



# Deep Learning Framework for Real-Time Estimation of *in-silico* Thrombotic Risk Indices in the Left Atrial Appendage

Xabier Morales Ferez<sup>1\*</sup>, Jordi Mill<sup>1</sup>, Kristine Aavild Juhl<sup>2</sup>, Cesar Acebes<sup>1</sup>, Xavier Iriart<sup>3</sup>, Benoit Legghe<sup>3</sup>, Hubert Cochet<sup>3</sup>, Ole De Backer<sup>4</sup>, Rasmus R. Paulsen<sup>2</sup> and Oscar Camara<sup>1</sup>

<sup>1</sup> Physense, BCN Medtech, Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain, <sup>2</sup> DTU Compute, Technical University of Denmark, Kongens Lyngby, Denmark, <sup>3</sup> IHU Liryc, University Hospital of Bordeaux, Bordeaux, France, <sup>4</sup> Department of Cardiology, Rigshospitalet, University of Copenhagen, Copenhagen, Denmark

## OPEN ACCESS

### Edited by:

Linwei Wang,  
Rochester Institute of Technology,  
United States

### Reviewed by:

Cristiana Corsi,  
University of Bologna, Italy  
Guang Yang,  
Imperial College London,  
United Kingdom

### \*Correspondence:

Xabier Morales Ferez  
xabier.morales@upf.edu

### Specialty section:

This article was submitted to  
Computational Physiology and  
Medicine,  
a section of the journal  
Frontiers in Physiology

Received: 14 April 2021

Accepted: 13 May 2021

Published: 28 June 2021

### Citation:

Morales Ferez X, Mill J, Juhl KA,  
Acebes C, Iriart X, Legghe B,  
Cochet H, De Backer O, Paulsen RR  
and Camara O (2021) Deep Learning  
Framework for Real-Time Estimation  
of *in-silico* Thrombotic Risk Indices in  
the Left Atrial Appendage.  
*Front. Physiol.* 12:694945.  
doi: 10.3389/fphys.2021.694945

Patient-specific computational fluid dynamics (CFD) simulations can provide invaluable insight into the interaction of left atrial appendage (LAA) morphology, hemodynamics, and the formation of thrombi in atrial fibrillation (AF) patients. Nonetheless, CFD solvers are notoriously time-consuming and computationally demanding, which has sparked an ever-growing body of literature aiming to develop surrogate models of fluid simulations based on neural networks. The present study aims at developing a deep learning (DL) framework capable of predicting the endothelial cell activation potential (ECAP), an *in-silico* index linked to the risk of thrombosis, typically derived from CFD simulations, solely from the patient-specific LAA morphology. To this end, a set of popular DL approaches were evaluated, including fully connected networks (FCN), convolutional neural networks (CNN), and geometric deep learning. While the latter directly operated over non-Euclidean domains, the FCN and CNN approaches required previous registration or 2D mapping of the input LAA mesh. First, the superior performance of the graph-based DL model was demonstrated in a dataset consisting of 256 synthetic and real LAA, where CFD simulations with simplified boundary conditions were run. Subsequently, the adaptability of the geometric DL model was further proven in a more realistic dataset of 114 cases, which included the complete patient-specific LA and CFD simulations with more complex boundary conditions. The resulting DL framework successfully predicted the overall distribution of the ECAP in both datasets, based solely on anatomical features, while reducing computational times by orders of magnitude compared to conventional CFD solvers.

**Keywords:** geometric deep learning, left atrial appendage, convolutional neural network, thrombus-atrial fibrillation, computational fluid dynamics, principal component analysis

## 1. INTRODUCTION

Atrial fibrillation (AF) is the most common clinically significant arrhythmia, with a cumulative lifetime development risk above 30% in individuals of European ancestry (Benjamin et al., 2019). AF is defined by a quivering or irregular heartbeat (arrhythmia) caused by chaotic electric activity, which leads to irregular contraction and wall rigidity of the left atrium (LA), preventing effective

flow of the blood to the ventricles. Such hemodynamic alterations, alongside factors such as endothelial or endocardial dysfunction, including a state of hypercoagulability, increase the risk of cerebrovascular accidents by allowing thrombus formation within the LA (Watson et al., 2009); if dislodged, such thrombi can occlude the cerebral circulation, causing a thromboembolic (ischemic) stroke. In fact, non-valvular AF is responsible for 15–20% of all cardioembolic ischemic strokes, which preferentially form at the left atrial appendage (LAA) (Cresti et al., 2019), an heterogeneous, tubular structure derived from the anterior wall of the LA.

In this regard, researchers have explored the correlation between LAA morphology and the risk of stroke (Yaghi et al., 2020; Dudzińska-Szczerba et al., 2021; Słodowska et al., 2021). Nonetheless, so far the results have been ambiguous, as the current classifications and associated morphological parameters of the LAA are often entirely subjective, hand-crafted features; there is a need for more systematic shape analysis of the LAA with advanced and observer-independent computational tools such as statistical atlases (Slipsager et al., 2019).

Besides, due to the critical role of blood stasis in thrombogenesis, the interest in the analysis of LA hemodynamics is gaining momentum. Yet, the intricate behavior of the left atrium as a modulator of left ventricular filling (reservoir, conduit, and booster pump function; Vieira et al., 2014), coupled to a substantial anatomical heterogeneity, makes modeling left atrial hemodynamics a notoriously difficult task. Consequently, computational fluid dynamics (CFD) analyses have emerged as an invaluable tool in analyzing the mechanistic relationship between patient-specific organ morphology and blood stasis (García-Isla et al., 2018; Masci et al., 2019; García-Villalba et al., 2021). Nevertheless, conventional CFD methods are renowned for their large memory requirements and long computing times (Liang et al., 2018), which also involve extensive pre-processing of each patient-specific mesh, resulting in studies with very limited sample sizes and severely hindering its suitability for time-sensitive clinical applications.

As a response, neural networks have increasingly been employed in complex dynamical systems such as fluid dynamics, resulting in highly accurate surrogate models that can be evaluated with significantly less computational resources and several orders of magnitude faster than conventional finite element solvers (Hennigh, 2017). Recently, deep learning (DL) has made its way into biological fluid modeling, aiming at predicting blood velocity vector fields or derived hemodynamic parameters that play a crucial role in the diagnosis and development of several cardiovascular diseases (Liang et al., 2018; Li et al., 2021). Nevertheless, most studies have mostly focused on structures such as the aorta, which present a less complex morphology and hemodynamic profile than the LA and LAA. That being said, applying conventional DL models to non-Euclidean domains, such as graphs and meshes, in which medical data is often best represented, is not a trivial task, as most widespread neural networks can only operate over regular data such as images (Fey et al., 2018). In this regard, geometric deep learning approaches, which are tailored to operate over graph data, have already been applied to biomedical meshes,

especially in cardiac electrophysiological models (Grandits et al., 2021; Meister et al., 2021).

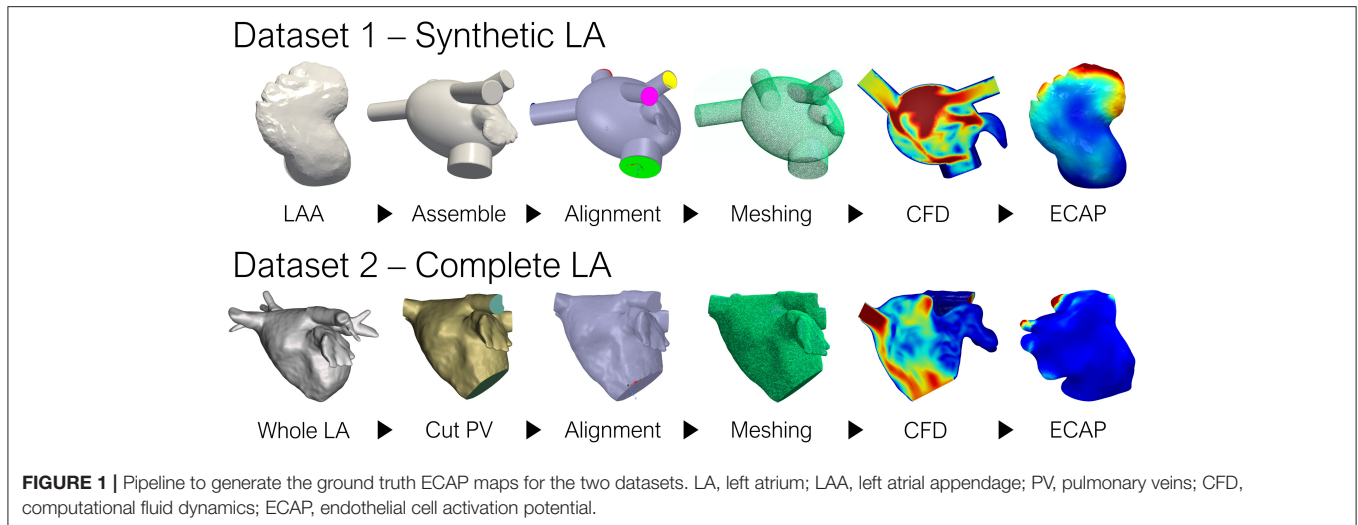
Hence, in the present study, we have leveraged a collection of distinct DL models, which are well-tailored to deal with mesh data, to develop a CFD surrogate capable of learning the complex relationship between the heterogeneous LAA morphology and the endothelial cell activation potential (ECAP), parameter linked to an increased risk of thrombosis. By employing neural networks, there is no need to manually craft morphological features, ensuring that the model only learns the most relevant anatomical characteristics toward the automatic prediction of ECAP. Moreover, once trained, neural networks allow the prediction of ECAP maps in new unseen patients, orders of magnitude faster than it is possible with current CFD solvers. The implemented DL approaches included principal component analysis (PCA) based shape analysis coupled to fully connected layers, flattening the LAA morphology to a UV space to leverage convolutional neural networks (CNN) and geometric deep learning, which is perfectly suited to non-Euclidean data such as meshes. All the mentioned methods were first tested on a simplified LA model containing 256 real and synthetic LAA (dataset 1). In addition, the best performing model was further tested on a second, more realistic 114 patient dataset, which incorporated the entire patient-specific LA anatomy (dataset 2).

## 2. METHODS

The overall pipeline employed to generate the ground-truth data (i.e., the *in-silico* ECAP index from CFD simulations in the whole dataset of 370 geometries) is shown in **Figure 1**. Preprocessing slightly differed between the two datasets: LAA comprising dataset 1 were all assembled to an oval LA, while in dataset 2, which considered the whole patient-specific LA anatomy, all pulmonary veins (PV) were trimmed at the first branching to define the inlets and outlets. Later, tetrahedral volumetric meshes were generated to run the CFD simulations and compute the ground truth ECAP maps. For the networks to capture the most relevant morphological features, the triangular meshes employed to describe the anatomy of each LAA had to be transformed according to the prerequisites of each of the implemented DL methods. Lastly, the neural networks were trained to learn the arbitrary non-linear function linking the geometry of the LAA and its corresponding ECAP maps. Model prototyping and fine-tuning were completed on the synthetic LA dataset (dataset 1). Afterward, the best-performing network was further tested in the more realistic, complete LA dataset (dataset 2).

### 2.1. Data

The first dataset (dataset 1) was derived from computed tomography (CT) images provided by the Department of Radiology at Rigshospitalet (Copenhagen, Denmark) acquired as part of the Copenhagen General Population Study (Nordestgaard et al., 2012). It was comprised of 256 LAA, combining 54 real patients and 202 synthetic LA geometries. The latter, being borrowed from a preceding study (Morales et al., 2020), stem from a statistical shape model (SSM) based on 103 real LAA surfaces (Slipsager et al., 2019). In this synthetic LA model, only



the geometry of the LAA was considered, as incorporating the highly heterogeneous LA anatomy would qualitatively increase the inter-subject variability of the hemodynamic parameters. Thus, all appendages were assembled to a common oval approximation of the LA (García-Isla et al., 2018), reducing the complexity of the maps to be predicted, and ensuring that the LAA morphology remained the only independent variable affecting the ECAP values.

Conversely, the second dataset considered the complete patient-specific LA morphology during CFD simulations (dataset 2). The data was provided by Hospital Haut-Lévêque (Bordeaux, France), originating from pre-procedural high-quality CT scans from 114 AF patients that underwent a left atrial appendage occlusion (LAAO) intervention. Both studies were approved by the local Institutional Ethics Committees, and all patients provided informed consent.

## 2.2. CFD Simulations

A total of 370 CFD simulations were run to generate the ground-truth data, 256 of which corresponded to the synthetic LA dataset, while the remaining 114 were part of dataset 2. All synthetic LAA simulations on dataset 1 were borrowed from a preceding study, adjusting the setup of the simulations for the remaining 54 real morphologies accordingly (Morales et al., 2020). First, an input velocity profile was set in the PV, based on clinical observations (Fernández-Pérez et al., 2012). Second, the mitral valve (MV) was considered as a wall boundary during diastole, while an outlet pressure of 1,067 Pa was set through systole. The motion of the LA was based on a diffusion-based dynamic mesh emanating from the MV ring plane, adjusted according to literature (Veronesi et al., 2008; Mill et al., 2019). Only a single heartbeat was completed for each simulation.

On the other hand, simulations from dataset 2 featured more complex boundary conditions (BC), with the inlet being defined at the PV based on pressure wave measurements from an AF patient; the velocity outlet was set on the MV based on Doppler ultrasound velocity profiles derived from a single

patient. Therefore, while the LA morphology was completely patient-specific, the same boundary conditions were shared among all cases. However, all BCs were synchronized to their corresponding patient's electrocardiogram. The dynamic mesh governing LA motion was changed to a spring-based model. Unlike in dataset 1, three full heartbeats were completed for each simulation, aiming to reach a steady state. Only the last heartbeat was considered when computing the risk indices of thrombosis. Lastly, whereas final tetrahedral volumetric meshes for dataset 1 consisted of ~350 k elements, each mesh from dataset 2 doubled that figure at around 800 and 900 k elements, after a mesh convergence study that included meshes up to 1 M elements.

Simulations were computed on Ansys Fluent 19 R32 (ANSYS Inc, USA)<sup>1</sup> and automatized leveraging the MATLAB AAS toolbox,<sup>2</sup> while post-processing was performed in Paraview<sup>3</sup> alongside in-house python scripts. The blood was treated as a Newtonian fluid, with a density of 1,060 Kg/m<sup>3</sup> and a viscosity of 0.0035 Pa/s, while using time-steps of 0.1 s.

The endothelial cell activation potential (ECAP), proposed by Di Achille et al. (2014), was the parameter chosen to evaluate the risk of thrombosis in the LAA. Since the pathophysiology of thromboembolism in AF is based upon the formation of mural thrombi, the ECAP focuses on hemodynamic behavior in the proximity of the vessel wall. More precisely, the ECAP is defined as the ratio between the oscillatory shear index (OSI) and the time-averaged wall shear stress (TAWSS).

$$ECAP = \frac{OSI}{TAWSS}. \quad (1)$$

Thereby, a dimensionless parameter related to thrombosis is obtained, thus, avoiding the need for a more complicated neural network architecture capable of handling temporal sequential data. High ECAP values result from low TAWSS and high OSI values, indicating the presence of low velocities and high flow

<sup>1</sup><https://www.ansys.com/products/fluids/ansys-fluent>

<sup>2</sup><https://es.mathworks.com/products/matlab.html>

<sup>3</sup><https://www.paraview.org/>

complexity, which is associated with endothelial susceptibility and risk of thrombus formation.

## 2.3. Deep Learning Architectures

### 2.3.1. Principal Component Analysis—Fully Connected Model

Although it is gradually being replaced by more sophisticated non-linear models, principal component analysis (PCA), has long been employed to learn a linear latent space of 3D registered meshes, for tasks such as compression, reconstruction, and animation (Zhou et al., 2020). Nevertheless, PCA requires all meshes to be registered to a common template so that the same topology and connectivity are shared among them. In our case, this step was completed through non-rigid volumetric registration of signed distance fields, based on the work by Slipsager et al. (2019), registering all meshes to a common template comprised of 2,466 vertices. For the PCA analysis, the spatial coordinate of the nodes composing the mesh were employed as the input features. Thus, the morphology of each LAA can be expressed by a small set of scalar values through truncated PCA following:

$$X = \bar{X} + \sum_{i=1}^M \alpha_i \sqrt{\lambda_i} W_i, \quad \rightarrow \quad \alpha_i = \frac{W_i^T (Y - \bar{Y})}{\sqrt{\lambda_i}}, \quad (2)$$

where  $\bar{X}$  is the mean shape,  $W_i$  and  $\lambda_i$  being the set of eigenvectors and eigenvalues of the covariance matrix for the retained number,  $i$ , of principal components (PC). Hereby, if the variability of the dataset is explained by a small set of PCs, each LAA anatomy can be expressed by a number  $i$ , of  $\alpha_i$  scalars that can be fed directly to any regular neural network. For our dataset, a total of  $n = 32$  PCs were kept, provided that the training dataset was large enough. By doing so, 97.6% of the morphological variability was retained for the synthetic cases and 94.1% for the real LAA geometries. Afterward, the non-linear mapping between the low dimensional representation of LAA morphology and its corresponding ECAP maps was completed through a fully connected feed-forward neural network (FCN). It comprised five hidden layers, as shown in **Figure 2A**, sequentially increasing the size of each layer. The whole model was implemented in Keras,<sup>4</sup> using TensorFlow<sup>5</sup> as backend.

### 2.3.2. UV Mapping—U-Net Model

Although PCA models have been extremely successful (Tewari et al., 2017), they often struggle to capture finer details, since the resulting latent space is a linear combination of the input features. Alternative models such as convolutional neural networks (CNN) are widely employed to capture spatial features in regular grids (Zhou et al., 2020), which owing to a combination of desirable properties, such as local connectivity, weight sharing, and displacement invariance, became the backbone of fields such as computer vision. That being said, due to the irregular nature of mesh data, spatially-shared convolution kernels cannot be

directly leveraged, unless the 3D mesh data is mapped to a UV space, also known as flattening.

Consequently, the LAA were “flattened” based on the approach described in Acebes et al. (2021). First, each LAA was divided into an equivalent number of isolines, based on the geodesic distance from the ostium to the LAA apex, which was computed through a heat equation method (Crane et al., 2013). Subsequently, an equivalent number of vertices were sampled from each isoline, through an angular mapping performed by pivoting around the centroid of each isoline. Meanwhile, the points closest to the position of the circumflex artery, which was manually marked from the CT images, were chosen as the reference 0–360° angle. Once polar coordinates had been derived, each LAA mesh was represented as a 2D image either as a circumferential polar plot, also known as a Bull’s eye plot (Cerqueira et al., 2002), with the apex of the LAA located in the center of the circumference (see **Figure 2B**), or as a rectangular image whose two axes consist in the apical-ostium distance and angular mapping. Even though the outer ring of the bull’s eye plot undergoes distortion, it better preserves the LAA topology avoiding the cut-off produced by the flattening process. Conversely, the Cartesian grid representation faces much stronger warping in the area close to the apex, which is far overrepresented relative to its actual surface area in the 3D mesh. Therefore, both flattened representations were included to weigh up their trade-offs. Lastly, as shown in **Figure 3**, the bull’s eye plot was padded to a rectangular image before feeding it to the neural network. During training, the padded regions were not taken into account when performing the loss calculation and subsequent accuracy measurements.

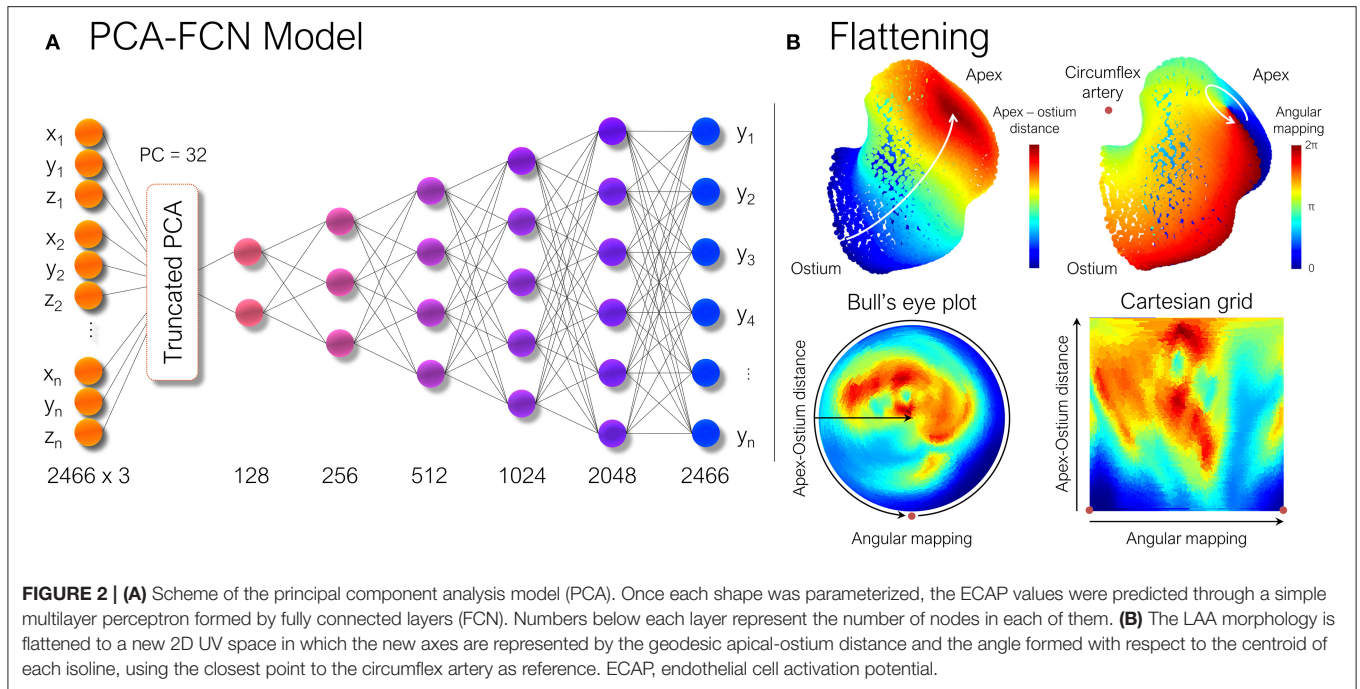
With regards to the DL model, we opted for a conventional U-Net architecture which is comprised of overlapping convolutional layers arranged in a typical encoder-decoder bowtie structure (Ronneberger et al., 2015), consisting in sequential pooling operations that ensure that multi-scale features are learned from the input data. In addition, skip connections, encourage the network to reuse low-level features in the decoding layers, which result in state-of-the-art performance in several tasks such as medical segmentation. The vanilla U-Net provided by Thurey et al. (2020) was leveraged, which was implemented to predict turbulent flow over a set of distinct airfoils. In our case, the input features consisted in the spatial coordinates of the vertices sampled from each LAA mesh during UV mapping, arranged as a three-channel depth tensor, analogous to an RGB image in computer vision tasks. Similarly to the original paper, seven convolutional blocks were employed (**Figure 3**), each including batch normalization layers and ReLU activations after convolution layers, as is the standard practice.

### 2.3.3. Geometric Deep Learning

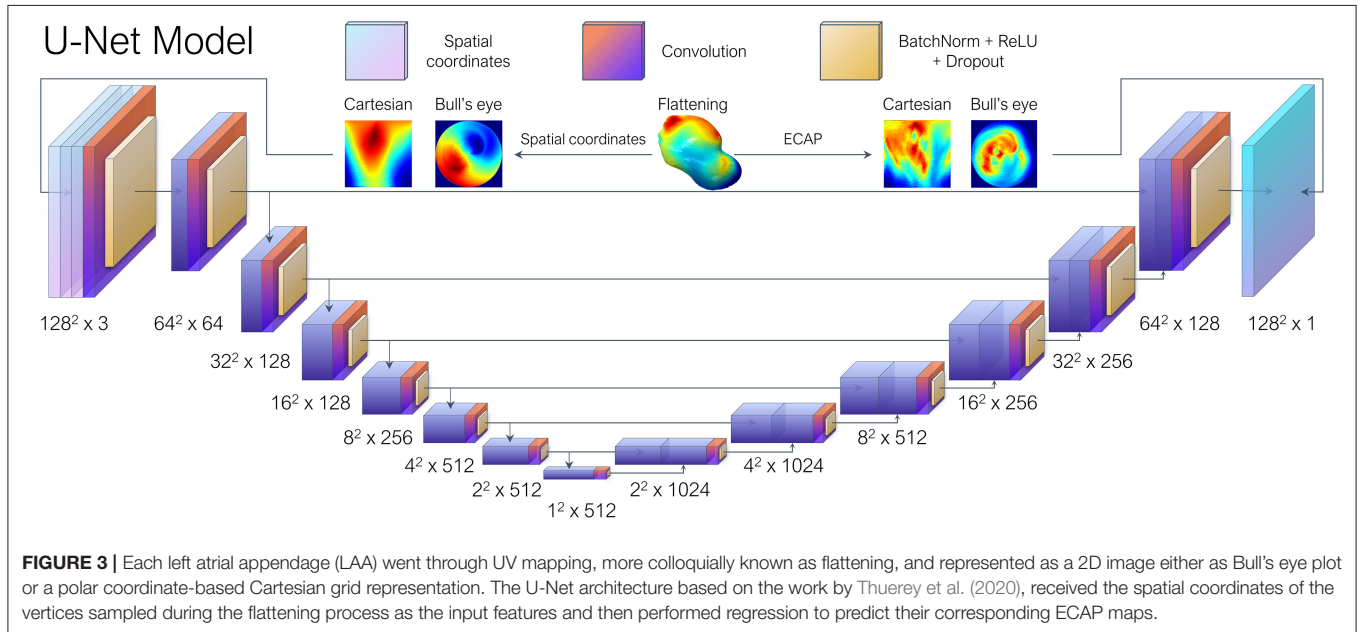
While UV mapping enables the direct leveraging of CNNs over meshes, the resulting 2D image suffers from distortion as the original mesh data must be cut and warped. Moreover, it cannot be easily extended to more general volumetric data. Whilst alternative workarounds such as voxelization exist, the most efficient way of representing 3D surface shapes and topologies is through polygonal meshes (Hanocka et al., 2019). As a

<sup>4</sup><https://keras.io/>

<sup>5</sup><https://www.tensorflow.org/>



**FIGURE 2 | (A)** Scheme of the principal component analysis model (PCA). Once each shape was parameterized, the ECAP values were predicted through a simple multilayer perceptron formed by fully connected layers (FCN). Numbers below each layer represent the number of nodes in each of them. **(B)** The LAA morphology is flattened to a new 2D UV space in which the new axes are represented by the geodesic apical-ostium distance and the angle formed with respect to the centroid of each isoline, using the closest point to the circumflex artery as reference. ECAP, endothelial cell activation potential.



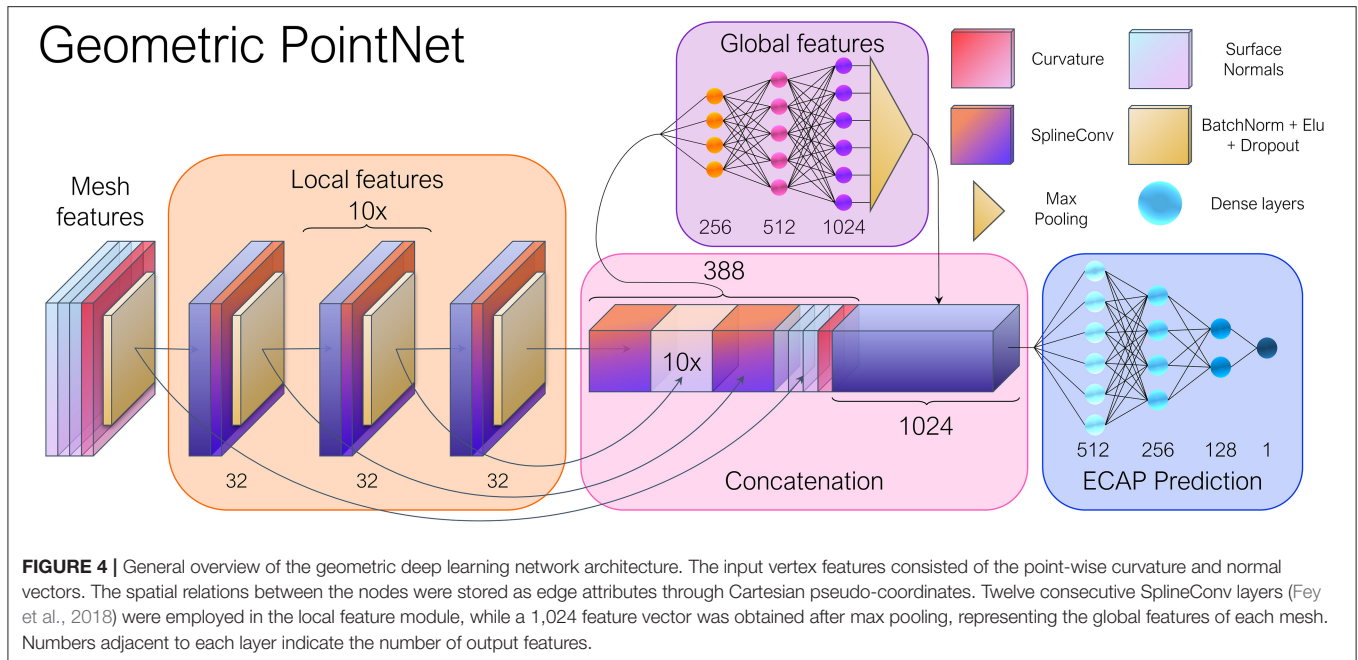
**FIGURE 3 |** Each left atrial appendage (LAA) went through UV mapping, more colloquially known as flattening, and represented as a 2D image either as Bull's eye plot or a polar coordinate-based Cartesian grid representation. The U-Net architecture based on the work by Thurey et al. (2020), received the spatial coordinates of the vertices sampled during the flattening process as the input features and then performed regression to predict their corresponding ECAP maps.

response, a set of emerging methods, under the umbrella term of Geometric DL, are succeeding in generalizing DL models to non-Euclidean domains such as polygonal meshes, seamlessly extending operations such as convolutions to the native form of the data (Bronstein et al., 2017).

Among the array of available graph CNN layers, we opted for SplineCNN (Fey et al., 2018), since being a spatial method, it offers several advantages when dealing with meshes. In particular, it avoids the need of establishing mesh correspondence. Additionally, defining the spatial relations between vertex

features becomes trivial by employing pseudo-coordinates. In our use case, pseudo-coordinates were obtained by computing the relative distance in Cartesian coordinates between the vertices of each edge. During the training process, these edge attributes define how the input features will be aggregated in the neighborhood of a given node. Additionally, the vertex-wise curvature and normals were fed to the network as vertex features.

As aforementioned, besides convolution, operations such as pooling and strided convolutions play a key role in the success of CNNs, by allowing the network to sequentially extract larger



scale and abstract features. Consequently, a PointNet-inspired (Qi et al., 2016) architecture was implemented, in which a series of consecutive layers focus on learning local features. Subsequently, the resulting feature arrays are concatenated and fed to a multilayer perceptron that generates a vector of global features using max-pooling, as shown in **Figure 4**. An almost identical model was employed by Meister et al. (2021) to estimate left ventricular depolarization times, albeit using a distinct convolutional operator. By swapping the multilayer perceptrons employed on the original PointNet (Qi et al., 2016) for graph convolutional layers, we expect to better exploit local correlations and weight sharing, while providing topological information to the network, which is explicitly absent in point cloud data.

The model was constructed by using PyTorch Geometric (PyG),<sup>6</sup> a Geometric DL extension of PyTorch.<sup>7</sup> PyG offers a broad set of convolution and pooling operations that extend the capabilities of traditional CNNs to irregularly structured data such as graphs and manifolds. With this in mind, the mesh dataset resulting from the simulations were converted into individual graphs. Together with PyVista,<sup>8</sup> we converted each mesh to a graph represented by  $G = (V, E)$ , with  $V = 1, \dots, N$  being the set of nodes, and  $E$  corresponds to the set of edges of the triangular faces. For each vertex, we computed the curvature and surface normal vectors, totaling four input feature channels.

## 2.4. Hyperparameter Tuning

A thorough grid search was carried out to fine-tune the models by iteratively swapping several hyperparameters while keeping a fixed seed in the dataset split. In the PCA-FCN and U-Net

model, ReLU activations were employed coupled with a learning rate of  $lr = 0.01$  and  $lr = 0.0005$  and trained during 150 and 300 epochs, respectively. Concerning general hyperparameters of the geometric DL model, the exponential linear unit (ELU) provided the best results among all activation functions, as it is standard in many mesh-related tasks. In addition, the training loop was carried out through 300 epochs with a batch size of 16 and a learning rate of 0.001. In regards to the parameters of the SplineConv layer, a B-spline basis of degree 1 and a kernel size of  $k = 5$  were chosen, following suggestions by the authors (Fey et al., 2018). All models employed a dropout of 0.1 and included Adam as the optimizer. Similarly to Thuerey et al. (2020), while alternative loss functions such as  $L_2$  Loss and smooth- $L_1$  yielded similar results, the  $L_1$  loss still had an edge over them.

## 2.5. Experimental Setup

All of the above-presented models were first tested on the synthetic LA dataset (dataset 1). Some of the experiments aimed to determine whether the synthetic data resulting from the statistical shape model were sufficiently representative of real patient data. If so, synthetic data could be of help with the data-hungry nature of neural networks in the face of data scarcity, which is a recurring issue in the medical field. On the one hand, k-fold cross-validation was performed, first training in the real ( $n = 54$ ) and synthetic ( $n = 202$ ) cases separately, and later combining both datasets. We have called these experiments “Cross real,” “Cross synth,” and “Cross,” respectively. The experiment in the real dataset was divided into 6-folds while the latter two experiments run 8-folds to ensure the groups were even. In addition, as only the areas of high ECAP values are pro-thrombotic, we wanted to assess the capability of the models in predicting the areas with the highest ECAP. For this purpose, a binary classification was performed taking the 90th percentile

<sup>6</sup>[https://github.com/rusty1s/pytorch\\_geometric](https://github.com/rusty1s/pytorch_geometric)

<sup>7</sup><https://pytorch.org/>

<sup>8</sup><https://docs.pyvista.org/>

**TABLE 1** | Prediction accuracy results in terms of mean absolute error (MAE) and true positive rate (TPR) for the cross-validation experiments.

Model	Cross	Cross real	Cross synth	Cross (%)	Cross real (%)	Cross synth (%)
	Mean absolute error (MAE)			True positive rate (TPR)		
PCA-FCN	0.608 ± 0.021	0.591 ± 0.023	0.603 ± 0.008	69	41	69
Cartesian	0.651 ± 0.007	0.661 ± 0.028	0.617 ± 0.017	55	20	64
Bull's eye	0.654 ± 0.009	0.582 ± 0.027	0.628 ± 0.009	64	33	65
Geometric	<b>0.521 ± 0.013</b>	<b>0.519 ± 0.021</b>	<b>0.514 ± 0.017</b>	<b>77</b>	<b>57</b>	<b>79</b>

PCA, principal component analysis; FCN, fully connected network. Bold signals the model with the best value for a given experiment and metric.

of the distribution as the threshold, which roughly equated to 4, following a similar approach to Di Achille et al. (2014) in abdominal aortic aneurysms.

On the other hand, the second set of experiments was conducted in which the amount of training data was sequentially scaled, to monitor the generalization and accuracy improvement (or lack thereof) on the testing dataset. Therefore, the testing scheme from the cross-validation experiments was maintained, but several runs were completed for each testing fold, changing the amount of available training data on each. For the first two experiments, which we deemed “Sequential Real” and “Sequential synthetic,” the real and synthetic morphologies were trained and tested separately. These experiments aimed to learn which models performed better with few amounts of data. Alternatively, in a third experiment, all the 202 synthetic LAA were employed as the training baseline. On top of this baseline, real geometries were sequentially added while testing on the remaining real cases. The objective of this experiment, deemed as “Sequential real+synthetic,” focused on the number of real cases required to build a model just trained on synthetic data, being able to properly generalize to patient-specific LAA morphologies. In all the aforementioned experiments, 10% of all the training data was employed as validation and used to select the best performing model. In addition, due to the stochastic nature of the training process in neural networks, the presented results have been averaged across several runs.

Lastly, the best-performing model from the previous experiments underwent further testing on dataset 2. While the complete patient-specific LA morphology was included during simulation, solely the LAA anatomy was fed to the neural network during the prediction of the ECAP maps. A single 10-fold cross-validation experiment was completed in this dataset along with the binary classification. The 90th percentile equalled 16 in this case.

### 3. RESULTS

The ECAP distributions resulting from both simulations were distinct due to the different geometry and boundary conditions. The ECAP maps from the LA synthetic dataset had a mean value of  $\mu = 2.14 \pm 1.41$ , whereas dataset 2 exhibited a far more lopsided distribution, with a mean value of  $\mu = 34.82 \pm 251.68$  but a median equal to 0.492.

Each simulation in the synthetic LA dataset lasted around 3–4 h, whereas it took at least 24 h to complete every single dataset

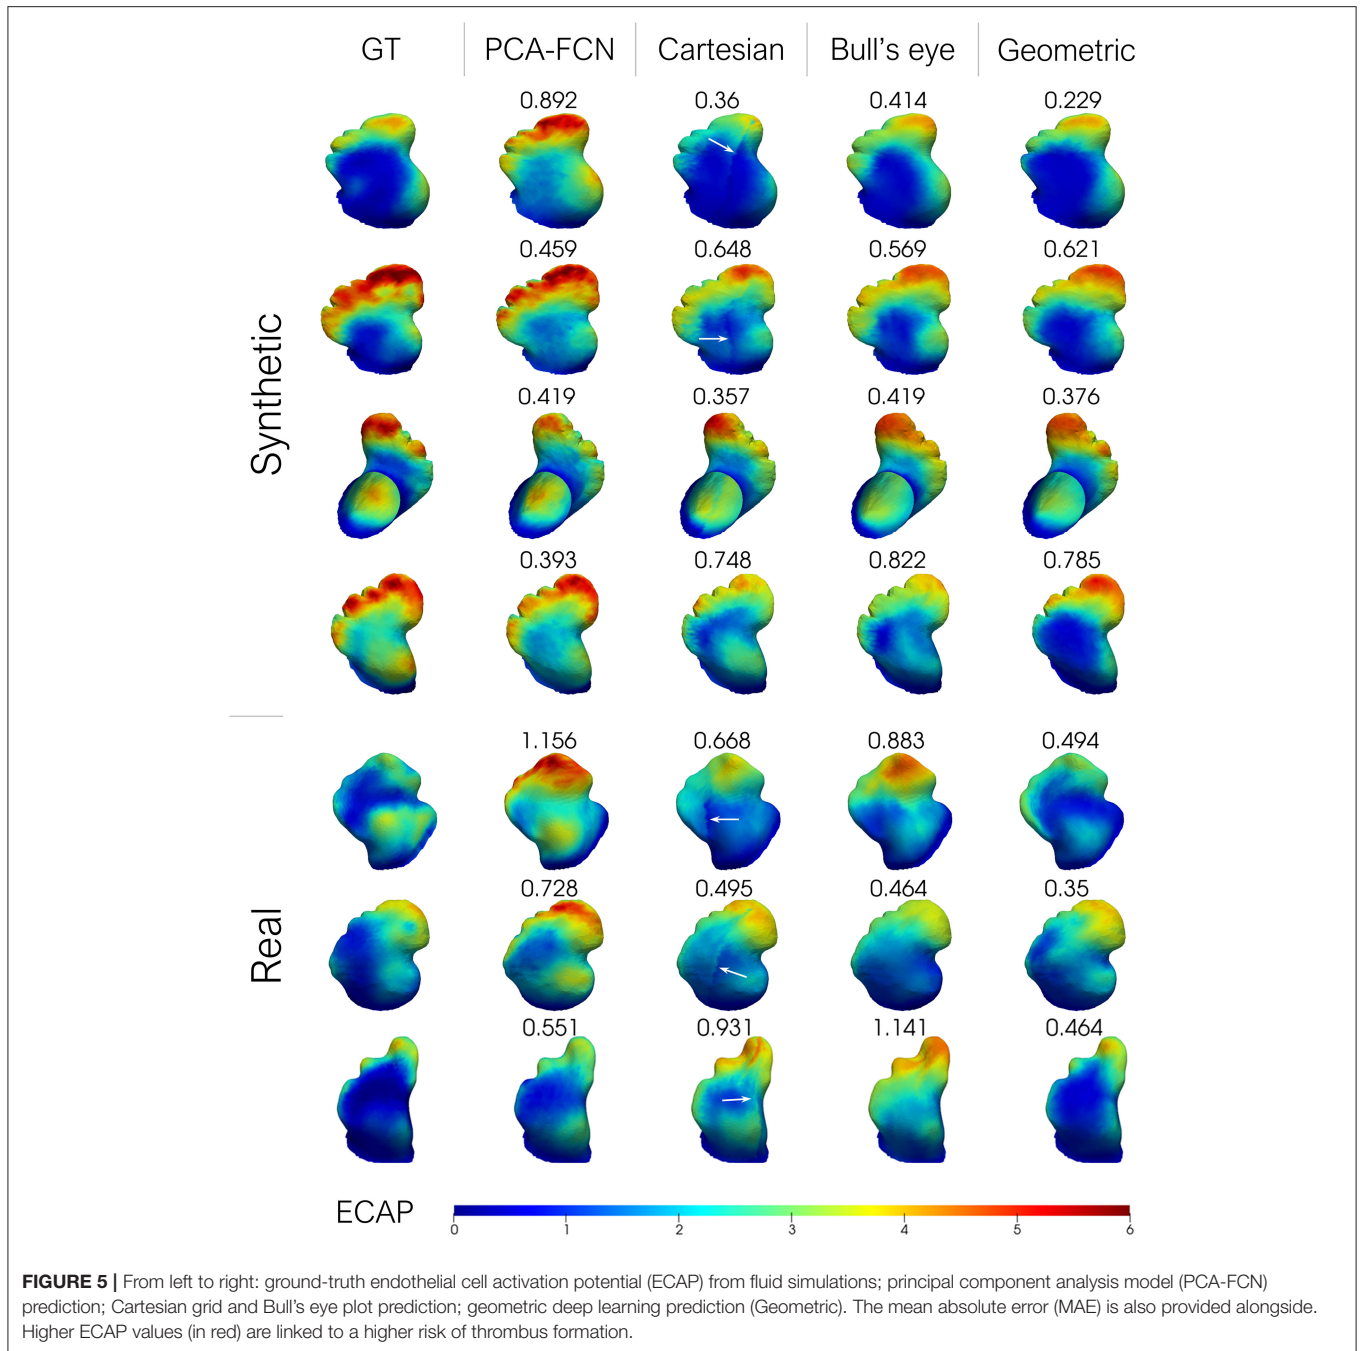
2 simulation, some requiring up to 48 h. Conversely, the PCA model was the fastest training DL network by a long margin, only requiring an average of 2 min to train. The training runs of the remaining two networks (i.e., UV mapping—U-net and Geometric DL) by contrast, lasted around 15–20 min. Once the models were trained, the prediction of ECAP maps pertaining to new unseen cases was instantaneous. On the other hand, the graph-based network was the lightest, with a total of 1.686.097 trainable parameters, in comparison to the 7.846.178 weights in the PCA model, and 9.304.833 for the U-Net.

The accuracy results for the cross-validation experiments are provided in **Table 1**, both in terms of the mean absolute error (MAE) of the ECAP and the true positive error (TPR), that is, the percentage of areas above the 90th percentile that have been predicted as such by the network. The geometric DL network outperformed the remaining approaches in all cross-validation experiments for both metrics. Nonetheless, there is a noticeable disparity between the MAE and classification results, given that even though the cross-validation on real data has provided accuracy on par to the other two scenarios in terms of MAE, it has a significantly lower TPR among all models.

Additionally, a small batch of seven testing geometries from one random fold of the “Cross” cross-validation experiment is shown in **Figure 5**. Cases in rows 1–4 of the figure were derived from the SSM model while the remaining three represent real patient cases. The results from the remaining test samples in the fold are provided as **Supplementary Material**. In order to visually compare with the rest of the approaches, the results derived from the flattening models were interpolated back to the original mesh. Furthermore, the mean absolute error is provided for each prediction instance.

In **Figure 6**, the results from the sequential experiments are provided. Once again the graph-based model outscored its counterparts by some margin. Conversely, the PCA model struggled whenever few data was available as the maximum number of PCs had to be lowered in folds with <32 training instances. Interestingly, the bull's eye representation also had an edge over the rectangular Cartesian grid in the majority of tasks. Furthermore, the addition of synthetic cases in the training dataset for experiment “Sequential Real + Synthetic,” did not improve upon the results of models solely trained on real data for the “Sequential Real” experiment.

Lastly, some exemplary results from the Geometric PointNet on the more complex dataset 2 are showcased in **Figure 7**. The remaining test subjects are provided as **Supplementary Material**.



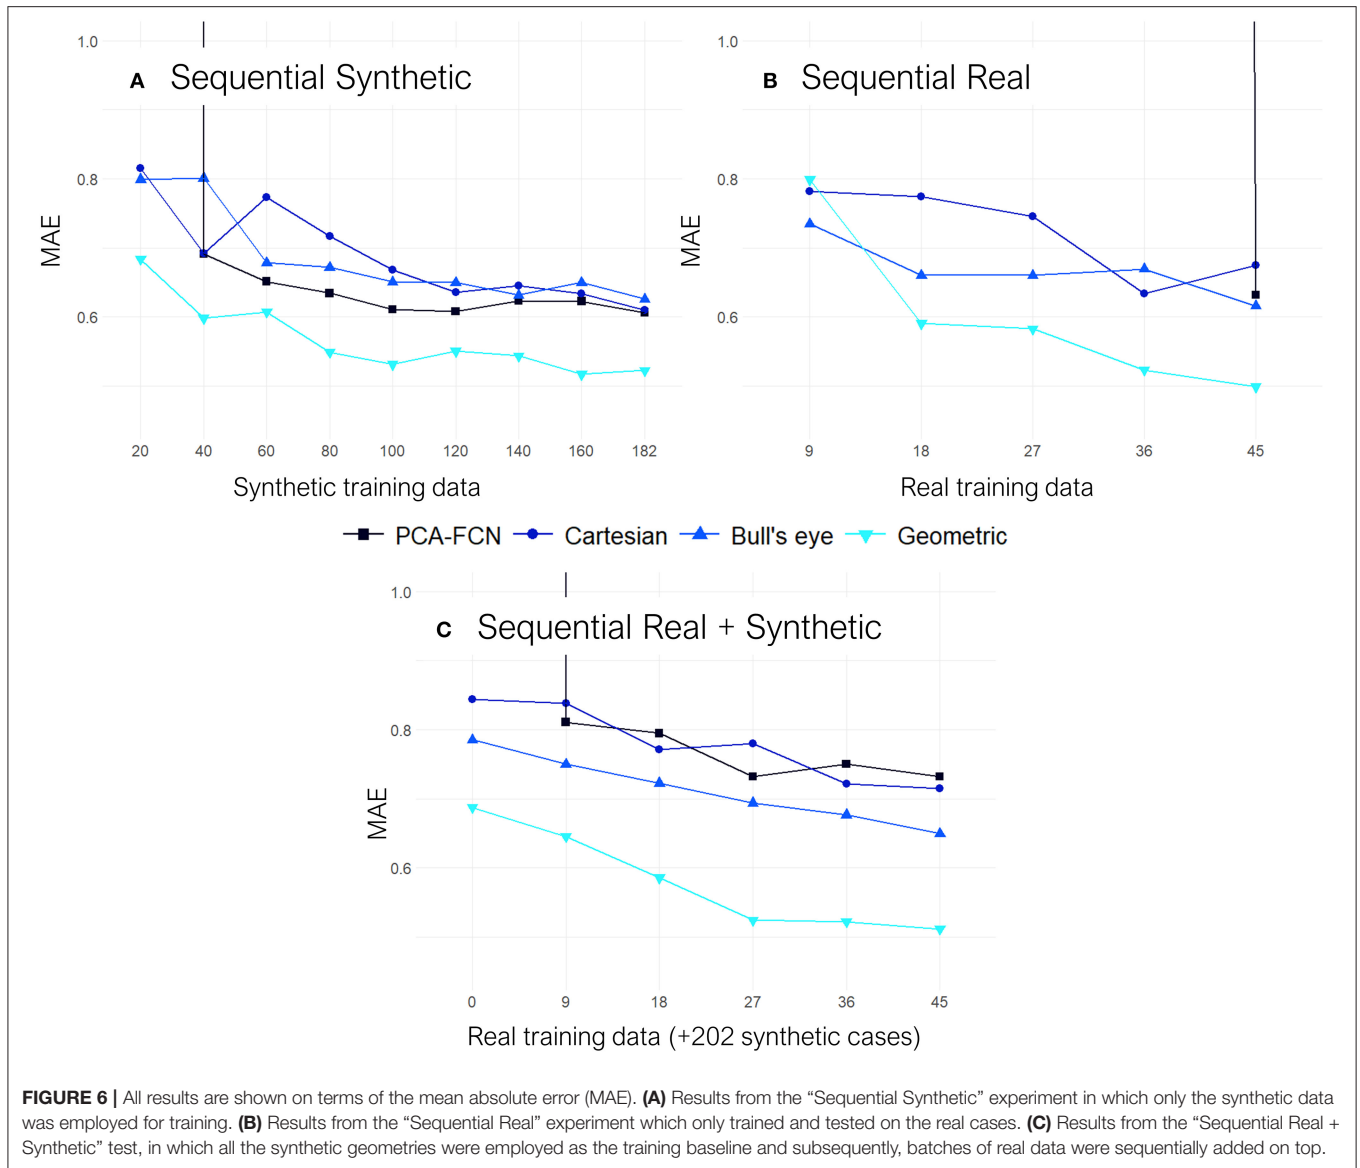
The cross-validation resulted in a MAE = 1,506 ± 0,543, while a TPR of 70% was achieved on the binary classification task.

#### 4. DISCUSSION

The primary goal of this work was to accurately estimate the ECAP, an *in-silico* thrombosis risk index, with a set of distinct deep learning approaches, thus, being able to instantaneously predict the risk indices when presented with new morphologies, without the need of running time-intensive and computationally

demanding simulations. It is evident from **Figures 5, 7** that the developed framework, especially the one based on geometric DL, successfully mimicked the behavior of two distinct sets of CFD simulations with different boundary conditions, managing to capture the global ECAP distributions solely on the basis of LAA morphology. Moreover, the geometric DL model seamlessly extends to realistic data without the need for template registration or 2D mapping. More importantly, the training and prediction of the *in-silico* index were completed orders of magnitude faster than conventional solvers of fluid simulations





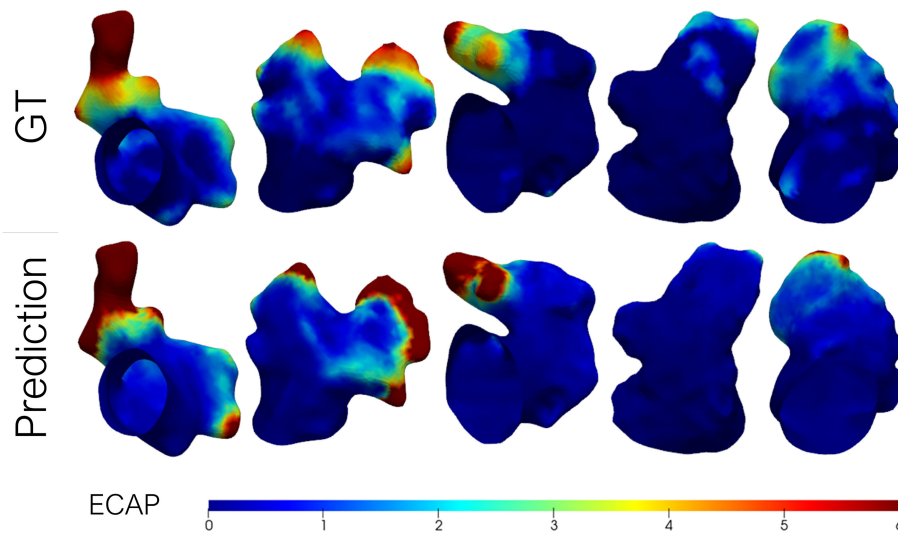
**FIGURE 6** | All results are shown on terms of the mean absolute error (MAE). **(A)** Results from the “Sequential Synthetic” experiment in which only the synthetic data was employed for training. **(B)** Results from the “Sequential Real” experiment which only trained and tested on the real cases. **(C)** Results from the “Sequential Real + Synthetic” test, in which all the synthetic geometries were employed as the training baseline and subsequently, batches of real data were sequentially added on top.

(i.e., tens of minutes vs. several hours). Furthermore, once trained inference can be completed instantaneously.

In this regard, a proper understanding of the data to be learned by the network was imperative in simplifying the space of results and achieving good accuracy. In our case, proper scaling of the data turned out to be crucial in improving the results. For instance, the input tensor to the U-Net model, containing the spatial coordinates of the vertices was standardized. Similarly, power transforming of the curvature data in the graph neural network also offered superior performance. The clearest example, however, involved the ECAP maps obtained from dataset 2 (e.g., with patient-specific LA data) which had a very marked positive skew compared to dataset 1 (e.g., with synthetic LA), rendering the model completely unable to learn. This issue was easily resolved by *log* transforming the ECAP maps for training, resulting in an almost symmetrical data distribution which could

then be reconstructed back to visualize the results. The resulting distribution densities are provided as **Supplementary Material**.

By far, the most laborious and time-consuming aspect of the study consisted in setting up and running the 370 CFD simulations. Several of the steps typically involved in a geometry-specific fluid modeling pipeline (i.e., medical image segmentation, mesh building, the definition of boundary conditions, simulation execution), often necessitate manual intervention (e.g., mesh corrections). This lack of automatization represents a major bottleneck when simulating large datasets, hence most fluid dynamics studies end up including <10 cases when focused on complex morphologies such as the LA and LAA. By automating several of the aforementioned procedures we managed to streamline most of the simulation workflow, thus enabling the formation of a dataset large enough to train neural networks.



**FIGURE 7 |** Predicted endothelial cell activation potential (ECAP) maps in a batch of left atrial appendages (LAA) from dataset 2 (i.e., patient-specific LA) alongside the ground-truth (GT) from fluid simulations. Higher ECAP values (in red) are linked to a higher risk of thrombus formation.

#### 4.1. Dataset 1—Synthetic LA Dataset

Careful inspection of the results presented in **Table 1** indicates that the geometric DL model outperformed all its counterparts in the three designated tasks. A look at **Figure 5** further supports this hypothesis, as the geometric DL network obtains a better accuracy than the rest of the models in 5 out of 7 of the shown cases. Interestingly, despite having a more rudimentary DL architecture, the PCA model was the best non-geometric DL approach, even when trained on the real cases alone. A strong performance of the PCA network on the synthetic dataset was to be expected since the geometries were sampled from a statistical shape model based on the same methodology. Regardless, the PCA model is second only to the geometric DL network on the “Cross real” experiment in terms of TPR and is very close to the Bull’s eye regarding MAE, highlighting the strength of PCA as a shape analysis tool. As for the flattening approaches, the results obtained were ambiguous: while both the Cartesian and Bull’s eye model perform similarly in terms of MAE in the experiments including synthetic data (“Sequential Synthetic” and “Sequential Real + Synthetic”), the circumferential approach generalized far more effectively to the real dataset. In fact, the Cartesian grid method was the only model to worsen its accuracy in the realistic LAAs. A possible explanation may be related to the cut-off introduced in the Cartesian grid representation when performing the flattening, which results in the loss of the original mesh topology. This gives rise to a discontinuity when performing the convolution over the flattened mesh, which produces a very prominent cut (see white arrows in **Figure 5**). As real geometries are far more heterogeneous, the position of the reference 0–360° line marked by the circumflex artery localization might fluctuate more often, which we hypothesize leads to inconsistent learning of the morphological features for the Cartesian method.

On the other hand, the disparity observed between the MAE and the TPR in the real and synthetic cases seems to stem from

a distinct distribution of the data. In this respect, the ECAP in the real LAA dataset has a lower  $\mu = 1.664$  and a 90th percentile equal to 3.313, whereas the synthetic cases have a  $\mu = 2.266$  and a 90th percentile of 4.343. With fewer training data encompassing vertices with ECAP values above the 90th percentile threshold, the model is more prone to fail when confronted with higher ECAP values, leading to far worse TPR (see **Table 1**). Concerning the gap in the distribution of ECAP values between real and synthetic cases, we observed a prevalence of “Cauliflower” like appendages in the synthetic geometries. These morphologies are characterized by having several lobes, such as cases [2,3,4] in **Figure 5**. As observed in these three anatomies, the ECAP in these lobes tends to be quite high, probably due to increased blood stasis, which might explain the disparity in the values of ECAP between the two populations. In turn, this is the reason that seems to jeopardize the potential accuracy improvement from including synthetic data.

Moving on to the sequential experiments, all models keep improving as more training data is added, which suggests that further increasing the size of the simulation datasets could be highly beneficial for the overall accuracy. Similarly to the previous experiments, the geometric DL model continues to exhibit superior performance over all the other approaches. This time, it is the PCA-FCN network that struggles the most, as appreciated in the “Sequential Real” and “Sequential Synthetic” experiments, shown in **Figure 6**, since the amount of initial training data in both of these experiments is well below the 32 principal components that yielded the best results with the PCA model. In fact, the PCA approach was not able to obtain good predictions until the training dataset amounted to about 40 geometries. More interestingly, even though the baseline training dataset already comprised 202 synthetic geometries in the “Sequential Real + Synthetic” experiment, the PCA model did not perform well (i.e.,  $MAE \gg 1$ ) until a minimum number

of real geometries were provided. Finally, the “Sequential Real + Synthetic” experiment suggests that the inclusion of the synthetic data was not of particular help in improving the accuracy of patient-specific LAA. One would expect that, as the amount of real training data increases, the accuracy achieved would eventually exceed that of the model solely trained in the 54 real cases or alternatively, that a similar level of accuracy would be obtained but utilizing a smaller number of real geometries. Neither of these two scenarios held true for any of the models, as accuracy actually worsened overall. Only the geometric DL model managed to achieve accuracy on par with that obtained in the “Sequential Real” experiment, so it appears to have learned more relevant and universal morphological features.

All things considered, the graph-based neural network was superior not only in terms of performance but also regarding the ease of deployment, while the PCA-FCN and flattening models each had their strengths and weaknesses. First, the PCA-FCN model showcased good robustness with regards to real data and it was the faster training model by far. Nonetheless, the need for registration was a major handicap during mesh processing, given that mesh connectivity had to be preserved, which greatly restricted employing tools such as remeshing, vital to avoid mesh quality problems. Not to mention the employed registration itself, which entailed a degree of deformation in the mesh being registered. In regards to flattening, although altogether bypasses the need for registration and template selection, it only succeeded in overcoming the above approach in cases where a very small amount of data was available for training. Besides, UV mapping can not be easily extended to other topologies should we consider including the full LA geometry. All that being said, flattening representations are still very useful for visualization and comparison of large LAA populations. Concerning the geometric DL model, it delivered the best results while working directly over the native form of the data, and using almost an order of magnitude fewer weights than its counterparts. In addition, as no correspondences were required the initial pre-processing was minimal, thus facilitating the editing of the meshes and avoid mesh quality issues. For all these reasons, the graph-based model was chosen for further testing on the second dataset.

## 4.2. Dataset 2—Complete LA Dataset

As aforementioned, an inspection of **Figure 7** shows that the distribution of ECAP maps in dataset 2 differs from the synthetic LA dataset (dataset 1). The more complex boundary conditions used in dataset 2 have strengthened the washout in the proximal portion of the LAA. Only in those recesses and cavities in which the inflow fails to reach, the ECAP is higher than in the first dataset. On the other hand, the incorporation of the entire LA geometry during simulation signifies that the ECAP no longer solely depends on the variation of LAA morphology; other anatomical features such as the orientation of the pulmonary veins (García-Isla et al., 2018) will play a role in shaping the variability of the risk index. Despite the added complexity of the second dataset the geometric DL network effectively learned the abstract set of anatomical features related to blood stagnation. Unfortunately, owing to the “black box” nature of

neural networks, it is difficult to pinpoint what the model is learning, whether it is a combination of the distance to the ostium along with local curvature on a given bulge or some other arrangement of features that might be challenging for humans to grasp. Although the MAE results effectively tripled in this second dataset relative to dataset 1, it was to be expected given the skewed nature of the data and the extremely high values at given spots.

## 4.3. Limitations and Future Work

Despite the promising results, the presented study has several limitations that must be addressed before it can be of any use in a clinical setting.

First, the choice of the ECAP as a thrombosis risk index may be a subject of contention, since its validity has yet to be proven on the LAA. At first, the ECAP was chosen as it provides a dimensionless scalar field that captures some of the most relevant hemodynamic characteristics related to the formation of thrombi in the LAA, which in turn, allows simplifying the required DL model architecture. Moreover, even though the ECAP index was originally developed in carotid and abdominal aorta fluid models (Di Achille et al., 2014), it has already seen some use in clinical studies exploring device-related thrombus formation in LAA occlusion surgeries (Aguado et al., 2019; Mill et al., 2020). In any case, the underlying mechanisms of thrombus formation in the aforementioned situations always involve some degree of blood stagnation or re-circulation at low velocities that the ECAP should be able to grasp to some extent. Furthermore, there is mounting evidence challenging the utility of standard clinical scores such as the CHAD<sub>2</sub>DS<sub>2</sub>-VASc, which has been long held as the main guide for anticoagulation therapy in AF patients, highlighting the need for more advanced risk indexes accounting for AF-specific factors such as hemodynamic alterations (Siddiqi et al., 2021). In this regard, the geometric DL model could seamlessly extend to 3D data allowing to predict more recently adopted indexes of blood stagnation such as the residence time, which offers an approximate measure of blood stagnation time scale, based on LA flow velocity vector fields (García-Villalba et al., 2021).

On the other hand, the hemodynamic variability arising from the heterogeneous anatomy of the LA was completely neglected when training the network for the sake of simplicity. Nonetheless, since the geometric DL framework does not involve any kind of mesh correspondence it should be fairly trivial to include the complete LA anatomy. Moreover, the network should be capable of learning the ECAP fluctuations caused by factors such as the interaction between the orientation of the pulmonary veins. Yet, an increase in the size of the input graph could render the current local convolution scheme insufficient. In this sense, the network could greatly benefit from widespread approaches in computer vision such as strided convolutions or pooling, aimed at extracting multi-scale features. Nonetheless, although we tested several of the available approaches to construct an encoder-decoder-like architecture such as in Hanocka et al. (2019) and Zhou et al. (2020), for the time being, we have not been able to successfully integrate any of them in the graph-based model. Future work should also be focused on the interpretability of the models, as learning the features that the network is focusing

on is a crucial step before being able to deploy the model in a clinical environment.

At the moment, the model is completely agnostic to flow dynamics and distinct boundary conditions that play a key role in the process of thrombogenesis. To address this challenge, we intend on capitalizing on the rapid advances in the field of physics-informed neural networks, with examples such as the study by Pfaff et al. (2021). This may enable the full exploitation by artificial neural networks of the rich Spatio-temporal data available within CFD data, which may pave the way toward the real-time prediction of the full velocity vector field in the LA without the need for hour-long fluid dynamics simulations.

The ground truth from fluid simulations could also be substituted by 4D flow magnetic resonance imaging (MRI), which enables a full non-invasive mapping of the intravascular 3D velocity field over time. Nevertheless, for the time being, most of currently available 4D flow MRI acquisitions employ velocity encodings ( $V_{enc}$ ) better suited to higher velocity vessels and continue to suffer from poor signal-to-noise ratio and spatiotemporal resolution (Jiang et al., 2015). As a result, reliable imaging of the LAA flow field is extremely challenging, especially in the proximity of the vessel wall, making it nearly impossible to obtain accurate values of derived hemodynamic indices such as the wall shear stress or the ECAP (Pettersson et al., 2012). In this regard, attempts have already been made to tackle said limitations such as the development of Dual- $V_{enc}$  acquisition sequences (Callahan et al., 2019) or leveraging CFD simulations to obtain 4D flow super-resolution (Ferdian et al., 2020).

Lastly, to get the full picture of the risk of thrombus formation the inclusion of factors such as endothelial damage/dysfunction could be of particular interest. Scar segmentation in AF patients can be performed automatically by employing neural networks over MRI acquisitions (Li et al., 2020; Yang et al., 2020), allowing detection of left atrium wall fibrosis which is independently associated with LAA thrombogenesis (Guo et al., 2012).

## 5. CONCLUSION

In the present study, we have successfully leveraged a set of deep learning frameworks to instantaneously predict the ECAP mapping in the LAA solely from its anatomical mesh, effectively skipping the need to run CFD simulations at inference time. All models were effective in a simplified LA model, the graph-based geometrical DL network repeatedly outscoring its competitors. Moreover, this same model exhibited good predictive capability even in more advanced simulations with improved boundary conditions and including the entire LA anatomy. These results could lay the foundation for real-time monitoring of LAA thrombosis risk in the future and open exciting avenues for

future preoperative applications and interfaces in which a clinical user could interactively change settings of a left atrial appendage occluder device and instantaneously assess the associated risk of device-related thrombus generation.

## DATA AVAILABILITY STATEMENT

The code to reproduce the experiments of this study is available on [github.com/Xtaltec/DL-based-Estimation-of-Thrombotic-Risk](https://github.com/Xtaltec/DL-based-Estimation-of-Thrombotic-Risk). Simulation data, including solver configurations and results, are available under request.

## ETHICS STATEMENT

The studies involving human participants were reviewed and approved by the ethics Committee of Rigshospitalet, Copenhagen, Denmark and the ethics committee of Hôpital Haut-Lévêque, Bordeaux, France. The patients/participants provided their written informed consent to participate in this study.

## AUTHOR CONTRIBUTIONS

XM and OC designed the research and wrote the manuscript. XI, BL, HC, and OD collected and analyzed the data. XM, JM, KJ, CA, HC, BL, XI, OD, RP, and OC did the research and reviewed the manuscript. All authors contributed to the article and approved the submitted version.

## FUNDING

This work was supported by the Agency for Management of University and Research Grants of the Generalitat de Catalunya under the Grants for the Contracting of New Research Staff Programme—FI (2020-FI-B-00690) and the Spanish Ministry of Economy and Competitiveness under the Programme for the Formation of Doctors (PRE2018-084062), the Maria de Maeztu Units of Excellence Programme (MDM-2015-0502) and the Retos Investigación project (RTI2018-101193-B-I00). Additionally, this work was supported by the H2020 EU SimCardioTest project (Digital transformation in Health and Care SC1-DTH-06-2020; grant agreement No. 101016496).

## SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fphys.2021.694945/full#supplementary-material>

## REFERENCES

Acebes, C., Morales, X., and Camara, O. (2021). “A cartesian grid representation of left atrial appendages for a deep learning estimation of thrombogenic risk

predictors,” in *Statistical Atlases and Computational Models of the Heart* (Lima: Springer International Publishing), 35–43. doi: 10.1007/978-3-030-68107-4\_4  
Aguado, A. M., Olivares, A. L., Yagüe, C., Silva, E., Nuñez-García, M., Fernández-Quilez, Á., et al. (2019). *In silico* optimization of left atrial appendage occluder

- implantation using interactive and modeling tools. *Front. Physiol.* 10:237. doi: 10.3389/fphys.2019.00237
- Benjamin, E. J., Muntner, P., Alonso, A., Bittencourt, M. S., Callaway, C. W., Carson, A. P., et al. (2019). Heart disease and stroke statistics 2019 update: a report from the American heart association. *Circulation* 139, e56–e528. doi: 10.1161/CIR.0000000000000659
- Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., and Vandergheynst, P. (2017). Geometric deep learning: going beyond Euclidean data. *IEEE Signal Process. Mag.* 34, 18–42. doi: 10.1109/MSP.2017.2693418
- Callahan, S., Singam, N. S., Kendrick, M., Negahdar, M., Wang, H., Stoddard, M. F., et al. (2019). Dual-venic acquisition for 4d flow MRI in aortic stenosis with spiral readouts. *J. Magn. Reson. Imaging* 52, 117–128. doi: 10.1002/jmri.27004
- Cerqueira, M. D., Weissman, N. J., Dilsizian, V., Jacobs, A. K., Kaul, S., Laskey, W. K., et al. (2002). Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart. *Circulation* 105, 539–542. doi: 10.1161/hc0402.102975
- Crane, K., Weischedel, C., and Wardetzky, M. (2013). Geodesics in heat. *ACM Trans. Graph.* 32, 1–11. doi: 10.1145/2516971.2516977
- Cresti, A., García-Fernández, M. A., Sievert, H., Mazzone, P., Baratta, P., Solari, M., et al. (2019). Prevalence of extra-appendage thrombosis in non-valvular atrial fibrillation and atrial flutter in patients undergoing cardioversion: a large transeosophageal echo study. *EuroIntervention* 15, e225–e230. doi: 10.4244/EIJ-D-19-00128
- Di Achille, P., Tellides, G., Figueroa, C. A., and Humphrey, J. D. (2014). A haemodynamic predictor of intraluminal thrombus formation in abdominal aortic aneurysms. *Proc. R. Soc. A* 470:2172. doi: 10.1098/rspa.2014.0163
- Dudzinska-Szczerba, K., Michalowska, I., Piotrowski, R., Sikorska, A., Paszkowska, A., Stachnio, U., et al. (2021). Assessment of the left atrial appendage morphology in patients after ischemic stroke—the Assam study. *Int. J. Cardiol.* 330, 65–72. doi: 10.1016/j.ijcard.2021.01.001
- Ferdian, E., Suinesiaputra, A., Dubowitz, D. J., Zhao, D., Wang, A., Cowan, B., et al. (2020). 4DflowNet: super-resolution 4d flow MRI using deep learning and computational fluid dynamics. *Front. Phys.* 8:138. doi: 10.3389/fphys.2020.00138
- Fernández-Pérez, G., Duarte, R., Corral de la Calle, M., Calatayud, J., and Sánchez González, J. (2012). Analysis of left ventricular diastolic function using magnetic resonance imaging. *Radiologia* 54, 295–305. doi: 10.1016/j.rxeng.2011.09.003
- Fey, M., Lenssen, J. E., Weichert, F., and Müller, H. (2018). “SplineCNN: Fast geometric deep learning with continuous b-spline kernels,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Salt Lake City, UT). doi: 10.1109/CVPR.2018.00097
- García-Isla, G., Olivares, A. L., Silva, E., Nuñez-García, M., Butakoff, C., Sanchez-Quintana, D., G. et al. (2018). Sensitivity analysis of geometrical parameters to study haemodynamics and thrombus formation in the left atrial appendage. *Int. J. Numer. Methods Biomed. Eng.* 34:e3100. doi: 10.1002/cnm.3100
- García-Villalba, M., Rossini, L., Gonzalo, A., Vigneault, D., Martínez-Legazpi, P., Durán, E., et al. (2021). Demonstration of patient-specific simulations to assess left atrial appendage thrombogenesis risk. *Front. Physiol.* 12:49. doi: 10.3389/fphys.2021.596596
- Grandits, T., Pezzuto, S., Lubrecht, J. M., Pock, T., Plank, G., and Krause, R. (2021). “Piemap: Personalized inverse eikonal model for cardiac electro-anatomical maps,” in *Statistical Atlases and Computational Models of the Heart. Me&Ms and EMIDEC Challenges*, eds E. Puyol Anton, M. Pop, M. Sermesant, V. Campello, A. Lalande, K. Lekadir, A. Suinesiaputra, O. Camara, and A. Young (Cham: Springer International Publishing), 76–86. doi: 10.1007/978-3-030-68107-4\_8
- Guo, Y., Lip, G. Y., and Apostolakis, S. (2012). Inflammation in atrial fibrillation. *J. Am. Coll. Cardiol.* 60, 2263–2270. doi: 10.1016/j.jacc.2012.04.063
- Hanocka, R., Hertz, A., Fish, N., Giryas, R., Fleishman, S., and Cohen-Or, D. (2019). MeshCNN. *ACM Trans. Graph.* 38, 1–12. doi: 10.1145/3306346.3322959
- Hennigh, O. (2017). Lat-Net: compressing lattice boltzmann flow simulations using deep neural networks. *arXiv:1705.09036*.
- Jiang, J., Kokeny, P., Ying, W., Magnano, C., Zivadinov, R., and Mark Haacke, E. (2015). Quantifying errors in flow measurement using phase contrast magnetic resonance imaging: comparison of several boundary detection methods. *Magn. Reson. Imaging* 33, 185–193. doi: 10.1016/j.mri.2014.10.009
- Li, G., Wang, H., Zhang, M., Tupin, S., Qiao, A., Liu, Y., et al. (2021). Prediction of 3d cardiovascular hemodynamics before and after coronary artery bypass surgery via deep learning. *Commun. Biol.* 4:99. doi: 10.1038/s42003-020-01638-1
- Li, L., Wu, F., Yang, G., Xu, L., Wong, T., Mohiaddin, R., et al. (2020). Atrial scar quantification via multi-scale CNN in the graph-cuts framework. *Med. Image Anal.* 60:101595. doi: 10.1016/j.media.2019.101595
- Liang, L., Liu, M., Martin, C., and Sun, W. (2018). A deep learning approach to estimate stress distribution: a fast and accurate surrogate of finite-element analysis. *J. R. Soc. Interface* 15:20170844. doi: 10.1098/rsif.2017.0844
- Masci, A., Alessandrini, M., Forti, D., Menghini, F., Dedé, L., Tomasi, C., et al. (2019). A proof of concept for computational fluid dynamic analysis of the left atrium in atrial fibrillation on a patient-specific basis. *J. Biomech. Eng.* 142:011002. doi: 10.1115/1.4044583
- Meister, F., Passerini, T., Audigier, C., Lluch, É., Mihalef, V., Ashikaga, H., et al. (2021). “Graph convolutional regression of cardiac depolarization from sparse endocardial maps,” in *Statistical Atlases and Computational Models of the Heart. Me&Ms and EMIDEC Challenges*, eds E. Puyol Anton, M. Pop, M. Sermesant, V. Campello, A. Lalande, K. Lekadir, A. Suinesiaputra, O. Camara, and A. Young (Lima: Springer International Publishing), 23–34. doi: 10.1007/978-3-030-68107-4\_3
- Mill, J., Olivares, A., Silva, E., Genua, I., Fernández, Á., Aguado, A., et al. (2019). “Optimal boundary conditions in fluid simulations for predicting occlude related thrombus formation in the left atria,” in *Computational and Mathematical Biomedical Engineering, Sixth International Conference* (Sendai).
- Mill, J., Olivares, A. L., Arzamendi, D., Agudelo, V., Regueiro, A., Camara, O., et al. (2020). Impact of flow dynamics on device-related thrombosis after left atrial appendage occlusion. *Can. J. Cardiol.* 36, 968.e13–968.e14. doi: 10.1016/j.cjca.2019.12.036
- Morales, X., Mill, J., Juhl, K. A., Olivares, A., Jimenez-Perez, G., Paulsen, R. R., et al. (2020). “Deep learning surrogate of computational fluid dynamics for thrombus formation risk in the left atrial appendage,” in *Statistical Atlases and Computational Models of the Heart*, eds M. Pop, M. Sermesant, O. Camara, X. Zhuang, S. Li, A. Young, T. Mansi, and A. Suinesiaputra (Shenzhen: Springer International Publishing), 157–166. doi: 10.1007/978-3-030-39074-7\_17
- Nordestgaard, B. G., Palmer, T. M., Benn, M., Zacho, J., Tybjaerg-Hansen, A., Smith, G. D., et al. (2012). The effect of elevated body mass index on ischemic heart disease risk: causal estimates from a mendelian randomisation approach. *PLoS Med.* 9:e1001212. doi: 10.1371/journal.pmed.1001212
- Petersson, S., Dwyerfeldt, P., and Ebbens, T. (2012). Assessment of the accuracy of MRI wall shear stress estimation using numerical simulations. *J. Magn. Reson. Imaging* 36, 128–138. doi: 10.1002/jmri.23610
- Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., and Battaglia, P. (2021). “Learning mesh-based simulation with graph networks,” in *International Conference on Learning Representations* (Vienna).
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. (2016). Pointnet: deep learning on point sets for 3d classification and segmentation. *CoRR, abs/1612.00593*. doi: 10.1109/CVPR.2017.16
- Ronneberger, O., Fischer, P., and Brox, T. (2015). “U-net: convolutional networks for biomedical image segmentation,” in *Lecture Notes in Computer Science*, eds N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi (Munich: Springer International Publishing), 234–241. doi: 10.1007/978-3-319-24574-4\_28
- Siddiqi, T. J., Usman, M. S., Shahid, I., Ahmed, J., Khan, S. U., Ya’qoub, L., et al. (2021). Utility of the CHA2DS2-VASc score for predicting ischaemic stroke in patients with or without atrial fibrillation: a systematic review and meta-analysis. *Eur. J. Prevent. Cardiol.* doi: 10.1093/eurjpc/zwab018. [Epub ahead of print].
- Slipsager, J. M., Juhl, K. A., Sigvardsen, P. E., Kofoed, K. F., De Backer, O., Olivares, A. L., et al. (2019). “Statistical shape clustering of left atrial appendages,” in *Statistical Atlases and Computational Models of the Heart*, eds M. Pop, M. Sermesant, J. Zhao, S. Li, K. McLeod, A. Young, K. Rhode, and T. Mansi (Granada: Springer International Publishing), 32–39. doi: 10.1007/978-3-030-12029-0\_4
- Slodowska, K., Szczepanek, E., Dudkiewicz, D., Holda, J., Bolechała, F., Strona, M., et al. (2021). Morphology of the left atrial appendage: introduction of a new simplified shape-based classification system. *Heart Lung Circul.* doi: 10.1016/j.hlc.2020.12.006. [Epub ahead of print].
- Tewari, A., Zollhofer, M., Kim, H., Garrido, P., Bernard, F., Perez, P., et al. (2017). “MoFA: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction,” in *2017 IEEE International*

- Conference on Computer Vision Workshops (ICCVW)* (Venice). doi: 10.1109/ICCVW.2017.153
- Thuerey, N., Weißenow, K., Prantl, L., and Hu, X. (2020). Deep learning methods for reynolds-averaged navier–stokes simulations of airfoil flows. *AIAA J.* 58, 25–36. doi: 10.2514/1.J058291
- Veronesi, F., Corsi, C., Sugeng, L., Caiani, E. G., Weinert, L., Mor-Avi, V., et al. (2008). Quantification of mitral apparatus dynamics in functional and ischemic mitral regurgitation using real-time 3-dimensional echocardiography. *J. Am. Soc. Echocardiogr.* 21, 347–354. doi: 10.1016/j.echo.2007.06.017
- Vieira, M. J., Teixeira, R., Gonçalves, L., and Gersh, B. J. (2014). Left atrial mechanics: echocardiographic assessment and clinical implications. *J. Am. Soc. Echocardiogr.* 27, 463–478. doi: 10.1016/j.echo.2014.01.021
- Watson, T., Shantsila, E., and Lip, G. Y. (2009). Mechanisms of thrombogenesis in atrial fibrillation: Virchow’s triad revisited. *Lancet* 373, 155–166. doi: 10.1016/S0140-6736(09)60040-4
- Yaghi, S., Chang, A. D., Akiki, R., Collins, S., Novack, T., Hemendinger, M., et al. (2020). The left atrial appendage morphology is associated with embolic stroke subtypes using a simple classification system: a proof of concept study. *J. Cardiovasc. Comput. Tomogr.* 14, 27–33. doi: 10.1016/j.jcct.2019.04.005
- Yang, G., Chen, J., Gao, Z., Li, S., Ni, H., Angelini, E., et al. (2020). Simultaneous left atrium anatomy and scar segmentations via deep learning in multiview information with attention. *Fut. Gener. Comput. Syst.* 107, 215–228. doi: 10.1016/j.future.2020.02.005
- Zhou, Y., Wu, C., Li, Z., Cao, C., Ye, Y., Saragih, J., et al. (2020). “Fully convolutional mesh autoencoder using efficient spatially varying kernels,” in *Advances in Neural Information Processing Systems*, Vol. 33, eds H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Vancouver, BC: Curran Associates, Inc.), 9251–9262.

**Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2021 Morales Ferez, Mill, Juhl, Acebes, Iriart, Legghe, Cochet, De Backer, Paulsen and Camara. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.