatic and split into two main steps: localization of the far wall and actual IMC segmentation (Fig. 2).

³ <https://data.mendeley.com/datasets/fpv535fss7/1>

⁴ <https://data.mendeley.com/datasets/m7ndn58sv6/1>

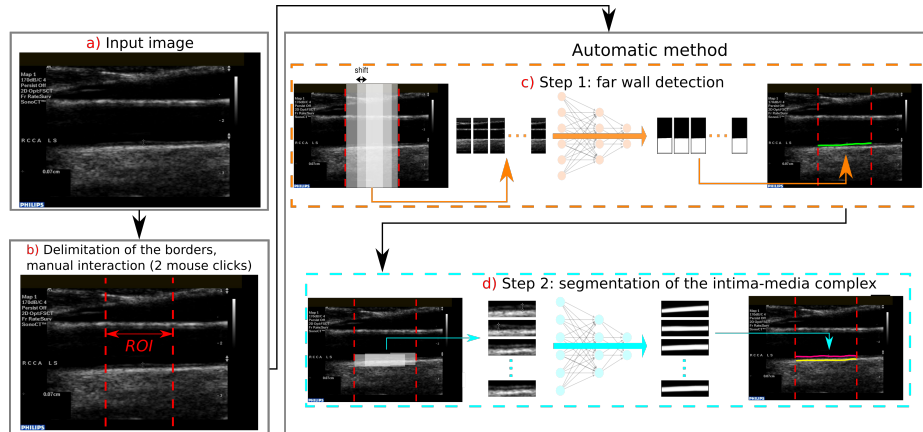


Fig. 2. Flowchart of the proposed method: a) Input image. b) User delimitation of the left and right borders of the ROI. c) Far wall detection: patches are extracted through a sliding window with overlap within ROI borders; post-processing of predicted overlapping masks leads to extraction of the median axis (green). d) IMC segmentation: overlapping patches are picked along the median axis and post-processing of the predicted masks leads to extraction of the LI (red) and MA (yellow) interfaces.

The proposed solution is based on U-net architecture [12], which has widely demonstrated its ability to produce accurate results in medical image segmentation with limited annotated data available for supervised training. We have kept standard components, such as ReLU activation function, and experimentally selected the following settings: number of layers 5, initial number of convolution kernels 32, and size of the kernel 3×3 (with batch normalization and bias). We implemented dilated convolutions on the bottleneck, to increase the receptive field [9]. The outputs of the filters with different dilation factors $\{1, 2, 3, 4, 5, 6\}$ are concatenated into one single tensor. The U-net operates in fixed-size (128 -pixel width W , 512 -pixel height H) overlapping patches distributed within the ROI. The two first layers contract only the height, so as to achieve a square shape (128×128) of the feature space starting from the third layer. The horizontal overlap between patches is equal to $W - \Delta x$, where Δx is the horizontal shift. A post-processing combines the predictions made within the patches to extract smooth contours over the entire ROI regardless of its width. The core of the method consists of two steps: approximately localizing the far wall (Fig. 2c), and precisely segmenting the IMC around this location (Fig. 2d).

3.1 Detection of the far wall

Like in many state-of-the-art methods [14, 13, 8, 11], the initial localization of the far wall is performed to focus the subsequent actual segmentation. In this step, the algorithm attempts to separate the ROI in two regions: above and below the median axis of the IMC, respectively. Here, the patches are of full image height,

and the corresponding U-net will be referred to as Θ_{FW} . We first describe how the images were pre-processed and how Θ_{FW} was trained, then we specify the post-processing chosen to obtain the curve approximately localizing the far wall on the entire ROI width from patch-wise predictions inferred using Θ_{FW} .

Pre-processing and training: All images of the database were resampled to a constant height of 512 pixels; as the native height of all images in the database is around 600 pixels, the distortion thus introduced was minimal (regions containing alphanumerical information and/or ECG curve overlaid onto the actual image were not clipped). For training data, the median axis of the IMC was defined as the line halfway between LI and MA annotations, interpolated across the entire width of the ROI, and a reference mask (M_{ROI}) was generated by setting all pixels below the median axis to one and the others to zero. Then the ROI and M_{ROI} were identically cut into patches; a 100-pixel overlap ($\Delta x = 28$) between patches aimed at data augmentation. Thus obtained patches with their associated masks (Fig. 3) were fed into the training process, which used the ADAM optimizer and the sum of the binary cross-entropy and the Dice loss as the loss function. The latter was experimentally chosen to minimize the boundary distance and maximize the overlay with respect to the reference masks. Training was performed with on-the-fly data augmentation: horizontal and vertical flip applied on 50% of the images, and affine transformation with random rotation $[-2, 2]$ degrees, shearing $[-2, 2]$ degrees, and translation (vertical $[-20, 20]$ pixels, horizontal $[-5, 5]$ pixels). The batch size was of 32, and the number of epochs 300 with early stopping switched on, so that the process stopped if the loss value on the validation subset did not improve during 50 epochs.

Inference and post-processing: Prior to inference, each image is resampled as described above, and then the corresponding ROI is cut into overlapping 128×512 -pixel patches. Next, all patches are segmented using Θ_{FW} . Knowing the location and the size of each patch, two maps are created:

- **prediction map:** contains, for each pixel, the sum of values predicted by Θ_{FW} in all patches.
- **overlay map:** contains, for each pixel, the number of overlapping patches it belonged to.

Dividing the prediction map by the overlay map provides, for each pixel, an average value in the range $[0, 1]$, which is then binarized by using a threshold of 0.5, to obtain the segmentation map. The latter is cleaned by retaining the largest connected component. The median axis we seek is the upper boundary of thus segmented region. Eventually, this boundary is smoothed by using a third-order polynomial regression.

3.2 Segmentation of the IMC

The above-described approximation of the far-wall median axis is used to initialize the actual segmentation of the IMC. This segmentation process presents




Fig. 3. Data used during the far wall detection training phase: patch cut from the ROI (left) and the corresponding patch from M_{ROI} . Red and cyan dots represent the annotations for LI and MA interfaces, respectively. The green curve is the median axis calculated from the interpolated annotations.

many similarities with concepts explained in Section 3.1: overlapping patches of 128×512 pixels, an overlay map, a prediction map, a similar post-processing except that two contours are extracted (the LI and MA interfaces), as well as the same optimizer, loss function, and U-net architecture. The dilated U-net trained here will be referred to as Θ_{IMC} . Hereafter, we emphasize the specific choices made for this step.

Pre-processing and training: The segmentation task has to be as accurate as possible, hence the algorithm works at a sub-pixel resolution. To this purpose, the vertical pixel size of the images was homogenized to $5 \mu\text{m}$ using a linear interpolation. According to this physical size, the patch height $H = 512$ pixels roughly corresponds to 2.6 mm, which aims to encompass the IMC, knowing that the average IMC thickness is about 0.8 mm. For training, the ground truth was then deduced from thus interpolated images (Fig. 4): each pixel located between the annotated LI and MA interfaces was set to one, and the others to zero. Unlike the far wall detection, the patches were extracted along the median axis and, in addition to the 100-pixel horizontal overlap (*i.e.*, $\Delta x = W - 100 = 28$ -pixel shift), at each abscissa x_i , the mean ordinate y_i of the median axis was computed on the patch width $W = 128$, and three patches were extracted, respectively centered at y_i and $y_i \pm \Delta y$ ($\Delta y = 128$). This data augmentation attempted to cope with possibly inaccurate far-wall approximation as well as with tilted arteries.

Inference and post-processing: During inference, the patches are extracted

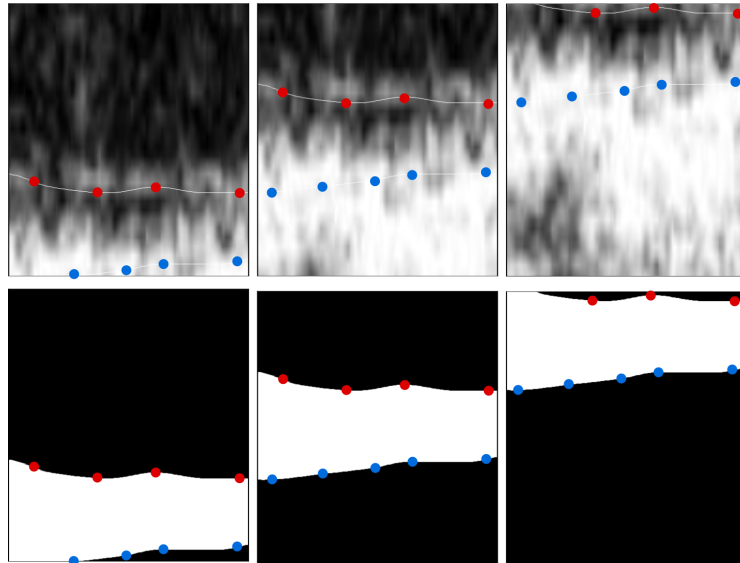


Fig. 4. Data used during the training phase for IMC segmentation. Patches and their associated masks located at: a) $(x_i, y_i - 128)$, b) (x_i, y_i) , and c) $(x_i, y_i + 128)$. Red and blue dots represent the corresponding annotations for LI and MA interfaces, respectively. The contours were obtained by interpolating the annotations. Of note: the actual size of the patches is 128×512 pixels, but they have been vertically compressed for display purposes.

along the far wall approximation resulting from the first step (Section 3.1). At each abscissa x_i three or more patches are captured at different ordinates, depending on the tilt of the median axis, with the goal to cover all the expected extent of the IMC. The predictions made by Θ_{IMC} in all patches are combined into a prediction map, and then the segmentation map is derived thereof, as described above. Finally, the LI and MA interfaces are respectively defined as the upper and lower boundaries of the region thus segmented.

4 Results

The evaluation was carried out using 5-fold cross-validation, so as to assess each network on data not seen during its training. In each fold, the database (combined data from the open-access datasets CUBS1 [7] and CUBS2 [5]) was split into training (60%), validation (20%), and testing (20%) subsets. Thus, five pairs of networks Θ_{FW} and Θ_{IMC} were trained and tested independently, and the results reported here are the merging of the test sets of these five pairs, thus evaluating the method on the entire database.

In the proposed cascade approach, a failure of the first step (far wall detection) will trigger a failure of the second step (*IMC* segmentation). To conduct

Table 1. Errors of segmentation results (mean \pm standard deviation): mean absolute difference (MAD) for thickness quantification (IMT) and Hausdorff distance (HD) for contour locations (LI and MA). The comparisons are reported for the entire database and separately for CUBS1 and CUBS2 datasets. For each dataset, the following differences are reported: the proposed method against reference annotations, the inter- and intra-observer variabilities.

Dataset	Comparison	Measure		
		IMT MAD (μm)	LI HD (μm)	MA HD (μm)
CUBS1 \cup CUBS2	Method vs. A1	100 \pm 89	317 \pm 193	287 \pm 152
	A2 vs. A1	204 \pm 168	370 \pm 197	348 \pm 162
	A1' vs. A1	147 \pm 127	356 \pm 194	324 \pm 161
CUBS1	Method vs. A1	99 \pm 89	320 \pm 193	287 \pm 153
	A2 vs. A1	206 \pm 168	380 \pm 207	351 \pm 161
	A1' vs. A1	144 \pm 123	357 \pm 204	319 \pm 155
CUBS2	Method vs. A1	106 \pm 89	305 \pm 197	289 \pm 147
	A2 vs. A1	192 \pm 166	327 \pm 138	338 \pm 184
	A1' vs. A1	160 \pm 140	352 \pm 140	346 \pm 185

a fair evaluation of both steps, we first quantified the success rate of the first step alone, and then we quantified the accuracy of the second step by manually enforcing valid initial conditions when needed.

Robustness of the far wall detection: After visual inspection, 36 out of 2676 predicted median axes (1.3% of the database) were considered as failures, *i.e.*, curves unusable to initialize the IMC segmentation step. Hence, the success rate was of 98.7% and, in the 36 images with failures, the median axis was manually redrawn using a home-made graphical interface.

Accuracy of the IMC segmentation: The segmentation error was quantified in two ways. To enable a fair comparison of our method with five state-of-the-art methods evaluated on the CUBS1 and CUBS2 datasets in the previously mentioned recent studies [6, 4], we asked the first authors of the latter (see Acknowledgements) to apply onto the LI and MA contours extracted by our method exactly the same evaluation protocol and metrics as used in their publications. Hence, the IMT values and all errors were calculated on a region restricted to a common support where all participating methods succeeded in extracting the contours. The IMT was measured as the polyline distance between LI and MA contours, and its error was calculated as the mean absolute difference (MAD) between the method output and the reference values calculated from the expert’s (A1) annotations. The (worst case) errors for LI and MA were separately estimated by calculating the Hausdorff distances between the respective contours extracted by the method and the annotations performed by A1. These results are summarized in Table 1.

The reported values were obtained with the same implementation and parameter settings as in the previous study restricted to the CUBS2 dataset [4]. In that study, the processing time was the only drawback of our method: 0.92 sec for far-wall localization and 1.77 sec for IMC segmentation, *i.e.*, a total of 2.69 sec

per image, and was achieved with a very large horizontal overlap of 124 pixels (*i.e.*, shift $\Delta x = 4$). Therefore, without modifying the implementation, we varied $\Delta x \in \{8, 16, 32, 64, 96\}$ to check how much the horizontal overlap – and thereby the number of processed patches – can be decreased without degrading the results, and how much the processing time can thus be improved. The results obtained on the entire database are summarized in Table 2. Here the IMT and the errors in contour extraction were assessed using in-house tools, by measuring the column-wise distances instead of the polyline and Hausdorff distances. Also, the comparison was not restricted to the common support for all methods, and used the full extent of the extracted contours. In addition, the failures of the automatic far-wall localization step were not manually corrected. For all these reasons, the results for $\Delta x = 4$ differ from those in Table 1, but they can be consistently compared with those obtained for the remaining shifts.

Table 2. Estimation errors (mean of absolute differences \pm their standard deviation) with respect to reference annotations A1, and processing time (FW stands for far-wall localization and IMC for the actual segmentation), as a function of the horizontal shift Δx . The horizontal overlap between patches is equal to $128 - \Delta x$. The reported processing time was achieved on the following hardware: CPU Intel(R) Core(TM) i7-6700, 32 GB RAM, 3.40 GHz; GPU NVIDIA GeForce GTX 1070, 8 GB RAM.

Shift Δx (pixels)	4	8	16	32	64	96
IMT MAD \pm std (μm)	158 \pm 88	158 \pm 88	158 \pm 88	159 \pm 88	159 \pm 89	160 \pm 87
LI MAD \pm std (μm)	116 \pm 68	116 \pm 68	116 \pm 68	113 \pm 65	114 \pm 66	118 \pm 71
MA MAD \pm std (μm)	104 \pm 62	104 \pm 62	104 \pm 63	107 \pm 68	107 \pm 68	105 \pm 62
FW time (sec)	0.92	0.73	0.64	0.59	0.58	0.56
IMC time (sec)	1.77	1.14	0.80	0.65	0.58	0.55
total time (sec)	2.69	1.87	1.44	1.24	1.16	1.11

5 Discussion

We developed and assessed a deep-learning method to extract the contours of the intima-media complex in longitudinal B-mode ultrasound images of the carotid artery. The method first approximately localizes the far wall, and then segments the anatomical interfaces of interest. The proposed approach allows for variable width of the region of interest without rescaling the images. Robustness of the far-wall localization step is a prerequisite for overall correct segmentation. This step was successful in 98.7% of the images, therefore attesting of the method robustness. The actual segmentation step achieved good accuracy, with errors smaller than the inter-observer and intra-observer variability. The results reported in Section 4 were obtained on combined data from the open-access datasets CUBS1 [7] and CUBS2 [5] and can be compared with the results of five state-of-the-art methods evaluated on the same datasets in recent studies [6, 4]. On both datasets, our method outperformed the remaining ones. In particular,

on the CUBS1 dataset, the compared methods produced IMT errors (MAD) ranging from 114 ± 117 to $255 \pm 230 \mu\text{m}$, while for our method the MAD was of $99 \pm 89 \mu\text{m}$. Similarly, on the CUBS2 dataset, the other methods produced IMT errors ranging from 139 ± 118 to $224 \pm 178 \mu\text{m}$, while for our method the MAD was of $106 \pm 89 \mu\text{m}$.

On average, the MAD produced by the proposed method represented 1.6 pixels (which corresponds to 13.7% of the mean IMT) in CUBS1, and 1.8 pixels (12.4%) in CUBS2, respectively. Although these inaccuracies remain relatively large compared to the target measurement (IMT), they are respectively twice and 1.5 times smaller than the inter-observer and intra-observer variabilities.

The processing time was below 1 sec for the far-wall localization step and below 2 sec for the actual IMC segmentation step, when using a very large horizontal overlap and thereby a large number of patches. It can be seen that, on average, the errors were stable when decreasing the overlap, particularly up to the shift $\Delta x = 16$, but the resulting reduction of the processing time was limited. Of note, when segmenting a sequence of images, the far-wall localization step needs to be executed only once, at the beginning of the sequence, while the actual IMC segmentation needs to be repeated for each image. In the current implementation, the processing time of this step remains greater than 0.5 sec even with a strongly reduced overlap. Further reduction of this time would require an optimization of the implementation, which has not yet been attempted.

The largest errors occurred in the presence of atherosclerotic plaques. As the work presented here was oriented towards asymptomatic plaque-free subjects, images with plaques were not expected. Nevertheless, we anticipate that results might be improved by increasing the number of such images in the database and re-training the networks.

In conclusion, with a 98.7% success rate and errors smaller than the inter- and intra-observer variability, the proposed method likely is robust and accurate enough to be used in clinical practice, as well as to study the periodic compression-decompression of the arterial wall, which is another potential biomarker of atherosclerosis.

Code and examples are available at <https://github.com/nl3769/caroSegDeep>.

6 Acknowledgments

The authors thank K. M. Meiburger and F. Marzola for their help in calculating the comparisons with reference annotations on CUBS1 and CUBS2 datasets. This work was partly supported, via NL’s doctoral grant, by the LABEX PRIMES (ANR-11-LABX-0063) of Université de Lyon, within the program "Investissements d’Avenir" (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR).

7 Compliance with ethical standards information

The data from human subjects used in this work were obtained and treated in line with the principles of the Declaration of Helsinki. Approval was granted by the Ethics Committees of the institutions involved in creating the multicentric database, from which these data were accessed.

References

1. Biswas, M., Saba, L., Omerzu, T., Johri, A.M., Khanna, N.N., Viskovic, K., Mavrogeni, S., Laird, J.R., Pareek, G., Miner, M., Balestrieri, A., Sfikakis, P.P., Protogerou, A., Misra, D.P., Agarwal, V., Kitas, G.D., Kolluri, R., Sharma, A., Viswanathan, V., Ruzsa, Z., Nicolaides, A., Suri, J.S.: A review on joint carotid intima-media thickness and plaque area measurement in ultrasound for cardiovascular/stroke risk monitoring: Artificial intelligence framework. *Journal of Digital Imaging* 34(3), 581–604 (2021)
2. Loizou, C.P.: A review of ultrasound common carotid artery image and video segmentation techniques. *Medical & Biological Engineering & Computing* 52(12), 1073–1093 (2014)
3. Meiburger, K.M., Acharya, U.R., Molinari, F.: Automated localization and segmentation techniques for B-mode ultrasound images: A review. *Computers in Biology and Medicine* 92, 210–235 (2018)
4. Meiburger, K.M., Marzola, F., Zahnd, G., Faita, F., Loizou, C., Lainé, N., Carvalho, C., Steinman, D., Gibello, L., Bruno, R.M., Clarenbach, R., Francesconi, M., Nikolaides, A., Liebgott, H., Campilho, A., Ghotbi, R., Kyriacou, E., Navab, N., Griffin, M., Panayiotou, A., Gherardini, R., Varetto, G., Bianchini, E., Pattichis, C., Ghiadoni, L., Rouco, J., Orkisz, M., Molinari, F.: Carotid ultrasound boundary study (CUBS): Technical considerations on an open multi-center analysis of computerized measurement systems for intima-media thickness measurement on common carotid artery longitudinal B-mode ultrasound scans. *Computers in Biology and Medicine* 144, 105333 (2022)
5. Meiburger, K.M., Marzola, F., Zahnd, G., Faita, F., Loizou, C., Lainé, N., Carvalho, C., Steinman, D., Gibello, L., Bruno, R.M., Clarenbach, R., Francesconi, M., Nikolaides, A., Liebgott, H., Campilho, A., Ghotbi, R., Kyriacou, E., Navab, N., Griffin, M., Panayiotou, A., Gherardini, R., Varetto, G., Bianchini, E., Pattichis, C., Ghiadoni, L., Rouco, J., Orkisz, M., Molinari, F.: DATASET for "Carotid Ultrasound Boundary Study (CUBS): Technical considerations on an open multi-center analysis of computerized measurement systems for intima-media thickness measurement on common carotid artery longitudinal B-mode ultrasound scans". *Mendeley Data*, V1 (2022)
6. Meiburger, K.M., Zahnd, G., Faita, F., Loizou, C.P., Carvalho, C., Steinman, D.A., Gibello, L., Bruno, R.M., Marzola, F., Clarenbach, R., Francesconi, M., Nicolaides, A.N., Campilho, A., Ghotbi, R., Kyriacou, E., Navab, N., Griffin, M., Panayiotou, A.G., Gherardini, R., Varetto, G., Bianchini, E., Pattichis, C.S., Ghiadoni, L., Rouco, J., Molinari, F.: Carotid ultrasound boundary study (CUBS): An open multicenter analysis of computerized intima-media thickness measurement systems and their clinical impact. *Ultrasound in Medicine & Biology* 47(8), 2442–2455 (2021)

7. Meiburger, K.M., Zahnd, G., Faita, F., Loizou, C.P., Carvalho, C., Steinman, D.A., Gibello, L., Bruno, R.M., Marzola, F., Clarenbach, R., Francesconi, M., Nicolaides, A.N., Campilho, A., Ghotbi, R., Kyriacou, E., Navab, N., Griffin, M., Panayiotou, A.G., Gherardini, R., Varetto, G., Bianchini, E., Pattichis, C.S., Ghadoni, L., Rouco, J., Molinari, F.: DATASET for "Carotid Ultrasound Boundary Study (CUBS): an open multi-center analysis of computerized intima-media thickness measurement systems and their clinical impact". Mendeley Data, V1 (2021)
8. Menchón-Lara, R., Sancho-Gómez, J., Bueno-Crespo, A.: Early-stage atherosclerosis detection using deep learning over carotid ultrasound images. *Applied Soft Computing* 49, 616–628 (2016)
9. Meshram, N., Mitchell, C., Wilbrand, S., Dempsey, R., Varghese, T.: Deep learning for carotid plaque segmentation using a dilated U-net architecture. *Ultrasonic Imaging* 42(4-5), 221–230 (2020)
10. Molinari, F., Zeng, G., Suri, J.S.: A state of the art review on intima-media thickness (IMT) measurement and wall segmentation techniques for carotid ultrasound. *Computer Methods and Programs in Biomedicine* 100, 201–221 (2010)
11. Qian, C., Su, E., Yang, X.: Segmentation of the common carotid intima-media complex in ultrasound images using 2-D continuous max-flow and stacked sparse autoencoder. *Ultrasound in Medicine & Biology* 46(11), 3104–3124 (2020)
12. Ronneberger, O., Brox: U-net: Convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W., Frangi, A. (eds.) *International Conference on Medical Image Computing and Computer Assisted Intervention – MICCAI*. vol. LNCS 9351, pp. 234–241. Springer, Cham (2015)
13. Wang, K., Pu, Y., Zhang, Y., Wang, P.: Fully automatic measurement of intima-media thickness in ultrasound images of the common carotid artery based on improved Otsu's method and adaptive wind driven optimization. *Ultrasonic Imaging* 42(6), 245–60 (2020)
14. Zahnd, G., Kapellas, K., van Hattem, M., van Dijk, A., Sérusclat, A., Moulin, P., van der Lugt, A., Skilton, M., Orkisz, M.: A fully-automatic method to segment the carotid artery layers in ultrasound imaging: Application to quantify the compression-decompression pattern of the intima-media complex during the cardiac cycle. *Ultrasound in Medicine & Biology* 43(1), 239–257 (2017)
15. Zahnd, G., Orkisz, M., Sérusclat, A., Moulin, P., Vray, D.: Simultaneous extraction of carotid artery intima-media interfaces in ultrasound images: assessment of wall thickness temporal variation during the cardiac cycle. *International Journal of Computer Assisted Radiology and Surgery* 9(4), 645–658 (2014)