MOUSE ARTERIAL WALL IMAGING AND ANALYSIS FROM SYNCHROTRON X-RAY MICROTOMOGRAPHY

Xiaowen Liang[1,2], *Aïcha Ben Zemzem*[2], *Sébastien Almagro*[2], *Jean-Charles Boisson*[3], *Luiz-Angelo Steffenel*[3], *Timm Weitkamp*[4], *Laurent Debelle*[2], *Nicolas Passat*[1]

[1] Université de Reims Champagne Ardenne, CReSTIC EA 3804, 51100 Reims, France
[2] Université de Reims Champagne Ardenne, CNRS, MEDyC UMR 7369, 51100 Reims, France
[3] Université de Reims Champagne Ardenne, LICIIS / LRC CEA DIGIT, 51100 Reims, France
[4] Synchrotron SOLEIL, 91192 Gif-sur-Yvette, France

ABSTRACT

Synchrotron X-ray microtomography (μCT) gives access to images with a micrometric resolution. In the context of vascular imaging, this allows the study of structural properties of arterial walls, even for small animals such as the mouse. However, the images available with μCT are non-usual, and there is no method specifically designed for their processing and analysis. This article describes a first pipeline dedicated to the segmentation of μCT images of mice aorta. This pipeline builds upon conventional image processing paradigms and more recent deep learning approaches, and tackles the issue of multiscale analysis of huge-sized, high-resolution data. It provides promising results, assessed by comparison with manual annotation of sampled data. This methodological framework is a step forwards to a finer analysis of the internal structure of the aortic walls, especially for understanding the consequences of ageing and/or disease (e.g. diabetes) on the vessels architecture.

***Index Terms*—** Synchrotron, X-ray phase-contrast imaging, segmentation, texture analysis, deep learning, mice aorta, diabetes.

1. INTRODUCTION

Vascular ageing is characterized by profound modifications of the large elastic arteries (mainly aorta), eventually leading to dissection, aneurysm formation, atheroma plaque deposition, fragmentation and stroke [1]. Age-related diseases can be precipitated by metabolic disorders, e.g. diabetes or chronic kidney diseases [2]. Vascular diseases involve morphological changes of the arterial extracellular matrix, particularly remodelling of the elastic lamellae found in the wall. Under normal circumstances, these elastic structures dampen the pulse wave due to the systolic output and thereby warrant a constant blood pressure. When a vascular disease is declared, this functioning is compromised with potentially fatal outcomes. Our current understanding of the time course of vascular diseases points out that they undergo an asymptomatic phase where the disease grows silently and unnoticed until the symptoms appear. To date, the initial and discrete alterations affecting the elastic lamellae and forecasting the diseased state are poorly understood.

These facts motivate the accurate investigation of the internal architecture of the vascular walls of the arteries, especially the aorta. Such studies require the availability of adequate images but also computational tools for their analysis.

The analysis of vascular structures from angiographic images has been intensively investigated for more than 25 years, mainly from conventional imaging modalities (magnetic resonance imaging, X-ray computed tomography...) with a focus on the human and on the only lumen extraction [3], with a recent trend to consider deep learning approaches [4]. Regarding the aorta, most efforts were geared towards segmentation of the lumen and arterial wall in the specific case of aortic dissection [5]. The aortic wall structure is then investigated at a mesoscopic scale.

In order to explore the internal architecture of the arterial wall and especially the aorta, it is mandatory to rely on observations at a microscopic scale. This requires high-resolution imaging of tissues *ex vivo*, which leads to rely on preclinical models. Mice are often considered, allowing relevant ageing and/or disease modelling. However, the usual imaging modalities present limitations. Two-photon and three-photon microscopy [6] are too limited in tissue thickness, whereas the resolution of optical coherence tomography, echography and magnetic resonance imaging varies from tens of micrometers to millimeters, which is insufficient.

By contrast, the synchrotron technology enables to design acquisition protocols that fulfil the following requirements: large field of observation, high resolution and 3D view. More precisely, synchrotron X-ray microtomography (μCT) with in-line phase contrast allows visualizing biological soft tissues with (sub-)micrometric resolution, several centimeters of tissue thickness and millimeter field of view [7, 8].

The counterpart of these good properties of μCT are manifold: the data are huge (many GB per image) and they contain a wide range of details at different scales, which induces semantic noise in addition to acquisition noise that may disturb the analysis. Consequently, μCT images are difficult to analyze, and there is no available state-of-the-art method currently dedicated to processing them.

In this article, we consider μCT images of mouse aorta, which visualize, among many other structures, the elastic lamellae that compose specific features inside the internal architecture of the medial and intimal parts of the aorta wall [9]. Recent works have emphasized important properties of these elastic lamellae [10]. An accurate extraction of these structures is then of high interest.

We provide a methodological pipeline dedicated to the analysis of μCT images of mouse aorta. It is mainly composed of four steps: (1) a preprocessing of the images, in order to remove the noise and artifacts that may disturb the subsequent steps; (2) a segmentation of

The authors acknowledge Synchrotron SOLEIL to have provided beamtime under project no. 20191028.

ANATOMIX is an Equipment of Excellence (EQUIPEX) funded by the Investments for the Future program of the French National Research Agency (ANR), project NanoimagesX (grant ANR-11-EQPX-0031).

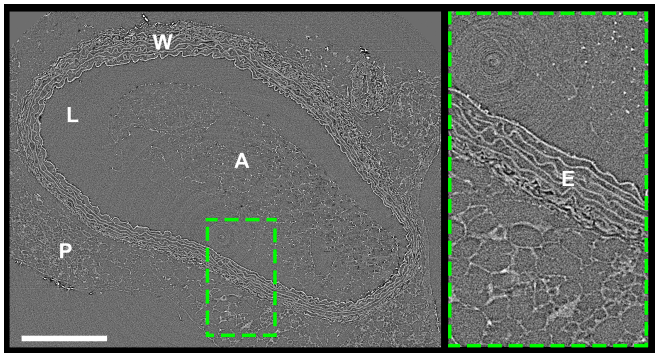

Fig. 1. Typical μCT image (2D slice). Raw acquisition of C57Bl6J mouse (control). The lumen (L) of the aorta contains agarose (A) and is surrounded by the arterial wall (W). The external part consists in a Perivascular tissue (P). Green frame magnifies the wall where elastic lamellae (E) are visible. Bar = 200 μm

the lumen, based on a mixed texture-based / mathematical morphology approach; (3) the computation of the normal field and distance maps from the lumen frontier, inside the arterial wall; and (4) a deep learning-based classification that aims to segment the arterial wall into two regions, including the area that contains the elastic lamellae. The methodological choices of this pipeline were mainly guided by the image properties and constraints, with the final purpose of automatically processing a large range of huge-sized images, with minimal requirements for parameters setting.

The remainder of this article is organized as follows. Sec. 2 provides information on the images. Sec. 3 describes the proposed methodological pipeline. Sec. 4 presents the experiments and the obtained results. Finally, Sec. 5 provides concluding remarks.

2. DATA

2.1. Samples

The mice[1], 6-month-old C57BL6J (control, $n = 4$) and 6-month-old *db/db* (diabetic, $n = 6$), were purchased from Charles River (Lyon, France). They were caged in temperature and humidity-controlled environment with a 12:12 hour light/dark cycle.

Mice aorta were collected after euthanasia; heart was injected by 2500 UI heparin. Heart and aorta were washed with 10 mL PBS to remove residual blood. The aorta was then prefixed by injection of 5 mL of 4% formalin. 6 mL of 1% low melting agarose were injected to keep the aorta open and to prevent collapse. Heart and aorta were both collected with surrounding tissues. The samples were fixed in 4% formalin for 24–48 hours, dehydrated and embedded in paraffin. The final samples were about 40-mm-long and 5-mm-wide paraffin rods containing the heart and aorta.

2.2. Images

Synchrotron X-ray μCT was performed on the ANATOMIX beamline [11] at the SOLEIL synchrotron (see Fig. 1). Samples were imaged with a polychromatic ("white") X-ray beam obtained from an undulator X-ray source set to a gap of 8.5 mm; the beam was filtered by a 0.6-mm-thick diamond plate and a 10-μm-thick layer of gold. The detector was an indirect lens-coupled system with a 20-μm-thick lutetium aluminum garnet single-crystal scintillator coupled to a CMOS-based scientific-grade camera with 2048^2 pixels via microscope optics ($10\times$ objective), resulting in an effective pixel size of 0.65 μm on the sample level. The distance between sample and scintillator was 22 mm. The exposure time for the camera was set to 100 ms per projection image. 1500 projections were taken over an angular range of 180°. The samples were positioned vertically with the heart in the lower part. Immediately after the acquisition, the imaged volume was reconstructed and checked. Tomographic reconstruction was performed using the standard processing pipeline at the beamline. The reconstructed volume stacks for each scan contained 2048^3 voxels of size $(0.65\ \mu\text{m})^3$, each represented by a 32-bit single-precision float value, i.e. a total of 32 GB per volume.

[1]Mouse procedures were realized in accordance with the Animal Subjects Committee of the Champagne-Ardenne Region (France).

3. METHODOLOGY

The following steps of the methodological pipeline are carried out slice by slice, thus dealing with 2D images. This choice is motivated by (1) reducing the computational cost and (2) considering the smooth evolution of the structures in the successive slices.

3.1. Preprocessing

The preprocessing includes denoising and normalization steps. The denoising steps aim to eliminate noise while highlighting useful details. The normalization step is needed due to imbalanced gray level distributions resulting from overexposure of specific regions (due to agarose) making other regions less contrasted than expected.

In order to avoid increasing the noise when improving the contrast during the further normalization, a first global filtering is applied, namely bilateral filtering [12]. This filter eliminates noise while preserving edges, which is mandatory due to the layered structure of the arterial wall. This filter is parameterized by σ_r (that controls signal proximity) and σ_d (that controls spatial proximity). The parameter σ_d is set to 50. As the global contrast of some images is often low, we consider a varying σ_r that depends on the standard deviation σ of the intensity distribution of the image I:

$$\sigma_r(I) = 2^{\sqrt{a\sigma(I)+b}} \tag{1}$$

where a and b were experimentally set to 0.5 and 10, respectively.

In a second step, we perform a normalization of the intensity distribution of each image I by adjusting its mean value to $\mu = 128$ and its standard deviation to $\sigma = 40$ if the initial one is less than 40.

A last, context-aware filtering is then applied, namely a non-local means filtering [13]. This filter has 3 parameters: the denoising strength h, the template window size k_1 and the search window size k_2. We iteratively apply 3 filterings with parameters (h, k_1, k_2) equal to $(20, 7, 21)$, $(20, 7, 21)$ and $(15, 3, 21)$, respectively.

3.2. Lumen segmentation

The first step of the segmentation consists of extracting the lumen, i.e. building the inner contour of the arterial wall. At this stage, the native 2048^2 images are processed at a lower resolution, in order to reduce the computational cost. Based on a Gaussian scale-space paradigm [14], we define a 1024^2 version I_{1024} of the image I, involving a 5×5 kernel built from the 1D kernel $\frac{1}{16}[1\ 4\ 6\ 4\ 1]$.

The lumen and the inner part of the vascular wall present specific textural properties. Then, we rely on a co-occurrence texture analysis [15] as basis for this first segmentation step. We compute

the co-occurrence matrix for each point x of the image I_{1024}, by considering a 9×9 window centered on x and the co-occurrences between the 8-neighbour pixels, with a set of values quantified into 8 bins, leading to a low-sized matrix $M(x) \in \mathcal{M}_{8\times8}(\mathbb{N})$. We build a feature map F of same dimension as I_{1024}, defined by:

$$F(x) = \max M(x) \tag{2}$$

The co-occurrence feature stored into this map tends to provide a high value located on the diagonal of the matrix in areas of constant signal (which is the case in the lumen region), while providing a much lower value due to the scattering of the co-occurrences in the non-constant, high-frequency signal regions (which is the case in the inner part of the arterial wall due to the presence of elastic lamellae).

A mean filtering is applied on F with a kernel size of 15^2, followed by a binary thresholding with a threshold value $\lambda = 150$. The step is finalized via mathematical morphology operations [16]. A dilation with a structuring element of size 15^2 is performed on the binary result, in order to eliminate high-frequency noise from the boundaries. An area opening (with an area threshold set to 20% of the size of I_{1024}) is finally carried out to discard the non-relevant connected regions and preserve only the regions corresponding to the lumen.

These operations provide a rough, underestimated boundary of the lumen, which is finalized via a region-growing [17] approach. The growing criterion is based on a global intensity difference (tolerance $\varepsilon = \pm 3$) between the candidate points and the initial boundary.

The boundary $\mathcal{B}$ of this region corresponds to the frontier between the lumen and the arterial wall.

3.3. Normal vector field and distance map computation

The purpose of this step is to prepare the data for the last, fine segmentation of the arterial wall into two classes of tissues, namely the *tunica media* (where the elastic lamellae are located) and *tunica adventitia* (the boundary with the external part). To this end, two pieces of information are required at each point x: the distance $\mathcal{D}(x)$ from the inner part of the arterial wall $\mathcal{B}$, and the vector $\vec{n}(x)$ that gives the normal orientation with respect to the arterial wall. These two fields $\mathcal{D}$ and $\vec{n}$ can be computed by carrying out a front propagation from the border $\mathcal{B}$.

The process is initialized as follows. For all pixels $x \in \mathcal{B}$, we set $\mathcal{D}(x) = 0$ and the normal $\vec{n}(x)$ is defined as [18]:

$$\vec{n}(x) = \frac{\sum_{y \in \overline{L} \cap B(x,\rho_0)} (x - y)}{|\sum_{y \in \overline{L} \cap B(x,\rho_0)} (x - y)|} \tag{3}$$

where $\overline{L}$ is the area outside the lumen and $B(x, \rho_0)$ is the ball of center x and radius ρ_0 (with $\rho_0 = 60$).

Each iteration of the process builds upon the information already computed on the region R to define the values on the pixels adjacent to R. For each such pixel x, the normal $\vec{n}(x)$ is defined as

$$\vec{n}(x) = \frac{\sum_{y \in R \cap B(x,\rho_1)} \vec{n}(y)}{|\sum_{y \in R \cap B(x,\rho_1)} \vec{n}(y)|} \tag{4}$$

where $\rho_1 = 5$, and the distance $\mathcal{D}(x)$ is defined as $1 + \mathcal{D}(x - \vec{n}(x))$, where $x - \vec{n}(x)$ is approximated on the discrete grid $\mathbb{Z}^2$.

3.4. Fine segmentation

The last step of this pipeline aims to discriminate *tunica media* and *tunica adventitia*. To reach that goal, we consider a Siamese neural network [19] in order to perform pixel-level classification. Our training process is inspired by [20], which involved homography in the context of image transformation.

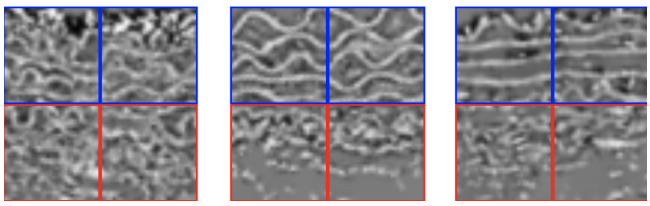

Fig. 2. Three examples of patches. All patches are oriented and positioned so that they are anatomically comparable. Each patch is subdivided into four subpatches that correspond to the two kinds of tissues: *tunica media* (red) and *tunica adventitia* (blue).

3.4.1. Training set

For training purposes, we did not have access to annotations for these two families of tissues. However, two priors were available: (1) the two kinds of tissues present distinct visual features; (2) the type of tissues is strongly correlated to the distance $\mathcal{D}$ from the inner contour of the arterial wall $\mathcal{B}$.

Consequently, although we cannot directly classify each pixel for designing training samples, we can design pairs of patches being "similar" or "dissimilar".

In our training process, samples are square image patches of size $2k \times 2k$ ($k = 32$). These patches are oriented to be vertically aligned with the normal vector field $\vec{n}$. The position of their barycenter is also defined with respect to the map $\mathcal{D}$, at a fixed value that maximizes the probability of being positioned at the frontier between the two tissue, with *tunica media* the closer to $\mathcal{B}$ and *tunica adventitia* the farther from $\mathcal{B}$.

Each patch P is then subdivided into 4 subpatches P_{ij} ($0 \leq i, j \leq 1$) of size $k \times k$. The two subpatches P_{0j} (resp. P_{1j}) on top (resp. bottom) are then assumed to belong to *tunica media* (resp. *tunica adventitia*). They provide two pairs of similar and two pairs of dissimilar ones. Examples of (sub)patches are given in Fig. 2.

3.4.2. Neural network structure

The structure of the neural network is the descriptor part defined in [20]. It is composed of a VGG-like encoder [21] and a decoder with 2 convolution layers; it has an output of 16 units. Convolution layers of the decoder and the first layer of the decoder (with kernel size of 3) are followed by ReLU non-linear activation and BatchNorm normalization, with sizes of 64-64-64-64-128-128-128-128; a max-pooling is added after each convolution layer. The output descriptors are normalized.

3.4.3. Loss function

The 4 subpatches P_{ij} obtained from one patch P enter the same network. This leads to 4 output (normalized) descriptors d_{ij} of length 16, combined as 6 pairs of (distinct) descriptors. The loss associated to a pair of subpatches is defined as:

$$\mathcal{L}(d_{ij}, d_{kl}) = \left\{ \begin{array}{l} \alpha.\delta_{ik}.\Delta^T.\Delta + \\ \beta.(1 - \delta_{ik}).[\max\{0, m_p - (\Delta^T.\Delta)^{1/2}\}]^2 \end{array} \right. \tag{5}$$

where $\Delta = d_{ij} - d_{kl}$, δ is the Kronecker symbol and $\alpha = 2$, $\beta = 1$, $m_p = 2$. This is a contrastive loss.

4. EXPERIMENTS

We run the neural network described in Sec. 3.4.2, with a batch size of 512 and Adam optimizer [22] with parameters of $l_r = 0.001$ and $\beta = (0.9, 0.999)$, the number of epochs was 5. The training was performed with a Quadro RTX 5000 GPU; the training runtime is about 164 seconds for 5 epochs.

4.1. Ground-truth

To our best knowledge, there is no state-of-the-art method dedicated to analyzing images such as those defined in Sec. 2. A comparative analysis of our method with respect to other approaches was then impossible. Consequently, we carried out an absolute, quantitative analysis based on the comparison of the computed results with manual annotations provided by a human expert. Four 2D slices (2048^2) were annotated to discriminate between the two regions: *tunica media* and *tunica adventitia* and the boundary $\mathcal{B}$.

4.2. Results

In order to carry out our experiments, we considered 49 3D (2048^3) images. From these images, 1457 2D slices (1024^2 after preprocessing) were processed for building patches required for the learning step (Sec. 3.4.1). From these data, we computed 47 118 image patches of size 64^2. This set was subdivided into two subsets: 80% of patches were used for training and 20% were used for validation.

4.2.1. Patch-based analysis (validation)

To start, we analyzed the ability of the method to correctly classify the two classes of patches within the 20% subset. Following the notations of Sec. 3.4.1, the classification for a given patch P was considered as correct if:

$$\|d_{00} - d_{01}\|_2 < \min\{\|d_{00} - d_{10}\|_2, \|d_{10} - d_{11}\|_2\} \qquad (6)$$

Under these hypotheses, the correct classification score was 99.14%.

4.2.2. Pixel-based analysis (testing)

In order to test the performance of the trained model in terms of segmentation at a pixel level, we firstly visualize segmentation results on patches. For a given patch P of size $2k \times 2k$ (with local coordinates $[0, 2k-1]^2$), we process the pixels x located in the central part of P of size $k \times k$. For each point x in this area, we compare the descriptor d_x at point $x = (x_1, x_2)$ with the two below / above descriptors d_- and d_+ at points $(x_1, \frac{k}{2})$ and $(x_1, \frac{3k}{2})$, that are assumed to correspond to the *tunica media* and *tunica adventitia*, respectively. If $\|d_x - d_+\|_2 < \|d_x - d_-\|_2$ then x is classified in the *tunica media*; otherwise it is classified in *tunica adventitia*. A partial result of this process is exemplified in Fig. 3.

Then we carried out the whole segmentation process on entire 2D slices. In particular, we processed 4 2D images endowed with a manual annotation of *tunica media* and *tunica adventitia*. To accelerate this process, we cut patches from points which distances are 0 in the distance map and rotated them according to corresponding normal vectors. In order to evaluate the segmentation of the lumen, we computed the average distance between the predicted and the annotated inner contour. This distance (in pixels) for the 4 annotated 2D images is 2.07, 1.94, 1.77 and 3.32. The mean error is then lower than 3 μm and acceptable with regard to the purpose of this segmentation that further defines the patches P.

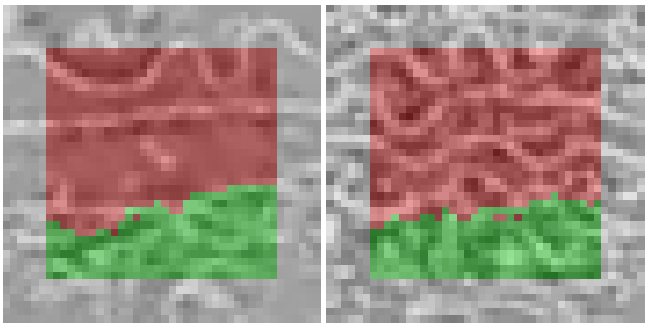

Fig. 3. Two samples of segmentation of *tunica media* (green) and *tunica adventitia* (red) within the arterial wall.

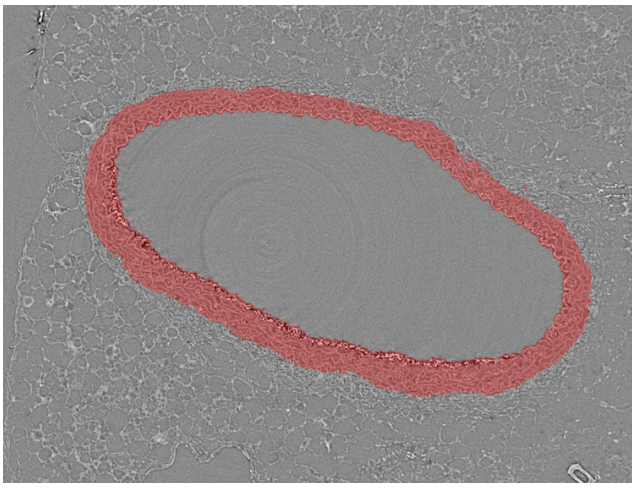

Fig. 4. Segmentation result: *tunica media* (in red).

In order to evaluate the segmentation of the *tunica media* vs. *tunica adventitia*, we relied on the Dice score [23]. This score for the 4 annotated 2D images is 0.926, 0.921, 0.933 and 0.917. A segmentation of *tunica media* is exemplified in Fig. 4.

5. CONCLUDING REMARKS

Chronic cardiovascular diseases are often discovered when it is too late, i.e. when irreversible damage has already occurred. Early identification of the very first stages of these diseases is therefore extremely important as they are the harbingers of its arrival. Synchrotron imaging has shown its power to detect these very localized events, but the complexity of the images requires dedicated developments to enable biologists to measure and interpret the different structures that make up these biological objects.

The proposed segmentation pipeline is a first step towards that goal. This pipeline mixes conventional image processing and deep learning paradigms, and provides promising results. The next steps of that work will consist, on the one hand, of improving the lumen segmentation step by reducing the parameter space to be handled and, on the other hand, of exploring more thoroughly the deep learning part of the process. In this context, these first results may be used to build a set of annotated data that would open the way to developing alternative deep learning architectures and more reliable testing procedures.

6. REFERENCES

[1] M. F. O'Rourke, M. E. Safar, and V. Dzau, "The cardiovascular continuum extended: Aging effects on the aorta and microvasculature," *Vascular Medicine*, vol. 15, no. 6, pp. 461–468, 2010.

[2] F. S. Vatner, J. Zhang, C. Vyzas, K. Mishra, R. M. Graham, and D. E. Vatner, "Vascular stiffness in aging and disease," *Frontiers in Physiology*, vol. 12, 2021.

[3] D Lesage, E. D. Angelini, I. Bloch, and G. Funka-Lea, "A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes," *Medical Image Analysis*, vol. 13, pp. 819–845, 2009.

[4] S. Moccia, E. De Momi, S. El Hadji, and L. S. Mattos, "Blood vessel segmentation algorithms - Review of methods, datasets and evaluation metrics," *Computer Methods and Programs in Biomedicine*, vol. 158, pp. 71–91, 2018.

[5] A. Pepe, J. Li, M. Rolf-Pissarczyk, C. Gsaxner, X. Chen, G. A. Holzapfel, and J. Egger, "Detection, segmentation, simulation and visualization of aortic dissections: A review," *Medical Image Analysis*, vol. 65, pp. 101773, 2020.

[6] X. Cheng, S. Sadegh, S. Zilpelwar, A. Devor, L. Tian, and D. A. Boas, "Comparing the fundamental imaging depth limit of two-photon, three-photon, and non-degenerate two-photon microscopy," *Optics Letters*, vol. 45, pp. 2934–2937, 2020.

[7] J. López-Guimet, L. Peña Pérez, R.S. Bradley, P. García-Canadilla, C. Disney, H. Geng, A.J. Bodey, P.J. Withers, B. Bijnens, M.J. Sherratt, and G. Egea, "MicroCT imaging reveals differential 3D micro-scale remodelling of the murine aorta in ageing and Marfan syndrome," *Theranostics*, vol. 8, pp. 6038–6052, 2018.

[8] G. Logghe, B. Trachet, L. Aslanidou, P. Villaneuva-Perez, J. De Backer, N Stergiopulos, M. Stampanoni, H. Aoki, and P. Segers, "Propagation-based phase-contrast synchrotron imaging of aortic dissection in mice: from individual elastic lamella to 3D analysis," *Scientific Reports*, vol. 8, pp. 2223, 2018.

[9] A. Ben Zemzem, A. Genevaux, A. Wahart, A. J. Bodey, S. Blaise, B. Romier-Crouzet, J. Jonquet, C. Bour, R. Cogranne, P. Beauseroy, M. Dauchez, M. J. Sherratt, L. Debelle, and S. Almagro, "X-ray microtomography reveals a lattice-like network within aortic elastic lamellae," *The FASEB Journal*, vol. 35, pp. e21844, 2021.

[10] B. Trachet, Ferraro M., G. Lovric, L. Aslanidou, G. Logghe, P. Segers, and N. Stergiopulos, "Synchrotron-based visualization and segmentation of elastic lamellae in the mouse carotid artery during quasi-static pressure inflation," *Journal of the Royal Society Interface*, vol. 16, pp. 20190179, 2019.

[11] T. Weitkamp, M. Scheel, J.L. Giorgetta, V. Joyet, V. Le Roux, G. Cauchon, T. Moreno, F. Polack, A. Thompson, and J.P. Samama, "The tomography beamline ANATOMIX at Synchrotron SOLEIL," *Journal of Physics: Conference Series*, vol. 849, pp. 012037, 2017.

[12] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *IEEE International Conference on Computer Vision (ICCV), Proceedings*, 1998, pp. 839–846.

[13] A. Buades, B. Coll, and J.-M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling and Simulation*, vol. 4, no. 2, pp. 490–530, 2005.

[14] T. Lindeberg, *Scale-space theory in computer vision*, vol. 256, Springer Science & Business Media, 2013.

[15] R. M. Haralick, K. Shanmugam, and I. Dinstein, "Textural features for image classification," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 3, pp. 610–621, 1973.

[16] L. Najman and H. Talbot, Eds., *Mathematical Morphology: From Theory to Applications*, ISTE/J. Wiley & Sons, 2010.

[17] R. Adams and L. Bischof, "Seeded region growing," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, pp. 641–647, 1994.

[18] G. Thürmer and C. A. Wüthrich, "Normal computation for discrete surfaces in 3D space," *Computer Graphics Forum*, vol. 16, pp. 15–26, 1997.

[19] G. Koch, R. Zemel, and R. Salakhutdinov, "Siamese neural networks for one-shot image recognition," in *ICML Deep Learning Workshop, Proceedings*, 2015.

[20] D. DeTone, T. Malisiewicz, and A. Rabinovich, "Superpoint: Self-supervised interest point detection and description," in *CVPR Workshops, Proceedings*, 2018, pp. 224–236.

[21] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *International Conference on Learning Representations (ICLR), Proceedings*, 2015.

[22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR), Proceedings*, 2015.

[23] L.R. Dice, "Measures of the amount of ecologic association between species," *Ecology*, vol. 26, pp. 297–302, 1945.