

Received 24 July 2023, accepted 9 September 2023, date of publication 15 September 2023, date of current version 20 September 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3315734

APPLIED RESEARCH

An HPC Pipeline for Calcium Quantification of Aortic Root From Contrast-Enhanced CCT Scans

GIADA MINGHINI¹, ARMANDO UGO CAVALLO², ANDREA MIOLA¹, VALENTINA SISINI¹, ENRICO CALORE³, FRANCESCA FORTINI⁴, RINO MICHELONI⁵, (Senior Member, IEEE), PAOLA RIZZO^{1,4}, SEBASTIANO FABIO SCHIFANO^{1,3}, FRANCESCO VIECELI DALLA SEGA⁴, AND CRISTIAN ZAMBELLI¹, (Member, IEEE)

¹Università degli Studi di Ferrara, 44121 Ferrara, Italy

²Istituto Dermopatico dell'Immacolata (IDI) IRCCS, 00167 Rome, Italy

³Istituto Nazionale di Fisica Nucleare (INFN) Ferrara, 44122 Ferrara, Italy

⁴Maria Cecilia Hospital, GVM Care and Research, 48033 Cotignola, Italy

⁵Avaneidi srl, 21047 Saronno, Italy

Corresponding author: Sebastiano Fabio Schifano (sebastiano.fabio.schifano@unife.it)

This work was supported in part by the Italian National Institute for Nuclear Physics (INFN); in part by the “Fondo per l’incentivo della Ricerca 2021” (FIR 2021) Program of the University of Ferrara; and in part by the Spoke 1 “FutureHPC and BigData” of the Italian Research Center on High-Performance Computing, Big Data and Quantum Computing (ICSC) funded by Ministero dell’Università e della Ricerca (MUR) Missione 4—Next Generation EU (NGEU).

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the “Comitato Etico Lazio 1” under Application No. 18/CE Lazio 1, in January 2021, and performed in line with the Declaration of Helsinki.

ABSTRACT Precise assessment of calcification lesions in the Aortic Root (AR) is relevant for the success of the Transcatheter Aortic Valve Implantation (TAVI) procedure. To this end, the radiologists analyze the Cardiac Computed Tomography (CCT) scans of patients, and detect the position and extent of the calcium deposits. In this contribution, we develop a computationally efficient High-Performance Computing (HPC) system to detect, segment, and quantify volumes of calcium in contrast-enhanced CCTs, embedding in a three-step pipeline two 3D Convolutional Neural Networks (CNN) and a threshold adaptive filter. The first step crops the images to a bounding-box around the AR keeping the original resolution, the second builds the segmentation, and the third detects and measures the volume of the calcium lesions. Our system is trained on high-resolution contrast-CCTs routinely planned for the TAVI manually annotated by expert radiologists, and evaluated on a test-set of patients with different levels of calcifications. The accuracy achieved in segmenting the AR is approximately 92% for the test-set, while the average difference of calcium lesion volumes with respect to the radiologists measurements is about 0.49 mm³. Running on a 4X NVIDIA-V100 and an 8X NVIDIA-A100 GPU systems, we achieve a remarkable inference throughput of 17 and 70 CCT/sec respectively, and a linear scaling of computing performance. Our contribution provides an HPC system suitable for hospital premises installation and is able to aid radiologists in assessing the calcification level of patients undergoing the TAVI, making this process automated, fast and more reliable.

INDEX TERMS Aortic root segmentation, calcifications assessment, contrast-enhanced CCT scans, GPU performance analysis.

I. INTRODUCTION

The Aortic Stenosis (AS) is the most frequent valvulopathy treated in the western world, especially in the elderly

The associate editor coordinating the review of this manuscript and approving it for publication was Ravibabu Mulaveesala¹.

population, with a prevalence of 2-9% in patients older than 65 [1]. Traditional treatment of severe AS is the Aortic Valve Replacement (AVR) by open-heart surgery, accounting for 60-70% of valve surgeries performed in the elderly [2]. However, at least 30% of patients are not advised for AVR due to high surgical risks related to advanced age or the presence

of various comorbidities [3], making the Transcatheter Aortic Valve Implantation (TAVI) more appropriate for them as it is less invasive [4]. The TAVI must be carefully planned to avoid issues such as the paravalvular regurgitation associated with high mortality probability [5], [6], occurring when the bio-prosthesis does not adhere properly to the aortic annulus due to multiple factors [7] among them the presence, position, and extension of calcification lesions within the aortic root.

A reliable method to assess the calcification level of the aortic root is the analysis of contrast-enhanced CCTs [8], [9]. To do this, delimitation of the aortic root and calcification lesions is routinely performed manually by expert radiologists on post-contrast CCT scans, resulting in a slow and time-consuming process prone to operator-dependent errors.

Making this process automatic is not trivial, and several issues have to be faced. In fact, since calcium is also contained in other internal organs, especially in bones, automatic analysis of CCT scans requires conceptually two steps, namely identification of the aortic root, and then of the calcium lesions within it. Moreover, in patients undergoing the TAVI the aortic root might be severely calcified, making it difficult to discover the bounds. For this reason, contrast-enhanced CCTs are usually performed to highlight the shapes of the tissues and other organ structures. However, using a contrast medium hampers the identification of the calcium lesions, since the luminance of voxels with the contrast is similar to that with low and medium levels of calcium. Finally, while for non-contrast CCTs standard methods have been already developed, like the Agatston score commonly used to evaluate the coronary artery calcifications [10], for contrast-enhanced no standard method has been yet implemented [7].

In this work, our aim is to develop a pipeline to automate and boost the segmentation of the aortic root, and make an assessment of calcium lesions within it. Automatic aortic root segmentation has several applications, among them the evaluation of aortic valve degeneration, and automatic aortic root geometrical measurements such as aortic valve or sinotubular junction dimensions. Here we focus on the quantification of the calcium, post-processing the segmentations using an adaptive filter. We do not aim to define any standard model for calcium assessment in contrast-enhanced CCT scans, which is out of the scope of our work, but we show how measures done by expert radiologists can be replicated by an automatic system.

In detail, our contribution is threefold: *i)* we design and implement a pipeline to make the segmentation of the aortic root in contrast-enhanced CCT scans, using two 3D U-Net CNNs running on HPC multi GPU systems; *ii)* we evaluate the performance of the CNNs in terms of segmentation accuracy, computing throughput, and scalability; *iii)* we validate the accuracy of the segmentations making automatic assessment of the aortic-root calcification level using a post-processing threshold adaptive filter, and comparing the results with that done by expert radiologists on a set of patients undergoing the TAVI. We underline that the success

of the latter point strongly depends on the precision of the aortic root segmentation. Moreover, we show how small-sized multi GPU-based HPC systems suitable to be hosted on hospital premises can be used to routinely support the work of radiologists in assessing the aortic root calcification level quickly and reliably. Besides this, the results achieved in this work in terms of throughput candidates our approach to be devised in retrospective studies characterized by a large amount of CCT images, and also in application scenarios where frequent scheduled CNN model retraining is foreseen to include new clinical cases and evidences.

The remainder of the paper is organized as follows: in the section II we make an overview of the related works and provide motivations for our work; in the section III we describe our dataset and how the image slices of CCTs have been manually annotated; in the section IV we give details about the methodologies used to make the segmentations through CNNs, and we describe the development of the adaptive threshold filter to post-process them to identify the calcium lesions; in the section V we present the results both in terms of accuracy and computational performance, and finally in section VI we draw our conclusions.

II. RELATED WORKS AND MOTIVATION

In the last decade, Deep Learning (DL) methods applied to diverse imaging techniques such as ultrasounds [11], Magnetic Resonance Imaging (MRI) [12], and CT [13], became the most widely used approach for cardiac image segmentation. In particular, the latter imaging technique is preferred since it generally leads to better segmentation accuracy ascribed to a higher image quality [14]. A revision of the literature [13] exposed four main branches where DL methods like Fully Connected Networks (FCN), CNNs, and more articulated models are devised for CCT imaging: *i)* cardiac substructure segmentation [15], [16], [17]; *ii)* coronary artery segmentation [18], [19]; *iii)* aortic root segmentation [20], [21], [22]; *iv)* calcium and plaque segmentation in coronaries and aortic valve [23], [24], [25]. Despite the confidence in the potentialities achievable in medical imaging by adopting DL-assisted segmentation procedures, there are still several challenges to address to improve the accuracy (even with image pre- and post-processing) and performance (i.e., images processed per second).

A. SEGMENTING THE AORTIC ROOT

The first goal of our work is to use DL to provide segmentation of the aortic root structure, which falls in the scope of *i)* and eventually *iv)*. Previous attempts on the same anatomical region of the heart relied on a two-step segmentation pipeline of neural networks [26], [27]. The first step consists of the extraction of the ROI to identify the heart and then feeding the second pipeline stage with images classified with a CNN [27]. Refined models use a localization network (Spatial Configuration Net) producing a coarse detection of the aortic landmarks and then applying

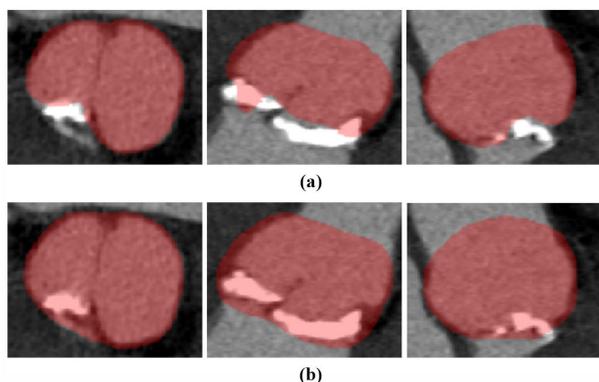


FIGURE 1. From left to right, Axial, Coronal, and Sagittal views of the segmentation of the aortic root. The images shown in (a) are extracted without applying the pre-processing filter for removing artifacts (see text), and evidence an incorrect anatomical structure segmentation since the areas of the aortic root under the white spots are not included in the segmentation; images in (b) show the correct segmentation achieved by applying the pre-processing filter.

a 3D CNN like UNet for segmentation [26]. A similar approach, although applied to ultrasound images, has been proposed in [28].

Most of the studies performed in the literature deal with non-contrast CCT images and consider patients affected by a low, yet null, level of calcification of the aortic valve. The works in [20] and [29] consider contrast-enhanced 2D and 3D CCT images on patients eventually undergoing the TAVI procedure (with a moderate/high level of calcification of the aortic valve), that is the scenario we are targeting in this work. Concerning the use of specific CNN models for, in [27] the authors have presented a CNN-based approach for detecting the aortic root in contrast-enhanced CTs, measured the aortic annulus diameter, and select the valvular prosthesis before TAVI. Meanwhile, in [20] authors have exposed an accurate and fast method to choose the TAVI device size by measuring the aortic annulus perimeter and area from manually segmented annulus planes using two U-Net models.

However, using standard DL models on our image dataset and in general with contrast-enhanced CCT would lead in some cases to the situation shown in Fig. 1a. Large calcified areas are not correctly picked up in the aortic root segmentation, thus leading to an incorrect evaluation of the anatomical structure and neglecting a large part of the calcium for quantification. We address this challenge by applying a filter pre-processing technique for artifact removal that will prove beneficial in the subsequent segmentation steps and in the overall quantification of the calcium on the aortic valve (see Fig. 1b).

B. QUANTIFYING THE CALCIUM ON THE AORTIC VALVE

The exact quantification of the calcium volume on the aortic valve is non-trivial when contrast-enhanced CCTs are considered. Indeed, there is a significant inherent inter- and intra-patient opacity (luminescence) variability of the dye [7] that hamper the development of fully-automated DL methods that exploit a constant luminescence threshold for calcium

extraction. This source of variation is ascribed to many factors such as the infusion rate of the contrast, the time elapsed between the injection and the start of the CCT exam, and the patient body habitus [7]. They cannot be ignored even following the guidelines on contrast usage for CCT exams [30]. Different approaches to quantify the calcium on the aortic valve have been proposed in literature, either using fixed or relative threshold values concerning the contrast medium luminescence (i.e., blood pool attenuation).

In [7], authors have measured the calcium volume score in CT angiography (CTA) to find the most accurate threshold to predict paravalvular regurgitation (PVR) after TAVI. They use two fixed thresholds (i.e., 650 and 850 HU) and four related to the luminal attenuation (LA) in the aortic annulus ($LA \times 1.25$, $LA \times 1.5$, $LA + 50$, $LA + 100$) for calcium detection, showing that with LA threshold cutoffs it is possible to achieve a discrimination accuracy (measured with the Area Under Curve method) between mild and severe PVR of 0.81.

In [31], authors have considered the contrast enhancement within the left ventricular outflow tract (LVOT), with a cut-off of 300 HU. The aortic calcium volume is measured with three fixed thresholds (450, 850, and probe+100 HU) and the results show that the best value (accuracy up to 70%) depends on the contrast in the LVOT. The developed method is not based on automated segmentation performed by DL methods, but rather applied to manual segmentations of the aortic valve. The authors point out that a standard reference for calcium assessment on contrast-enhanced CT series is currently lacking.

In [32], the authors applied commercial software to manually select the calcium elements in the CCT scan and obtain remarkable accuracy in scoring the aortic valve calcification. This work exposes a strong linear relationship between the blood pool attenuation and the threshold chosen in calcium volume quantification. Once again, this methodology is not included in a fully automated pipeline that performs this operation using an end-to-end approach with DL-assisted segmentation.

Current commercial solutions always require manual input in the segmentation and in the threshold choice, which introduces inter-operator variability [33]. Our work goes in the direction of providing an objective, fully repeatable, yet systematic framework with DL methods that is complemented with an adaptive filtering methodology to ease the calcium volume quantification when large luminescence variations between patients and within a single patient come into play.

C. UNCOMPROMISING ACCURACY WITH PERFORMANCE

The image processing time required by DL-assisted segmentation methods is usually an unexplored parameter in clinical studies. The accuracy of the segmentation appears to be the parameter of utmost importance in the development of fully automated methods and in some cases this is traded with the overall system performance. Despite some

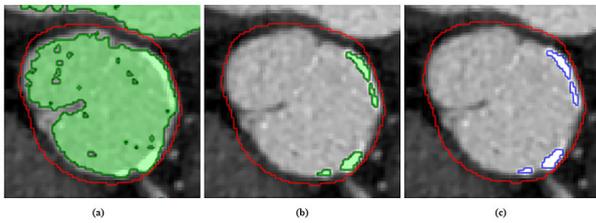


FIGURE 2. Steps to build segmentation of calcium lesions: (a) the 3D Threshold tool of MITK allows to highlight of all pixels with a luminance value equal to or greater than a threshold; (b) setting the appropriate threshold value, the operator highlights the calcium lesions within the aortic root; (c) application of the MITK Segmentation tool to create and export the ground-truth mask.

works addressing the achieved segmentation time and the required neural networks training time [27], [33], [34], [35], [36], [37], [38], none of them, to the best of our knowledge, neither investigate the scaling of such methods nor studied the role of the precision used in neural networks data representation in the context of HPC. Optimizing the segmentation performance becomes twofold relevant: *i*) can enable large volume studies considering many images from different sources; *ii*) provides fast re-training of the neural networks to improve the segmentation accuracy and calcium quantification even daily.

In Table 1, we report the characteristics of the studies that consider segmentation of the aorta for various applications using either 2D or 3D DL approaches based on CNN. We do not include in this comparison the possible pre- and post-processing times of the images, since it can be neglected to the overall inference time. Training time is nowadays in the range of several hours and even days for large datasets, whereas the inference time usually stays within one minute.

Our work investigates the impact on the training and inference segmentation performance of different small-sized HPC systems based on two generations of NVIDIA GPUs, namely the V100 and the most recent A100. We analyze the scaling in terms of GPUs adopted in the process and show how to achieve high accuracy and performance using different numerical precision in the DL data representation.

III. DATASET FEATURES AND IMAGING PROTOCOL

Our dataset comprises 27392 image slices corresponding to a total of 107 annotated gated contrast-enhanced CCTs acquired with a 256-slice scanner (Revolution CT, General Electric, Chicago, IL, USA), with prospective gating, setting the slice thickness at 0.625 mm, 120 kV, automatic mA, rotation time of 0.28 seconds, DFOV at 25 cm, and detector coverage at 160 mm. All patients have been subjected to administration of contrast material (Omnipaque 350 mg/mL, GE Healthcare, Chicago, IL, USA) at a rate of 5 mL/s followed by a saline chaser. The image acquisition was triggered after a threshold of 80 HU was reached in a region of interest placed in the left ventricle (bolus-tracking technique). Each CCT scan has 256 slices of 512×512 pixels, with a spacing of $0.488 \times 0.488 \times 0.625$ mm³, and the level of calcium regions ranges from 0 to 716 mm³, located

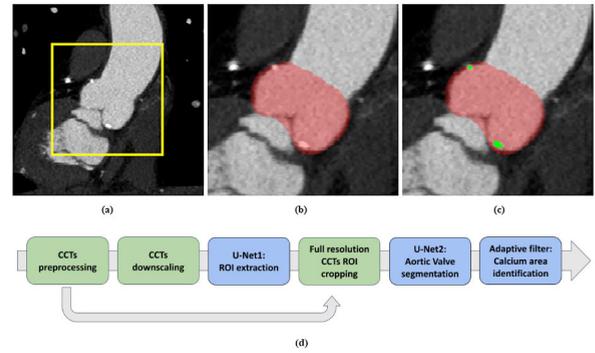


FIGURE 3. Processing steps to segment aortic root and calcifications: (a) The CCT image is scaled by a factor 2 along the coronal and sagittal axes and the ROI around the aortic root is extracted (yellow bounding box) by the U-Net1. (b) The aortic root is segmented with the U-Net2 within the ROI built by the U-Net1 on the image at full resolution. (c) The adaptive filter is applied on the aortic root mask to detect calcification areas (green regions). (d) The flowchart of our pipeline, where the green boxes identify the CCTs processing steps and the blue ones the U-Nets and filter phases.

at different places within the aortic root, including the valve cusps, the annulus, the sinotubular junction, and the coronaries junctions.

CCT analysis was performed by two radiologists in consensus with respectively 2 and 7 years of experience in cardiovascular imaging, using the *Medical Imaging Interaction Toolkit* (MITK), free open-source software for processing medical images [39]. Following the same approach used in [40], they first worked separately and, afterward, they reviewed the segmentations together. In case of disagreement, the consensus between the two is used as ground truth. The aortic root segmentation was built by making 3-4 contour planes along the axial, sagittal, and coronal anatomical axes, from the sinotubular junction to the aortic annulus. Then, a 3D interpolation has been applied to reconstruct the entire volume. To segment the calcification regions, the slices were inspected one by one, setting an appropriate threshold value to select the calcification voxels, as described in Fig. 2. Due to the differences in voxel luminance, the reference threshold is specific for each image [7]; for example, in our dataset, the mean reference threshold is 667 ± 127 HU. Once the calcium regions have been segmented, the corresponding volume has been measured using the Statistics tool of MITK.

IV. PIPELINE DESIGN FOR CALCIUM QUANTIFICATION

We propose a three-step pipeline as shown in Fig. 3. It is based on two CNN models to segment the aortic root, named U-Net1 and U-Net2, and on an adaptive threshold filter to detect the calcium regions. The entire process is complemented with image pre-processing, down-scaling, and cropping steps to increase the segmentation accuracy. The U-Net1 extracts the region of interest (ROI) around the aortic root (Fig. 3a) in low resolution obtained by scaling original images by a factor 2 along the coronal and sagittal axes, in order to fit images into the memory available on the GPUs used for running our models. This step gives, as a result, a bounding-box mask

TABLE 1. Summary of the accuracy and performance characteristics in state-of-the-art applications for aorta segmentation.

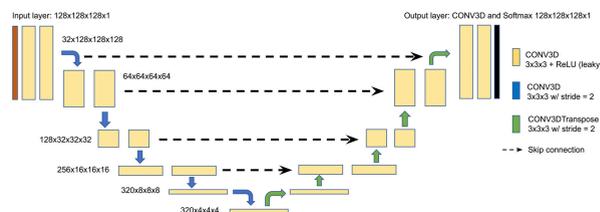
Study	Application	#Patients	DL method	Hardware	Accuracy	Training Time	Inference time
[34]	3D aortic segmentation	80	UNet2D pipeline	NVIDIA GeForce RTX 2080 Ti	0.93	18h	25s
[35]	Aortic measurement	3051	2 CNNs	24 GB GPU + 24-core CPU	0.96	19d	30s
[36]	Aortic aneurysm	72	Best of 3 CNNs	NVIDIA QUADRO P4000 (HPC)	0.88	-	15.23s
[37]	Aortic aneurysm	85	UNet2D pipeline	NVIDIA GeForce RTX 2080 Ti	0.89	21.5h	60s
[38]	General aortic disease	191	UNet3D	NVIDIA GeForce GTX 1080 Ti	0.95	-	60.5s
[33]	Aortic measurement	515	UNet3D	2×NVIDIA V100	0.94	-	180s
[27]	Aortic root analysis	126	CNN pipeline	4×NVIDIA Titan V	0.94 (F1)	20m	30s

that is used to crop the original image to a volume of $224 \times 256 \times 256$ voxels. Such a choice keeps the context around the aortic root as large as possible to help the second CNN detect it while fitting both the size of GPU memory and the requirements of the next CNN architecture. The U-Net2 (Fig.3-b) takes as input the cropped images at full resolution and returns a volume mask corresponding to the segmentation of the aortic root. The last step (Fig.3c) finds the calcification lesions within the region of the image corresponding to the aortic root segmented by the U-Net2. The whole flowchart is depicted in Fig. 3d, where the green boxes represent the preprocessing steps performed at the input of U-Net1 and U-Net2 and blue boxes the two U-Nets and the adaptive threshold filter.

A. AORTIC ROOT SEGMENTATION

U-Net1 and U-Net2 are CNNs based on the U-Net model, widely used in several studies for processing biomedical images [20], [41]. This neural network has been developed for image segmentation [42], with a U-shape architecture consisting of a contractive branch aiming to capture the image context, and an expansive branch to make precise localization of the segmentations. In particular, we have modified and adapted the 3D U-Net Medical CNN available at the NVIDIA DeepLearningExamples GitHub repository to our specific case [43]. This is an implementation of the model developed by [44] with improvements from [45], originally designed for the BRATS 2019 challenge to segment brain tumors.

The 3D U-Net architecture used in our work is configured with 5 down-sampling and 5 up-sampling steps as shown in Fig. 4. The steps in the down-sampling path consist of two convolutional blocks, each made of a convolution, a normalization, and an activation layer. The first block performs a stride down-sampling reducing all the spatial dimensions by a factor of two, while the latter consolidates the features learned at that depth level. In the up-sampling path, each step is made by a transposed convolution and two convolutional blocks. The first layer up-samples by a factor of two the spatial dimensions of the previous step, while the last two blocks keep as input the concatenation of the outputs of the transposed convolution and the skip connection. Skip connections are the outputs of the down-sampling blocks at the same depth level as the up-sampling block. They are particularly useful in image segmentation since they help to faithfully reconstruct the images using fine-grained details learned in the encoder part of the network. Lastly,

**FIGURE 4. Architecture of the 3D-UNet Medical CNN. It consists of 5-step contracting and expansive paths.**

the final output produced by the output block is processed by a softmax operation that is used for voxel classification. The model is implemented using TensorFlow [46] to leverage GPU acceleration and Horovod [47] for distributed learning on multi-GPU systems. Horovod is an open-source framework that provides tools to train deep learning models efficiently across multiple GPUs. By default, the model runs using 32-bit floating-point numerical precision, while mixed precision (AMP) computation, combining 16- and 32-bit operations, can be enabled to boost processing time both during the training and inference steps. Moreover, *Accelerated Linear Algebra* [48] (XLA) can be enabled to further boost the computation of mathematical operations. By default, the layers of the network coded using TensorFlow are processed independently. In contrast, using the XLA DL graph compiler, parts of the network are clustered into sub-graphs that can be optimized and compiled, providing performance benefits at the cost of some compilation overhead.

U-Net1 and U-Net2 are both trained to minimize the loss function $L = (1 - Dice) + CE$ and validated by measuring the Dice score on the inferences of the test set. The CE is the *Cross Entropy* computed as $-\sum_{i=1}^n t_i \log(p_i)$, where t_i is the truth label and p_i is the probability value for the i -th class output by the CNN for each voxel of the image; it gives a measure of the difference between the predicted probability distribution and the true distribution of the data. The Dice score is another commonly used metric in semantic segmentation returning a value between 0 and 1, being 1 when the segmentations match the ground truths.

At the input of both CNNs, the voxel intensities are clipped to 85% of the maximum luminance value to reduce the range and Z-score normalized, and stored as floating point numbers in the $[0 \dots 1]$ range, whereas labels are one-hot encoded for their later use in Dice or pixel-wise cross-entropy loss

TABLE 2. Accuracy of the filters in replicating the threshold reference value used by the radiologists to segment the calcium areas within the aortic root.

	F1	F2	F3	F4	F5	F6	F7
Coefficient R^2	-0.26	-0.11	-1.35	0.52	0.52	-2.57	0.89
Abs. mean error	94.90	107.96	162.83	76.77	72.42	215.27	29.58
Abs. error variance	96.81	66.25	86.42	33.33	42.01	71.90	25.81
Abs. error range	360.00	262.00	367.00	151.00	179.25	332.50	104.40

computation. Moreover, at the inputs of the U-Net2 a custom filter is applied to lower the luminance of the voxels in the ROI detected by U-Net1 and remove artifacts due to severe calcifications that may prevent the network from finding the bounds of the aortic root.

To increase the robustness of the network to different kinds of images, data augmentation techniques have been applied. In our case, at each training epoch, three different augmentations are randomly applied with a probability of 0.5 to the input images, namely random crop size, random brightness shifting, and horizontal flip. In this way, at each epoch the input samples are different from the previous epoch, helping in preventing over-fitting, and artificially increasing the number of samples processed by the CNNs.

B. DETECTION OF THE CALCIFICATION REGIONS

To segment and quantify calcifications, we have opted to filter the voxels within the aortic root with the luminance above a specific threshold. We developed different filters based on both fixed and variable thresholds computed with different algorithms, and tested them on a sample dataset of 7680 image slices of 30 CCTs randomly selected patients with different levels of calcifications. In Table 2, we assess for each filter the accuracy in replicating the reference threshold value (R.Thr.) set by the radiologists to segment the calcification regions. In detail, we report: *i*) the coefficient of determination R^2 to measure the goodness of the fit; this is a statistical indicator measuring how well the reference thresholds are replicated by the filter, based on the proportion of the total variation of the reference values explained by the model underlying the filter algorithm. The R^2 range is $[-\infty, 1]$, being 1 when the model perfectly fits the reference data; *ii*) the values of mean, standard deviation, and range of the absolute error of the thresholds predicted by the filter with respect to the R.Thr.

The filters F1, F2, and F3 use a fixed threshold set respectively to 600, 700, and 800 HU. In this case, the accuracy is low since the R^2 value is below zero, meaning that they are not able to replicate the reference thresholds. Additionally, the mean errors (94.90, 107.97, and 162.83 respectively) and the corresponding standard deviations (96.81, 66.25 and 86.41) are large.

To design a filter with a variable threshold, we measured the luminance L_m on a sample cube located within the aortic root of each patient. To this end, we first located the *geometric center* (GC) point of the aortic volume (see Fig.5a), and then constructed around it a cube of 20 voxels along each spatial dimension corresponding to approximately a volume

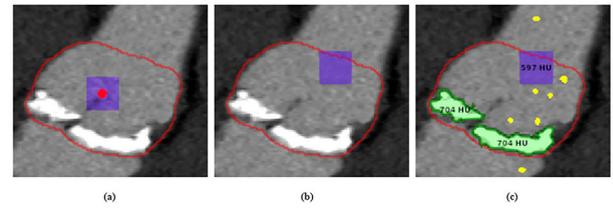


FIGURE 5. (a) Position of the geometric center (red point). (b) Position of the cube to sample the value of L_m at the sino-tubular junction. (c) Example where the value of $L_m = 597$ sampled within the cube is not high enough to avoid selection of voxels in the middle of the aortic root where calcifications are not possible (yellow regions); the luminance of calcification areas (green regions) is ≥ 704 HU.

of 10 mm^3 . The cube was then moved above the GC, near the sinotubular junction, to avoid contact with the valve leaves where calcification deposits may occur (see Fig.5b). In this way, the luminance of the voxels within the cube is that of blood and contrast medium. Furthermore, since the values of L_m and R.Thr. are strongly correlated (the Pearson correlation index measured with the Bivariate Correlation function is $r = 0.96$ with a P-value < 0.0001), we can expect improved results when using filters with variable thresholds computed as a function of L_m .

The threshold for filter F4 is set to the value of L_m , while the thresholds for filters F5 and F6 are respectively $L_m \times 1.25$ and $L_m \times 1.50$, as described in [7]. However, also in this case, none of these filters are able to accurately replicate the reference thresholds, since the R^2 values are very low (0.52, 0.52, and -2.57 , respectively), and the mean and standard deviation of the absolute errors are large. This is mainly due to the contrast medium that does not distribute uniformly within the aortic root, making some spots brighter than the maximum luminosity sampled within the cube. For example, as shown in Fig.5c, the value of L_m alone (597 HU in this case) is not high enough to avoid the selection of voxels in the middle of the aortic root (yellow spots with luminance greater than 597 HU), where we do not expect to have any calcification, and increasing the value by a fixed factor might bring to underestimate the calcium voxel count.

To further improve the accuracy, we have developed the filter F7 based on adaptive thresholds computed as $L_m + k(1000 - L_m)$. This formula is an implicit form of a linear regression where voxels within the region of the aortic root with a density above the upper bound of 1000 HU are assumed to be calcium. To this extent, the L_m needs to be increased by a corrective value that is lower as L_m approaches the upper limit of 1000 HU, otherwise, some areas of calcium are not identified. This allows the exclusion of the spots of contrast voxels within the aortic root. Applying this filter to our sample dataset, we have found that a value of $k = 0.2$ is enough to exclude the selection of spurious voxels. As reported in Table 2, the indicator values for the filter F7 are significantly better compared to the others, being the $R^2 = 0.89$ and the mean absolute error of 29.58 HU with a standard deviation of 25.81, making it the most accurate in predicting the reference thresholds.

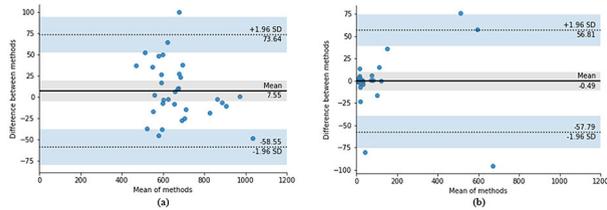


FIGURE 6. Bland-Altman difference plots for calcium thresholds in HU (a), and calcium volume in mm^3 (b). Limits of agreement are ± 1.96 times the standard deviation.

V. EXPERIMENTAL RESULTS

In this section, we discuss and assess the results achieved in terms of precision in segmenting the aortic root, quantification of the volume of calcium lesions, and computing performance.

A. AORTIC ROOT AND CALCIFICATIONS SEGMENTATION

Both U-Net1 and U-Net2 have been trained for 1250 epochs with a learning rate of 10^{-4} using the Adam optimizer. Using 18944 image slices of 74 CCTs for the train set and 8448 image slices of 33 CCTs for the test set, we have achieved for both CNNs a Dice score of ≈ 0.90 . To validate the results and to verify that no bias is present in the dataset, we also performed a K -fold cross-validation with $K = 5$, splitting the dataset into K random splits and training each model K times, where for each round four splits are used in turn for training and the last for testing.

Applying the filter F7 to the test set, the coefficient of determination R^2 between R.Thr. and the filter calculated threshold is 0.94, with a mean absolute error of 26.15 ± 21.81 HU and a range of 99 HU; regarding the volume, we get $R^2 = 0.97$, with a mean absolute error of 13.65 ± 25.35 mm^3 and a range of 95.46 mm^3 . In Fig. 6 we show two Bland-Altman plots; Fig. 6a refers to the agreement of the thresholds measured in HU for which the mean difference is 7.55 and the 95% confidence interval is $[-58.55, 73.64]$, while Fig. 6b to the calcium volume measured in mm^3 , where the mean error is -0.49 and the 95% confidence interval is $[-57.79, 56.81]$. In both cases, the agreement of the values predicted by F7 with the radiologists measures is accurate within the 95% of the confidence interval.

B. COMPUTING PERFORMANCE

We have run our application on two small-sized HPC systems, suitable to be installed in small data centers like those available today in several hospitals. The first system hosts 2X Intel Gold 6242 CPUs operating at 2.8 GHz, equipped with 394 GB of RAM and 4 NVIDIA V100-SXM2-32GB GPU accelerators. The latter is an NVIDIA DGX A100 based on dual AMD Rome 7742 128-core CPUs running at 2.25 GHz and 320 GB of memory, with 8 A100 GPUs. This is hosted in a standard 19 inch blade chassis of approximately 482 mm wide, 44 mm tall and 894 mm depth. The A100 architecture has introduced a novel math mode dedicated to AI training, namely the TensorFloat-32 (TF32). In this case, the numerical range is

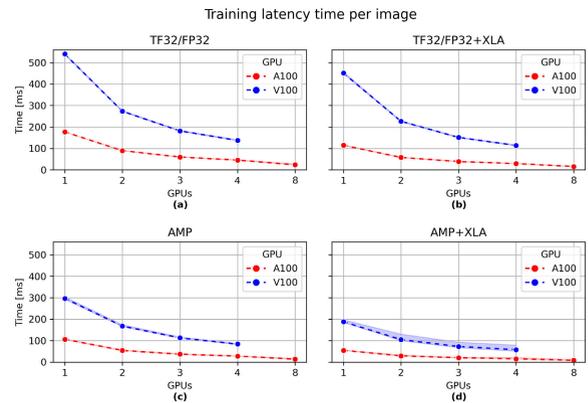


FIGURE 7. Train latency per CCT for the different models as a function of the number of GPUs.

the same as FP32 while the precision is that of FP16, resulting in a significant reduction in computation, memory, and memory bandwidth requirements without harmful impacts on prediction accuracy. This translates to an increase of operation throughput of up to $10\times$ or more compared to prior V100 generation of GPUs, and up to $5\times$ higher performance for DL workloads [49]. By default, when running on A100 GPUs, the TF32 numerical precision is used unless AMP or FP32 is explicitly selected. The DGX is hosted in a 19 inch box of approximately 482 mm width, 264 mm tall, and 897 mm depth.

In the following, we report the computing performance for the U-Net2 only without a lack of generality, since similar results have been measured also for the U-Net1.

In Fig. 7 we report for both systems the training latency per processed CCT (256 slice images) in units of milliseconds, achieved using up to four and eight GPUs respectively. We run the models using floating-point numerical precision, FP32 on the V100 and TF32 for the A100 system, and mixed precision (AMP) with and without XLA enabled. For the V100 system, using the FP32-based model (see Fig. 7a) we measure approximately 550 ms per CCT on one GPU. The time decreases using multiple GPUs, reaching nearly 150 ms per CCT with 4 GPUs, corresponding to an improvement of about 70%. Enabling XLA further improves the processing time by an additional 20% regardless of the number of GPUs (see Fig. 7b). By using the AMP precision, see Fig. 7c, the time decreases of $\approx 40\%$ compared to the FP32, and $\approx 30\%$ compared to FP32 with XLA, with an average execution-time of 306 ms using 1 GPU and 87 ms with 4 GPUs. Lastly, by combining AMP and XLA (see Fig. 7d), the processing time further decreases by $\approx 30\%$ reaching ≈ 60 ms per CCT using 4 GPUs, and running $9\times$ faster than the baseline (FP32-based model on 1 GPU). For the A100 system we have a similar behavior, achieving a lower training time. Using AMP with XLA enabled, we measure on average ≈ 55 ms per CCT on 1 GPU, while using 8 GPUs the training time decreases by 85% reaching approximately 8.65 ms per image.

Fig. 8 shows the throughput and the scalability achieved during the training process. As we see, the AMP model

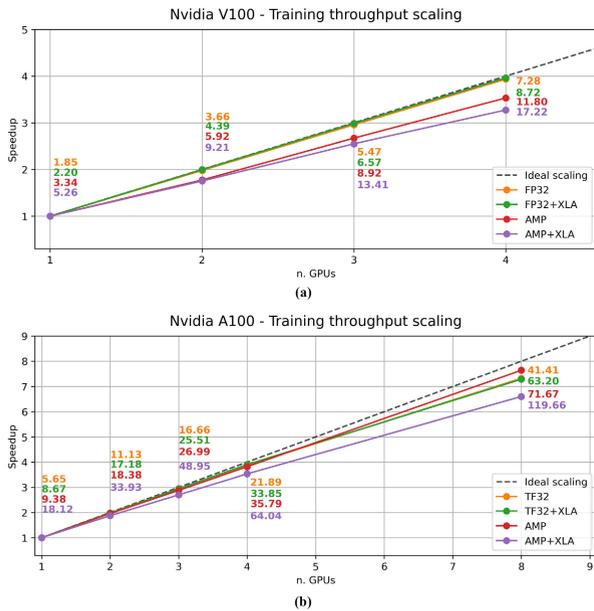


FIGURE 8. Average train throughput (a) and corresponding scaling (b) achieved with different precisions and optimizations as a function of the number of GPUs.

with XLA enabled performs significantly better than the others, reaching a remarkable throughput of ≈ 17 CCT/sec and ≈ 120 CCT/sec on the V100 and the A100 system, respectively. The scaling achieved is close to the ideal curve on both systems. In particular, on the V100 system, the FP32-based model scales better with and without the XLA enabled, while the AMP models perform somewhat worse due to the additional operations that each GPU needs to convert the computational graph to mixed precision. Additionally, the overheads for handling multi-GPU computation dominate the CCT processing time with a negative impact on the scaling. Similar behavior is observed for the A100 system, achieving a peak throughput of ≈ 120 CCT per second using all 8 GPUs.

In Fig.9 we report the inference throughput running on one V100 and one A100 GPU. As we see, on both systems, moving from TF32/FP32 to AMP the throughput slightly increases, giving the best performance using the AMP model with XLA enabled. Since, the inference can be easily parallelized running multiple instances of the trained model on different GPUs with different input images, running on all available GPUs on each system, we achieve ≈ 17 CCT per second on the V100 system, and ≈ 70 CCT per second on the A100 system.

To assess the impact of the use of AMP and XLA optimizations on the accuracy of the segmentation, we have run the K -fold cross-validation using different numerical precision, with and without XLA enabled. On the V100 system, using the FP32 the Dice score achieved is 0.915 ± 0.04 without XLA and 0.92 ± 0.02 with XLA enabled, while using the AMP we get 0.915 ± 0.04 and 0.926 ± 0.02 . This shows that moving from FP32 to an AMP-based model with and without XLA enabled we have no loss in terms of Dice

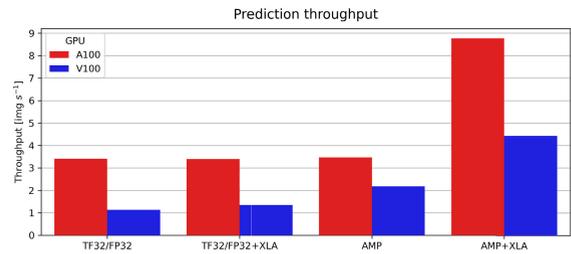


FIGURE 9. Throughput achieved for inference on one GPU with different precisions and optimizations enabled.

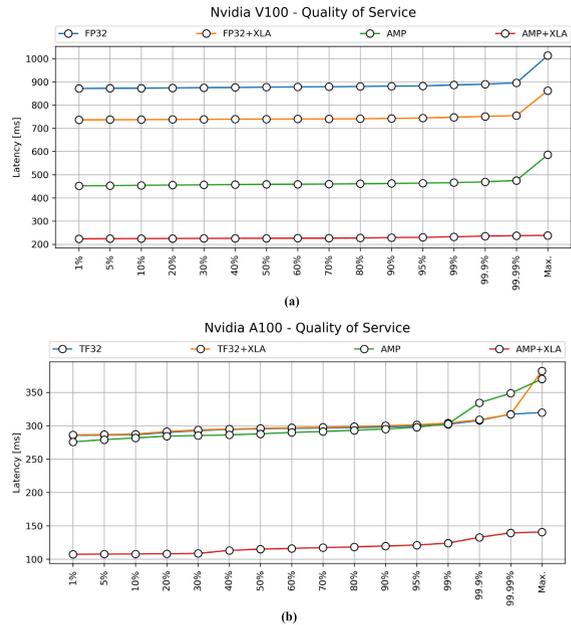


FIGURE 10. Quality of Service plots obtained running 10000 image inferences on the V100 (a) and A100 (b) systems.

score, but in terms of computing time, we run approximately $3 \times$ faster with XLA, and $2.5 \times$ without compared to the FP32 model. Similar behavior is measured on the A100 system.

Finally, in Fig.10 we report the quality of service (QoS) plots running 10000 inferences. On both systems we see that the time of inference is quite stable up to 99.99% of the inference latency distribution, exposing that the GPU inference time is largely predictable and the impact of potential distribution outliers is negligible.

VI. CONCLUSION

We have developed an HPC computing pipeline to automate the segmentation of the aortic root in contrast-enhanced CCTs using CNNs and GPU accelerators to quantify calcification volumes within it. The CNNs used achieve a Dice score of ≈ 0.90 , whereas the post-processing adaptive filter we have developed agrees with the radiologists measurements of calcification volumes with a $R^2 = 0.97$, and a mean error of ≈ 0.49 mm³. We have run our pipeline on two small footprint multi-GPU HPC systems, and in terms of computing performance the scaling as function of GPUs is quite ideal, achieving an aggregate throughput

of approximately 17 CCT/sec on the V100, and approximately 120 CCT/sec on the DGX A100 system. These results show that it is suitable to be hosted in hospital data centers to support radiologists in assessing the aortic root calcification level in routine clinical activities to plan the TAVI.

In addition, having a computational fast, and accurate system for processing CCT images can be beneficial for many reasons. For example, to make retrospective studies and find correlations between the calcification level of the aortic root and with success of the TAVI procedure, it is necessary to process hundreds or thousands of CCT scans. Also, it might open the possibility of making (pseudo) real-time assessments of the calcium burden volume by integrating the prediction system in the CT scanner. Fast training is also beneficial, for example, to re-train the models quickly on larger datasets including new cases in the train-set each time significant mispredictions are found. Finally, and this is a more general comment, reducing the time to solution is also beneficial in terms of energy to solution.

For future developments, we plan to validate our system on a larger dataset that includes CCTs from different scanners and use it to make retrospective studies to find correlations between aortic root calcification and TAVI procedure failures, as well as patients post-procedural prognosis.

DECLARATIONS

Informed consent was obtained from all the subjects involved in this study.

REFERENCES

- [1] P. Faggiano, F. Antonini-Canterin, F. Baldessin, R. Lorusso, A. D'Aloia, and L. D. Cas, "Epidemiology and cardiovascular risk factors of aortic stenosis," *Cardiovascular Ultrasound*, vol. 4, no. 1, pp. 5–27, Dec. 2006, doi: [10.1186/1476-7120-4-27](https://doi.org/10.1186/1476-7120-4-27).
- [2] R. R. Coeytaux, J. W. Williams, R. N. Gray, and A. Wang, "Percutaneous heart valve replacement for aortic stenosis: State of the evidence," *AIM*, vol. 153, no. 5, pp. 314–324, 2010, doi: [10.7326/0003-4819-153-5-201009070-00267](https://doi.org/10.7326/0003-4819-153-5-201009070-00267).
- [3] M. B. Leon, "Transcatheter aortic-valve implantation for aortic stenosis in patients who cannot undergo surgery," *New England J. Med.*, vol. 363, no. 17, pp. 1597–1607, Oct. 2010, doi: [10.1056/NEJMoa1008232](https://doi.org/10.1056/NEJMoa1008232).
- [4] C. R. Smith, "Transcatheter versus surgical aortic-valve replacement in high-risk patients," *New England J. Med.*, vol. 364, no. 23, pp. 2187–2198, Jun. 2011, doi: [10.1056/NEJMoa1103510](https://doi.org/10.1056/NEJMoa1103510).
- [5] K. Hayashida, T. Lefèvre, B. Chevalier, T. Hovasse, M. Romano, P. Garot, E. Bouvier, A. Farge, P. Donzeau-Gouge, B. Cormier, and M.-C. Morice, "Impact of post-procedural aortic regurgitation on mortality after transcatheter aortic valve implantation," *JACC, Cardiovascular Intervent.*, vol. 5, no. 12, pp. 1247–1256, Dec. 2012, doi: [10.1016/j.jcin.2012.09.003](https://doi.org/10.1016/j.jcin.2012.09.003).
- [6] G. Athappan, E. Patvardhan, E. M. Tuzcu, L. G. Svensson, P. A. Lemos, C. Fraccaro, G. Tarantini, J.-M. Sinning, G. Nickenig, D. Capodanno, C. Tamburino, A. Latib, A. Colombo, and S. R. Kapadia, "Incidence, predictors, and outcomes of aortic regurgitation after transcatheter aortic valve replacement," *J. Amer. College Cardiol.*, vol. 61, no. 15, pp. 1585–1595, Apr. 2013, doi: [10.1016/j.jacc.2013.01.047](https://doi.org/10.1016/j.jacc.2013.01.047).
- [7] N. Bettinger, O. K. Khalique, J. M. Krepp, N. B. Hamid, D. J. Bae, T. C. Pulerwitz, M. Liao, R. T. Hahn, T. P. Vahl, T. M. Nazif, I. George, M. B. Leon, A. J. Einstein, and S. K. Kodali, "Practical determination of aortic valve calcium volume score on contrast-enhanced computed tomography prior to transcatheter aortic valve replacement and impact on paravalvular regurgitation: Elucidating optimal threshold cutoffs," *J. Cardiovascular Comput. Tomogr.*, vol. 11, no. 4, pp. 302–308, Jul. 2017, doi: [10.1016/j.jcct.2017.04.009](https://doi.org/10.1016/j.jcct.2017.04.009).
- [8] M. Eberhard, M. Mastalerz, T. Frauenfelder, F. Tanner, F. Maisano, F. Nietlispach, B. Seifert, H. Alkadhhi, and T. Nguyen-Kim, "Quantification of aortic valve calcification on contrast-enhanced CT of patients prior to transcatheter aortic valve implantation," *EuroIntervention*, vol. 13, no. 8, pp. 921–927, Oct. 2017, doi: [10.4244/EIJ-D-17-00060](https://doi.org/10.4244/EIJ-D-17-00060).
- [9] M. K. Doris, W. Jenkins, P. Robson, T. Pawade, J. P. Andrews, R. Bing, T. Cartlidge, A. Shah, A. Pickering, M. C. Williams, Z. A. Fayad, A. White, E. J. van Beek, D. E. Newby, and M. R. Dweck, "Computed tomography aortic valve calcium scoring for the assessment of aortic stenosis progression," *Heart*, vol. 106, no. 24, pp. 1906–1913, Dec. 2020, doi: [10.1136/heartjnl-2020-317125](https://doi.org/10.1136/heartjnl-2020-317125).
- [10] D. Eilert and R. Goldenberg, "Fully automatic model-based calcium segmentation and scoring in coronary CT angiography," *Int. J. Comput. Assist. Radiol. Surg.*, vol. 9, no. 4, pp. 595–608, Jul. 2014, doi: [10.1007/s11548-013-0955-y](https://doi.org/10.1007/s11548-013-0955-y).
- [11] J. K. Dave, M. E. McDonald, P. Mehrotra, A. R. Kohut, J. R. Eisenbrey, and F. Forsberg, "Recent technological advancements in cardiac ultrasound imaging," *Ultrasonics*, vol. 84, pp. 329–340, Mar. 2018, doi: [10.1016/j.ultras.2017.11.013](https://doi.org/10.1016/j.ultras.2017.11.013).
- [12] A. S. Lundervold and A. Lundervold, "An overview of deep learning in medical imaging focusing on MRI," *Zeitschrift Für Medizinische Physik*, vol. 29, no. 2, pp. 102–127, May 2019, doi: [10.1016/j.zemedi.2018.11.002](https://doi.org/10.1016/j.zemedi.2018.11.002).
- [13] C. Chen, C. Qin, H. Qiu, G. Tarroni, J. Duan, W. Bai, and D. Rueckert, "Deep learning for cardiac image segmentation: A review," *Frontiers Cardiovascular Med.*, vol. 7, p. 25, Mar. 2020, doi: [10.3389/fcvm.2020.00025](https://doi.org/10.3389/fcvm.2020.00025).
- [14] X. Zhuang, "Evaluation of algorithms for multi-modality whole heart segmentation: An open-access grand challenge," *Med. Image Anal.*, vol. 58, Dec. 2019, Art. no. 101537, doi: [10.1016/j.media.2019.101537](https://doi.org/10.1016/j.media.2019.101537).
- [15] M. Zreik, T. Leiner, B. D. de Vos, R. W. van Hamersvelt, M. A. Viergever, and I. Išgum, "Automatic segmentation of the left ventricle in cardiac CT angiography using convolutional neural networks," in *Proc. IEEE 13th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2016, pp. 40–43, doi: [10.1109/ISBI.2016.7493206](https://doi.org/10.1109/ISBI.2016.7493206).
- [16] C. Wang and Ö. Smedby, "Automatic whole heart segmentation using deep learning and shape context," in *Statistical Atlases and Computational Models of the Heart. ACDC and MMWS Challenges*, M. Pop, M. Sermesant, P.-M. Jodoin, A. Lalonde, X. Zhuang, G. Yang, A. Young, and O. Bernard, Eds. Cham, Switzerland: Springer, 2018, pp. 242–249, doi: [10.1007/978-3-319-75541-0_26](https://doi.org/10.1007/978-3-319-75541-0_26).
- [17] C. Ye, W. Wang, S. Zhang, and K. Wang, "Multi-depth fusion network for whole-heart CT image segmentation," *IEEE Access*, vol. 7, pp. 23421–23429, 2019, doi: [10.1109/ACCESS.2019.2899635](https://doi.org/10.1109/ACCESS.2019.2899635).
- [18] M. C. H. Lee, K. Petersen, N. Pawlowski, B. Glocker, and M. Schaap, "TeTrIS: Template transformer networks for image segmentation with shape priors," *IEEE Trans. Med. Imag.*, vol. 38, no. 11, pp. 2596–2606, Nov. 2019, doi: [10.1109/TMI.2019.2905990](https://doi.org/10.1109/TMI.2019.2905990).
- [19] Y. Shen, Z. Fang, Y. Gao, N. Xiong, C. Zhong, and X. Tang, "Coronary arteries segmentation based on 3D FCN with attention gate and level set function," *IEEE Access*, vol. 7, pp. 42826–42835, 2019, doi: [10.1109/ACCESS.2019.2908039](https://doi.org/10.1109/ACCESS.2019.2908039).
- [20] P. Astudillo, P. Mortier, J. Bosmans, O. De Backer, P. de Jaegere, M. De Beule, and J. Dambre, "Enabling automated device size selection for transcatheter aortic valve implantation," *J. Interventional Cardiol.*, vol. 2019, pp. 1–7, Nov. 2019, doi: [10.1155/2019/3591314](https://doi.org/10.1155/2019/3591314).
- [21] F. Lalys, S. Esneault, M. Castro, L. Royer, P. Haïgron, V. Auffret, and J. Tomasi, "Automatic aortic root segmentation and anatomical landmarks detection for TAVI procedure planning," *Minimally Invasive Therapy Allied Technol.*, vol. 28, no. 3, pp. 157–164, May 2019, doi: [10.1080/13645706.2018.1488734](https://doi.org/10.1080/13645706.2018.1488734).
- [22] S. Saitta, F. Sturla, R. Gorla, O. A. Oliva, E. Votta, F. Bedogni, and A. Redaelli, "A CT-based deep learning system for automatic assessment of aortic root morphology for TAVI planning," *Comput. Biol. Med.*, vol. 163, 2023, doi: [10.1016/j.combiomed.2023.107147](https://doi.org/10.1016/j.combiomed.2023.107147).
- [23] S. Cowell, D. Newby, J. Burton, A. White, D. Northridge, N. Boon, and J. Reid, "Aortic valve calcification on computed tomography predicts the severity of aortic stenosis," *Clin. Radiol.*, vol. 58, no. 9, pp. 712–716, 2003, doi: [10.1016/S0009-9260\(03\)00184-3](https://doi.org/10.1016/S0009-9260(03)00184-3).
- [24] J. M. Wolterink, T. Leiner, B. D. de Vos, R. W. van Hamersvelt, M. A. Viergever, and I. Išgum, "Automatic coronary artery calcium scoring in cardiac CT angiography using paired convolutional neural networks," *Med. Image Anal.*, vol. 34, pp. 123–136, Dec. 2016, doi: [10.1016/j.media.2016.04.004](https://doi.org/10.1016/j.media.2016.04.004).

- [25] D. Eng, "Automated coronary calcium scoring using deep learning with multicenter external validation," *NPJ Digit. Med.*, vol. 4, no. 1, p. 88, Jun. 2021, doi: [10.1038/s41746-021-00460-1](https://doi.org/10.1038/s41746-021-00460-1).
- [26] G. Aoyama, L. Zhao, S. Zhao, X. Xue, Y. Zhong, H. Yamauchi, H. Tsukihara, E. Maeda, K. Ino, N. Tomii, S. Takagi, I. Sakuma, M. Ono, and T. Sakaguchi, "Automatic aortic valve cusps segmentation from CT images based on the cascading multiple deep neural networks," *J. Imag.*, vol. 8, no. 1, p. 11, Jan. 2022, doi: [10.3390/jimaging8010011](https://doi.org/10.3390/jimaging8010011).
- [27] N. Krüger, A. Meyer, L. Tautz, M. Hüllebrand, I. Wamala, M. Pullig, M. Kofler, J. Kempfert, S. Sündermann, V. Falk, and A. Hennemuth, "Cascaded neural network-based CT image processing for aortic root analysis," *Int. J. Comput. Assist. Radiol. Surg.*, vol. 17, no. 3, pp. 507–519, Mar. 2022, doi: [10.1007/s11548-021-02554-3](https://doi.org/10.1007/s11548-021-02554-3).
- [28] K. W. Lai, M. A. Shoaib, J. H. Chuah, M. H. A. Nizar, S. Anis, and S. L. W. Ching, "Aortic valve segmentation using deep learning," in *Proc. IEEE-EMBS Conf. Biomed. Eng. Sci. (IECBES)*, Mar. 2021, pp. 528–532, doi: [10.1109/IECBES48179.2021.9398738](https://doi.org/10.1109/IECBES48179.2021.9398738).
- [29] S. Queirós, C. Dubois, P. Morais, T. Adriaenssens, J. C. Fonseca, J. L. Vilaça, and J. D'hooge, "Automatic 3D aortic annulus sizing by computed tomography in the planning of transcatheter aortic valve implantation," *J. Cardiovascular Comput. Tomogr.*, vol. 11, no. 1, pp. 25–32, Jan. 2017, doi: [10.1016/j.jcct.2016.12.004](https://doi.org/10.1016/j.jcct.2016.12.004).
- [30] M. J. Budoff, S. Lakshmanan, P. P. Toth, H. S. Hecht, L. J. Shaw, D. J. Maron, E. D. Michos, K. A. Williams, K. Nasir, A. D. Choi, K. Chinnaiyan, J. Min, and M. Blaha, "Cardiac CT angiography in current practice: An American society for preventive cardiology clinical practice statement," *Amer. J. Preventive Cardiol.*, vol. 9, Mar. 2022, Art. no. 100318, doi: [10.1016/j.ajpc.2022.100318](https://doi.org/10.1016/j.ajpc.2022.100318).
- [31] M. Angelillis, G. Costa, O. De Backer, V. Mochi, A. Christou, C. Giannini, P. Spontoni, M. De Carlo, L. Søndergaard, M. Miccoli, and A. S. Petronio, "Threshold for calcium volume evaluation in patients with aortic valve stenosis: Correlation with Agatston score," *J. Cardiovascular Med.*, vol. 22, no. 6, pp. 496–502, Jun. 2021, doi: [10.2459/JCM.0000000000001145](https://doi.org/10.2459/JCM.0000000000001145).
- [32] T. R. Carlidge, R. Bing, J. Kwiecinski, E. Guzzetti, T. A. Pawade, M. K. Doris, P. D. Adamson, D. Massera, M. Lembo, F. E. C. M. Peeters, C. Couture, D. S. Berman, D. Dey, P. Slomka, P. Pibarot, D. E. Newby, M.-A. Clavel, and M. R. Dweck, "Contrast-enhanced computed tomography assessment of aortic stenosis," *Heart*, vol. 107, no. 23, pp. 1905–1911, Dec. 2021, doi: [10.1136/heartjnl-2020-318556](https://doi.org/10.1136/heartjnl-2020-318556).
- [33] S. Saitta, F. Sturla, A. Caimi, A. Riva, M. C. Palumbo, G. Nano, E. Votta, A. D. Corte, M. Glauber, D. Chiappino, M. M. Marrocco-Trischitta, and A. Redaelli, "A deep learning-based and fully automated pipeline for thoracic aorta geometric analysis and planning for endovascular repair from computed tomography," *J. Digit. Imag.*, vol. 35, no. 2, pp. 226–239, Apr. 2022, doi: [10.1007/s10278-021-00535-1](https://doi.org/10.1007/s10278-021-00535-1).
- [34] A. Fantazzini, M. Esposito, A. Finotello, F. Auricchio, B. Pane, C. Basso, G. Spinella, and M. Conti, "3D automatic segmentation of aortic computed tomography angiography combining multi-view 2D convolutional neural networks," *Cardiovascular Eng. Technol.*, vol. 11, no. 5, pp. 576–586, Oct. 2020, doi: [10.1007/s13239-020-00481-z](https://doi.org/10.1007/s13239-020-00481-z).
- [35] A. Bratt, D. J. Blezek, W. J. Ryan, K. A. Philbrick, P. Rajiah, Y. K. Tandon, L. A. Walkoff, J. C. Cai, E. N. Sheedy, P. Korfiatis, E. E. Williamson, B. J. Erickson, and J. D. Collins, "Deep learning improves the temporal reproducibility of aortic measurement," *J. Digit. Imag.*, vol. 34, no. 5, pp. 1183–1189, Oct. 2021, doi: [10.1007/s10278-021-00465-y](https://doi.org/10.1007/s10278-021-00465-y).
- [36] A. Comelli, N. Dahiya, A. Stefano, V. Benfante, G. Gentile, V. Agnese, G. M. Raffa, M. Pilato, A. Yezzi, G. Petrucci, and S. Pasta, "Deep learning approach for the segmentation of aneurysmal ascending aorta," *Biomed. Eng. Lett.*, vol. 11, no. 1, pp. 15–24, Feb. 2021, doi: [10.1007/s13534-020-00179-0](https://doi.org/10.1007/s13534-020-00179-0).
- [37] F. Brutti, A. Fantazzini, A. Finotello, L. O. Müller, F. Auricchio, B. Pane, G. Spinella, and M. Conti, "Deep learning to automatically segment and analyze abdominal aortic aneurysm from computed tomography angiography," *Cardiovascular Eng. Technol.*, vol. 13, no. 4, pp. 535–547, Aug. 2022, doi: [10.1007/s13239-021-00594-z](https://doi.org/10.1007/s13239-021-00594-z).
- [38] M. M. Sieren, C. Widmann, N. Weiss, J. H. Moltz, F. Link, F. Wegner, E. Stahlberg, M. Horn, T. H. Oecherting, J. P. Goltz, J. Barkhausen, and A. Frydrychowicz, "Automated segmentation and quantification of the healthy and diseased aorta in CT angiographies using a dedicated deep learning approach," *Eur. Radiol.*, vol. 32, no. 1, pp. 690–701, Jan. 2022, doi: [10.1007/s00330-021-08130-2](https://doi.org/10.1007/s00330-021-08130-2).
- [39] I. Wolf, M. Vetter, I. Wegner, M. Nolden, T. Boettger, M. Hastenteufel, M. Schobinger, T. Kunert, and H.-P. Meinzer, "The medical imaging interaction toolkit (MITK)—A toolkit facilitating the creation of interactive software by extending VTK and ITK," *Proc. SPIE*, vol. 5367, pp. 16–27, May 2004, doi: [10.1117/12.535112](https://doi.org/10.1117/12.535112).
- [40] N. Zhang, G. Yang, W. Zhang, W. Wang, Z. Zhou, H. Zhang, L. Xu, and Y. Chen, "Fully automatic framework for comprehensive coronary artery calcium scores analysis on non-contrast cardiac-gated CT scan: Total and vessel-specific quantifications," *Eur. J. Radiol.*, vol. 134, Jan. 2021, Art. no. 109420, doi: [10.1016/j.ejrad.2020.109420](https://doi.org/10.1016/j.ejrad.2020.109420).
- [41] S. R. Ravichandran, B. Nataraj, S. Huang, Z. Qin, Z. Lu, A. Katsuki, W. Huang, and Z. Zeng, "3D inception U-Net for aorta segmentation using computed tomography cardiac angiography," in *Proc. IEEE EMBS Int. Conf. Biomed. Health Informat. (BHI)*, May 2019, pp. 1–4, doi: [10.1109/BHI.2019.8834582](https://doi.org/10.1109/BHI.2019.8834582).
- [42] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention MICCAI 2015*. Cham, Switzerland: Springer, 2015, pp. 234–241, doi: [10.1007/978-3-319-24574-4_28](https://doi.org/10.1007/978-3-319-24574-4_28).
- [43] NVIDIA. (2020). *3D-UNet Medical Image Segmentation for TensorFlow 1.x*. Accessed: Apr. 2, 2023. [Online]. Available: https://github.com/NVIDIA/DeepLearningExamples/tree/master/TensorFlow/Segmentation/UNet_3D_Medical
- [44] O. Çiçek, A. Abdulkadir, S. Lienkamp, T. Brox, and O. Ronneberger, "3D U-Net: Learning dense volumetric segmentation from sparse annotation," in *Medical Image Computing and Computer-Assisted Intervention MICCAI 2016*. Cham, Switzerland: Springer, Oct. 2016, pp. 424–432, doi: [10.1007/978-3-319-46723-8_49](https://doi.org/10.1007/978-3-319-46723-8_49).
- [45] F. Isensee, P. Kickingereder, W. Wick, M. Bendszus, and K. Maier-Hein, *No New-Net*. Cham, Switzerland: Springer, 2019, pp. 234–244, doi: [10.1007/978-3-030-11726-9_21](https://doi.org/10.1007/978-3-030-11726-9_21).
- [46] M. Abadi. (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Accessed: Apr. 14, 2023. [Online]. Available: <https://www.tensorflow.org/>
- [47] A. Sergeev and M. Del Balso, "Horovod: Fast and easy distributed deep learning in TensorFlow," 2018, *arXiv:1802.05799*.
- [48] Tensorflow. (2023). *XLA: Optimizing Compiler for Machine Learning*. Accessed: Jul. 20, 2023. [Online]. Available: <https://www.tensorflow.org/xla>
- [49] V. Nguyen and C. Tekur. (2020). *Accelerating TensorFlow on NVIDIA A100 GPUs*. Accessed: Jul. 20, 2023. [Online]. Available: <https://developer.nvidia.com/blog/accelerating-tensorflow-on-a100-gpus/>



GIADA MINGHINI received the master's degree in biotechnology for the environment and health from the University of Ferrara, Ferrara, Italy, in 2021. Since 2022, she has been a Researcher at the University of Ferrara and her activity focuses on the analysis of cardiovascular images to develop a high-performance automatic system to detect, segment, and quantify the volume of calcium in the aortic valve.



ARMANDO UGO CAVALLO received the degree in medicine and surgery from the University of Salerno, in 2014, and the Ph.D. degree in medical biotechnology and translational medicine with a major focus on radiomics and artificial intelligence applications in cardiovascular and thoracic imaging. From 2015 to 2019, he was a Radiology Resident at the University of Rome "Tor Vergata." In 2018, he was a Research Fellow in cardiovascular imaging at Cleveland University Hospital/Case Western Reserve University, Cleveland, OH, USA. He has been a Staff Radiologist at Istituto Dermopatico dell'Immacolata, Rome, Italy, since 2022.



ANDREA MIOLA received the M.Sc. degree in astrophysics and cosmology and the second-level professional master's degree in space missions science, design, and applications (SPICES) from the Alma Mater Studiorum—University of Bologna, Italy, in 2019 and 2022, respectively.

He is currently pursuing the Ph.D. degree in physics with the University of Ferrara, Ferrara, Italy. He continues to explore the fields of deep learning and HPC as part of the Ph.D. research activity. In 2021, he started working in artificial neural networks at MEEO, Ferrara. Among his research interests, there are analysis of computer-aided medical images and satellite-based imagery.



VALENTINA SISINI received the bachelor's degree in computer science with a thesis on quantum computing, in 2020. She is currently pursuing the master's degree in computer engineering and is working on her thesis in the field of high-performance computing (HPC). She has experience as a software developer. She has contributed to the book *Reti neurali non supervisionate: il cognitrone di Fukushima: Con codice in linguaggio C*. Her research interest includes artificial intelligence, particularly in the linguistic aspect.



ENRICO CALORE received the graduate degree in computer engineering from the University of Padua, Italy, in 2010, and the Ph.D. degree in computer science from the University of Milan, Italy, in 2014. Later, he was a Postdoctoral Researcher with the Italian National Institute of Nuclear Physics (INFN) Ferrara and with the University of Ferrara, where he was also an Adjunct Professor with the Mathematics and Computer Science Department. He is currently a Research Engineer with INFN Ferrara. He is the author of more than 60 articles and conference papers. His research interests include HPC, parallel, and distributed computing, focusing in particular on scientific computing, parallel programming, performance engineering, and energy efficiency.



FRANCESCA FORTINI received the master's degree in applied biomolecular sciences from the University of Perugia, Perugia, Italy, in 2014, and the Ph.D. degree in molecular medicine and pharmacology from the University of Ferrara, Ferrara, Italy, in 2018. Since 2018, she has been a Scientific Researcher at the Translational Research Center of Maria Cecilia Hospital, GVM Care and Research, Cotignola, Ravenna, Italy. Her current research interests include the identification of predisposing factors for aortic valve stenosis and novel prognostic factors for the progression of peripheral artery disease in diabetic patients. She is also interested in studying the role of the Notch pathway in the cardiotoxicity induced by chemotherapeutic drugs.



RINO MICHELONI (Senior Member, IEEE) is with Avaneidi srl, Saronno, Italy. Before that, he was the Vice-President and a Fellow at Microsemi/Microchip Corporation, where he led a team working on NAND flash technology, machine learning, and error correction codes. Early in his career, he developed the world's first NOR flash memory with embedded ECC technology and the world's first NAND flash with embedded BCH. He has coauthored more than 100 publications and he holds 299 patents worldwide (including 143 U.S. patents). In 2020, he was selected for the European Inventor Award.



PAOLA RIZZO received the M.S. degree in molecular biology from the University of Naples, Naples, Italy, in 1989, and the Ph.D. degree in biomedical and endocrinological sciences from the University of Ferrara, Ferrara, Italy, in 2004. From 1995 to 2009, she was involved in cancer research, specifically on the role of viruses in human mesothelioma, and then on the role of the notch pathway in estrogen-responsive breast cancer as an Assistant Professor at the Loyola University Cancer Center, Chicago, IL, USA. Currently, she is an Associate Professor of applied biology at the University of Ferrara, where since 2010, she has been conducting cardiovascular research focused on the regulation of endothelial function in the context of atherosclerosis, on the cardiotoxicity of chemotherapeutic drugs, and more recently, on the molecular mechanisms of vascular calcification of the aortic valve. She has published more than 100 publications in peer-reviewed journals.



SEBASTIANO FABIO SCHIFANO received the degree in computer science from the University of Pisa, Italy. He spent his early scientific career at IEI-CNR and INFN, and now is a Professor of computer science at the University of Ferrara, Italy. He had a major role in several projects for the development of parallel systems optimized for lattice gauge theory, fluid dynamics, and spin-glasses. In 2007, he was the coauthor of the QPACE project to develop a massively parallel system based on IBM Cell-BE processors, interconnected by a custom 3-D-mesh network, awarded as the best Green500 system, in November and June. From 2005 to 2010, he contributed to the design and development of the Janus systems, an array of FPGAs for spin-glass simulations. He is the author of more than 100 articles and conference papers. More recently, his research activities have focused on the design and optimization of applications for multi- and many-core processor architectures.



FRANCESCO VIECELI DALLA SEGA received the master's degree in chemistry and pharmaceutical technologies and the Ph.D. degree in biochemistry from the University of Bologna. He is currently a Senior Researcher at the Translational Research Center, Maria Cecilia Hospital (MCH), where he is also with the Translational Research Center. He has 15 years of experience in basic and applied research in the fields of biochemistry, cellular and molecular biology, histology, and pathology. He has participated in numerous clinical studies in the field of cardiovascular diseases. His main research interests include molecular mechanisms underlying endothelial dysfunction and cardiovascular disease.



CRISTIAN ZAMBELLI (Member, IEEE) received the M.Sc. and Ph.D. degrees in electronic engineering from the University of Ferrara, Ferrara, Italy, in 2008 and 2012, respectively. Currently, he holds an Associate Professor at the University of Ferrara. His current research interests include the electrical characterization, physics, and reliability modeling of different non-volatile memory technologies for data storage applications and emerging in-memory computing concepts. He is also interested in the evaluation of the solid-state drives reliability/performance tradeoffs exposed by the integrated memory technology to develop computational storage accelerators for the HPC scenario.

...

Open Access funding provided by 'Università degli Studi di Ferrara' within the CRUI CARE Agreement