

Full-field surrogate modeling of cardiac electrophysiology encoding geometric variability

Martinez, Elena Sabdy; Moscoloni, Beatrice; Salvador, Matteo; Kong, Fanwei; Peirlinck, Mathias; Marsden, Alison Lesley

DOI

[10.1016/j.cma.2025.118444](https://doi.org/10.1016/j.cma.2025.118444)

Publication date

2026

Document Version

Final published version

Published in

Computer Methods in Applied Mechanics and Engineering

Citation (APA)

Martinez, E. S., Moscoloni, B., Salvador, M., Kong, F., Peirlinck, M., & Marsden, A. L. (2026). Full-field surrogate modeling of cardiac electrophysiology encoding geometric variability. *Computer Methods in Applied Mechanics and Engineering*, 448, Article 118444. <https://doi.org/10.1016/j.cma.2025.118444>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.



Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

**Green Open Access added to [TU Delft Institutional Repository](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.

Otherwise as indicated in the copyright section:
the publisher is the copyright holder of this work and the
author uses the Dutch legislation to make this work public.



Full-field surrogate modeling of cardiac electrophysiology encoding geometric variability

Elena Saby Martinez ^{id} ^{a,b}, Beatrice Moscoloni ^{id} ^{c,d}, Matteo Salvador ^{id} ^{a,b,e,f},
Fanwei Kong ^{id} ^{a,b,e}, Mathias Peirlinck ^{id} ^c, Alison Lesley Marsden ^{id} ^{a,b,e,g,*}

^a Institute for Computational and Mathematical Engineering, Stanford University, CA, USA

^b Cardiovascular Institute, Stanford University, CA, USA

^c Department of BioMechanical Engineering, Delft University of Technology, Delft, the Netherlands

^d IBiTech-bioMMeda, Ghent University, Ghent, Belgium

^e Pediatric Cardiology, Stanford University, CA, USA

^f Pasteur Labs, Brooklyn, NY, 11205, USA

^g Department of Bioengineering, Stanford University, CA, USA

ARTICLE INFO

Keywords:

Computational cardiology
Scientific machine learning
Statistical shape modeling
Branched latent neural maps
Tetralogy of fallot
Cardiac electrophysiology
Surrogate modeling

ABSTRACT

Combining physics-based modeling with data-driven methods is critical to enabling the translation of computational methods to clinical use in cardiology. The use of rigorous differential equations combined with machine learning tools allows for model personalization with uncertainty quantification in time frames compatible with clinical practice. However, accurate and efficient surrogate models of cardiac function, built from physics-based numerical simulation, are still mostly geometry-specific and require retraining for different patients and pathological conditions. We propose a novel computational pipeline to embed cardiac anatomies into full-field surrogate models. We generate a dataset of electrophysiology simulations using an accurate multi-scale mathematical model coupling partial and ordinary differential equations. We adopt Branched Latent Neural Maps (BLNMs) as a computationally efficient scientific machine learning method to encode activation maps extracted from physics-based numerical simulations into a neural network. Leveraging large deformation diffeomorphic metric mappings, we build a biventricular anatomical atlas and parametrize the anatomical variability of a small and challenging cohort of 13 pediatric patients affected by Tetralogy of Fallot. We propose a novel z-score sampling approach based on statistical shape modeling to generate a new synthetic cohort of 52 biventricular geometries that are compatible with the original geometrical variability. This synthetic cohort acts as the training set for BLNMs. Our surrogate model demonstrates robustness and great generalization across the complex original patient cohort, achieving an average adimensional mean squared error of 0.0034. The Python implementation of our BLNM model is publicly available under MIT License at <https://github.com/StanfordCBCL/BLNM>.

1. Introduction

For decades, cardiologists have based treatment decisions on a “one-size-fits-all” assumption that patients with standardized symptoms must share the same disease pathophenotype and thus exhibit similar responses to the conventional treatment plan [1]. However, the significant variation among heart physiology, geometry, and cell-to-organ scale properties often leads to diverse responses to the

* Corresponding author.

E-mail address: amarsden@stanford.edu (A.L. Marsden).

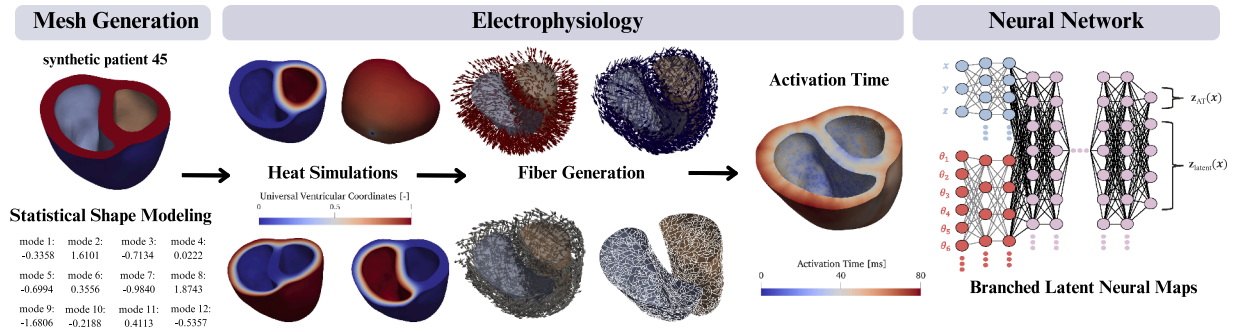


Fig. 1. Sketch of the computational pipeline illustrated for synthetic patient 45. We use SSM to generate a synthetic patient by sampling 12 modes. For this patient, we conduct heat simulations to define some universal ventricular coordinates [38]. Next, we create the corresponding fiber fields: fiber sheets (top left), longitudinal fibers (top right), normal fibers (bottom left), and Purkinje fibers (bottom right). We use these fiber orientations to run an electrophysiology simulation and generate the activation map. BLNMs predict activation maps by using two branches for modes and spatial coordinates, respectively.

same interventions, particularly in pediatric cardiology [2]. This variability has important implications for modeling approaches. For example, action potential models are often derived from averaged experimental data, producing a “typical” model that largely fails to accurately capture the variability observed across individual action potentials, even within a homogeneous population [3]. This limitation becomes particularly crucial in clinical practice. Treating cardiac arrhythmias using a general model rather than an individualized approach can result in incorrect defibrillation energy settings or medication doses [4,68]. This underscores the critical need for personalized models to improve patient safety and treatment efficacy [5,6].

Physics-based mathematical models can incorporate both patient-specific anatomic imaging and clinical data. These models employ rigorous differential equations to simulate cardiac dynamics while accounting for individual variations in heart structure [7]. Integrating data-driven methods with physics-based models enables the inclusion of diverse datasets obtained from various clinical modalities, such as imaging (e.g MRI or CT scans), electrophysiology recordings (e.g. ECG or EEG), and hemodynamic measurements (e.g blood flow or pressure data). This accelerates parameter calibration and facilitates more informed and timely decision-making in clinical settings [8–10]. Recent works demonstrate the efficacy of applying machine learning techniques to multi-scale cardiac models tailored to individual patients [11–18]. Various methods have aimed to reduce the computational cost of physics-based simulations while achieving accurate and efficient model performance specifically for electrophysiology [19]. For example, reduced-order models (ROMs) based on artificial neural networks (ANNs) can accurately generate active force in patient-specific geometries while maintaining computational efficiency [20]. Deep learning-based reduced order models (DL-ROMs) for cardiac electrophysiology can efficiently evaluate outputs of clinical interest such as activation maps and action potentials [13]. By incorporating the PDEs associated with activation mapping and neural networks, Physics-Informed Neural Networks (PINNs) can create activation maps based on a few sparse activation measurements [21]. These models, however, often overlook the significant interpatient variability in cardiac geometries, a crucial factor for accurate scalable predictions [22].

Statistical shape modeling (SSM) has been extensively applied to characterize anatomical variability across various cardiovascular applications [23–28]. Typically, SSM construction involves first mapping a cohort of anatomical meshes onto a common template to ensure vertex-to-vertex correspondence among subjects. Subsequently, dimensionality reduction techniques - most commonly principal component analysis (PCA) - are applied to derive a compact, low-dimensional representation of the cohort. PCA identifies an average anatomical configuration and principal components (or modes) that capture the majority of the anatomical variability present within the dataset. Following PCA, each anatomical mesh can be reconstructed as a linear combination of these modes, enabling the encoding of individual geometrical variability through mode-specific weights. Recently, several studies have combined statistical shape modeling with advanced sampling strategies to generate large synthetic cohorts of cardiac surface meshes [29–34]. The generation of synthetic anatomical cohorts holds significant potential for in silico studies and surrogate model training. However, realizing this potential critically depends on well-designed sampling strategies and rigorous acceptance criteria to ensure that the synthetically generated cohorts exclude anatomies that are physiologically implausible or improbable [29,35].

This paper proposes a novel computational pipeline for predicting activation maps from electrophysiology simulations by embedding cardiac anatomies into surrogate models using Branched Latent Neural Maps (BLNMs). BLNMs use feedforward partially-connected neural networks to separate contributions from distinct inputs, thereby maintaining training/inference computational efficiency and scalable memory footprints while accurately learning complex input-output relationships with significant variability [36]. These surrogate models are full-field, meaning they predict complete activation maps that represent the spatial distribution of electrical activity at each point of the computational domain. Our approach uniquely integrates innovative SSM approaches with BLNMs to capture the geometric and functional variability of highly variant cardiac geometries.

We consider a cohort of 13 pediatric patients affected by Tetralogy of Fallot. To expand our dataset, we generate a synthetic cohort of 52 additional biventricular geometries that align with the original distribution. For both the original and synthetic cohort, we produce a dataset of electrophysiology simulations using a multi-scale mathematical model based on differential equations. Using BLNMs, we encode the activation maps extracted from these physics-based numerical simulations. These activation maps represent

the timing of electrical activation in cardiac tissue and are key in guiding medical treatments and procedures for various cardiac conditions [37]. By testing this framework on a challenging cohort of patients, we demonstrate its robustness in handling diverse anatomies while preserving computational efficiency and without needing model modification (Fig. 1).

2. Methods

We introduce a methodological framework for predicting cardiac activation maps in highly heterogeneous biventricular anatomies by simulating cardiac electrophysiology driven by anatomical form. Although this framework is applicable to a wide range of healthy and diseased heart morphologies, we specifically highlight its application to a challenging cohort of Tetralogy of Fallot (ToF) patients. ToF is a complex congenital heart disease (CHD) characterized by four key anatomical abnormalities: a ventricular septal defect, an overriding aorta, stenosis of the right ventricular outflow tract, and right ventricular hypertrophy. Given the complex nature and varying severity of these defects, ToF presents an ideal pathophysiology to showcase the capabilities of our computational pipeline [39]. This process can be divided into three parts: first, generating a cohort of 52 synthetic geometries derived from 13 ToF patients; second, creating geometry-specific cardiac activation maps; and third, using Branched Latent Neural Maps to predict activation maps for different geometries.

2.1. Original cohort

We utilize ImageCHD, a dataset consisting of 110 3D computed tomography (CT) images that capture a wide range of congenital heart diseases in patients between 1 month and 40 years of age [40]. For this study, we focus specifically on the subset of patients with Tetralogy of Fallot (ToF). We generate high-resolution anatomical meshes from the ground truth myocardium segmentations of these 3D CT images. We focus on modeling only the ventricular region below the ventricular septal defect (VSD) in these patients to concentrate on areas most impacted by altered conduction patterns, while also reducing computational demand. Since our electrophysiology simulations focus on this particular region, we developed automated Python scripts to crop each myocardium mesh at a manually identified plane, perpendicular to the long axis of the ventricles and identify appropriate boundary faces. This process produces meshed biventricular geometries, encompassing the left and right ventricles, epicardium (the outer surface of the heart), and base (the plane where the ventricles and atria connect). To standardize and refine the meshes, we utilize the meshing tools available in SimVascular, an open-source software package specifically designed for cardiovascular modeling and simulation [41]. By using SimVascular we are able to remesh each geometry in the cohort with a consistent resolution while also ensuring seamless integration into our simulation pipeline.

We study a small and challenging cohort of $N_o = 13$ patients characterized by substantial anatomical variability. The myocardial tissue volumes across these biventricular geometries demonstrate substantial variation, with patient 5 exhibiting the highest volume at approximately 84.75 mL and patient 7 the lowest volume at 41.75 mL. We also observe significant variability in the endocardial surface areas of both ventricles, which ranged from 15.80 cm² (patient 7) to 35.37 cm² (patient 4) for the left ventricle (LV), and from 12.04 cm² (patient 7) to 51.24 cm² (patient 6) for the right ventricle (RV). The largest difference in surface area between the LV and RV is 16.62 cm², while the smallest difference is 0.31 cm². In our cohort, the LV is larger in 5 out of 13 patients, whereas the RV is larger in the remaining 8 (Fig. 2).

In addition to variability in overall size, the cohort also demonstrated significant variability in shape. The myocardial wall thickness is defined as the shortest distance from any point on the endocardium (the surface of the LV and RV) to the nearest point on the epicardium, or the nearest point on the endocardium of the opposite ventricle when it is closer. This definition ensures that we accurately capture the thickness of the septum. We calculate the average wall thickness by averaging the shortest distances of all myocardial points in a geometry. Across our cohort of patients, the average wall thickness varies from 0.52 cm to 1.16 cm. Furthermore, we define the heights of the biventricular geometries as the distance from the apex (the tip of the epicardial surface) to its nearest point on the base. The biventricular heights vary across geometries, ranging from 0.50 cm to 6.6 cm. These measurements emphasize the substantial anatomical variability across patients in both wall thickness and apicobasal (from apex to base) height (Fig. 2).

2.2. Synthetic cohort

We extend the original cohort by combining anatomical mapping, geometrical encoding, and a novel sampling strategy to generate a synthetic cohort that is consistent with the characteristics of the original patient cohort described in the previous section. Fig. 3 presents an overview of our synthetic cohort generation pipeline.

2.2.1. Large deformation diffeomorphic metric mapping

We perform mapping of the $N_o = 13$ anatomies of the original cohort using the large deformation diffeomorphic metric mapping (LDDMM) approach detailed in Moscoloni et al. [28]. In this approach, we map a set of target patient-specific biventricular geometries to a template biventricular geometry by means of a dynamic flow of smooth, invertible transformations called diffeomorphisms. More specifically, these diffeomorphisms $\Phi_t : \mathbb{R}^3 \mapsto \mathbb{R}^3$ entail a set of deformation vectors $(\mu_n)_{n=1, \dots, N_q}$, known as momenta, which are applied to an ambient space of N_q equally spaced control points $(\mathbf{q}_n)_{n=1, \dots, N_q}$. We keep the initial set of control points fixed

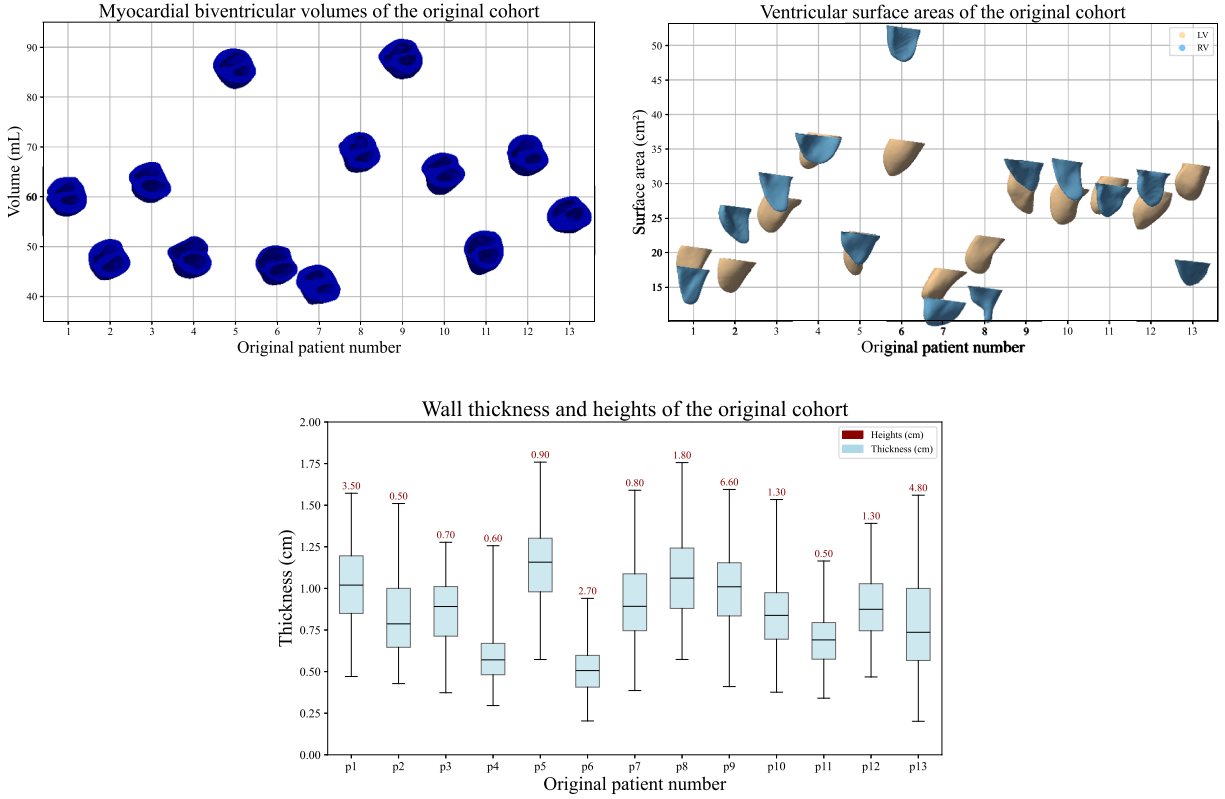


Fig. 2. Comparison of myocardial volumes, ventricular surface areas, and structural metrics (wall thickness and apicobasal heights) across the original patient cohort. Although each mesh is normalized to a common frame for visualization, the cohort exhibits variability in both size (volumes, surface areas) and shape (wall thickness, septal morphology).




Fig. 3. Overview of the synthetic cohort generation pipeline. We apply large deformation diffeomorphic metric mapping to map original anatomies onto a common template, followed by principal component analysis to encode the anatomical variability into z-scores. Using a novel sampling strategy, we generate synthetic anatomies based on their probabilistic proximity to the original samples in the multivariate space defined by these geometrical encodings. Specifically, we selected the 4 closest samples (dotted lines) to each original subject (colored lines).

throughout the large deformation diffeomorphic metric mapping optimization procedure. Consequently, each target geometry defined as a surface mesh Γ_i , is approximated by their unique set of momenta μ_i deforming the template geometry Γ :

$$\Gamma_i \approx \Phi_i(\Gamma) = \Phi_{\mathbf{q}, \mu_i}(\Gamma) \tag{1}$$

We construct diffeomorphisms Φ_i by mapping each point from an initial template geometry Γ to a corresponding point on a target geometry Γ_i . This mapping is defined through a spatiotemporally varying velocity field $\mathbf{v}(\mathbf{x}, t) = \mathbf{v}_t(\mathbf{x})$ over a *pseudo-time* $t \in [0, 1]$ [28,42]. We express the velocity field explicitly at each space location \mathbf{x} and time t as:

$$\mathbf{v}_t(\mathbf{x}) = \sum_{k=1}^{N_q} K_V(\mathbf{x}, \mathbf{q}_k(t)) \mu_k(t), \tag{2}$$

where K_V is a Gaussian kernel $K_V(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|^2 / \lambda_V^2)$ of kernel width λ_V . This kernel weighs the contribution of momenta $\mu_k(t)$ applied at control points $\mathbf{q}_k(t)$ to the instantaneous movement of the template vertex at position \mathbf{x} .

The control points themselves evolve dynamically under the same velocity field (Eq. (2)). Their trajectories follow Hamiltonian equations of motion [43], governed explicitly by their momenta $\mu(t)$:

$$\begin{cases} \dot{\mathbf{q}}(t) = K_V(\mathbf{q}(t), \mathbf{q}(t)) \cdot \mu(t) \\ \dot{\mu}(t) = -\frac{1}{2} \nabla_{\mathbf{q}} \{ K_V(\mathbf{q}(t), \mathbf{q}(t)) \cdot \mu(t)^\top \mu(t) \} \end{cases} \quad (3)$$

Given an initial set of control points $\mathbf{q}(0)$, we select time-varying momenta $\mu(t)$ that define deformation paths Φ_t that bring the template geometry as close to the target geometry of interest as possible at $t = 1$. Since multiple possible deformations paths exist, we prioritize paths that minimize the kinetic energy, resulting in geodesic trajectories. These geodesics are fully parametrized by the initial momenta $\mu_0 = \mu(0)$ [28,44].

To estimate the optimal template geometry Γ and deformation fields $\Phi_{\mathbf{q}, \mu_i}$, we minimize the following cost function with respect to the initial template geometry and the set of initial momenta:

$$C(\Gamma, \mu) = \sum_{i=1}^{N_\sigma} \frac{1}{\sigma^2} d_W(\Phi_{\mathbf{q}, \mu_i}(\Gamma), \Gamma_i)^2 + \sum_{i=1}^{N_\sigma} \mu_i^\top K_W(\mathbf{q}, \mathbf{q}) \mu_i, \quad (4)$$

where the first term minimizes the sum of the Varifold distances d_W between the target geometries Γ_i and their approximations through $\Phi_{\mathbf{q}, \mu_i}$, while the second term penalizes deformations with high kinetic energy. Here, σ balances the relative contributions of these two terms.

We optimize Eq. (4) iteratively through two alternating steps: (i) We first optimize the momenta μ for each target geometry, keeping the template fixed Γ . This step identifies the best deformation momenta to match the current template to each individual target geometry. (ii) Next, we optimize the template Γ itself, while keeping the previously obtained momenta μ fixed. This step updates the template to better represent the average of the current approximations of the target geometries derived from the current set of subject-specific momenta μ . We employ a gradient descent scheme for both optimization steps, and repeat these steps until convergence of the initial momenta and template is achieved [28,42,43].

After optimization, the final estimated template geometry Γ along with subject-specific momenta μ_i for each subject provides a complete mapping of each subject's anatomy within the cohort. Applying these subject-specific momenta to the vertices of the optimized template mesh explicitly reconstructed individual anatomical surface meshes. In practice, this mapping process entails the selection of an initial template, to which the surface meshes need to be aligned, and two user-defined hyperparameters that influence the accuracy of the mapping: the Gaussian kernel width of λ_V from Eq. (3), and the grid spacing parameter λ_W defining the density of the control point grid. Large λ_V values promote global, smooth deformations, while large λ_W values allow capturing fine anatomical details within target meshes [45].

2.2.2. Statistical shape modeling pipeline

Mesh alignment and anatomical mapping. To apply the LDDMM strategy outlined in the previous section to our cohort, we first align the target meshes and extract an initial template geometry. Starting with a cohort of N_σ surface meshes, we align the meshes in two steps. First, we align each mesh to an initial reference mesh, chosen as the mesh with the lowest number of vertices, using the iterative closest point (ICP) algorithm between each mesh and the initial target [46]. Briefly, ICP finds the optimal rigid body transformation minimizing the distances between corresponding points based on spatial proximity. Following this preliminary alignment, we select the medoid mesh - the instance which is closest to the cohort's mean - as our updated reference mesh and perform a second ICP alignment to this updated reference mesh, which subsequently serves as the initial template for the anatomical mapping. With this initial template selected, we perform LDDMM anatomical mapping using Deformetrica [43]. We optimize two hyperparameters, related to the Gaussian kernel width λ_V and the control point grid resolution λ_W (see Section 2.2.1). Following hyperparameter tuning, we set $\lambda_V = \lambda_W = 6$ mm. This hyperparameter choice yields a total of 1100 control points, which adequately balances global deformation and local anatomical detail. The initial template, control points, and model hyperparameters are used to estimate the final template mesh and patient-specific momenta for the cohort. The combination of control points, optimized template, and patient-specific momenta allows us to reconstruct each patient-specific geometry with consistent point-to-point correspondence.

Evaluating reconstruction accuracy. We quantify the accuracy of our anatomical mapping through reconstruction errors, measured as the average vertex-to-vertex distance between original mesh and reconstructed meshes. Specifically, let $(\mathbf{p})_{j=1, \dots, N_v}$ be the set of vertices in the original mesh for patient i , and $(\mathbf{r})_{j=1, \dots, N_v}$ be the corresponding vertices in the reconstructed mesh, where N_v is the total number of points. We define the reconstruction error ℓ_i for patient i as:

$$\ell_i = \frac{1}{N_v} \sum_{j=1}^{N_v} \|\mathbf{p}(j)_i - \mathbf{r}(j)_i\| \quad (5)$$

where $\|\mathbf{p}(j)_i - \mathbf{r}(j)_i\|$ is the Euclidean distance between the j th vertex point on the original mesh and the j th vertex point on the reconstructed mesh for patient i . The overall reconstruction accuracy for the cohort of N_σ patients is quantified by the average total

reconstruction error ℓ_{tot} across all individuals:

$$\ell_{\text{tot}} = \frac{1}{N_{\mathcal{O}}} \sum_{i=1}^{N_{\mathcal{O}}} \ell_i \quad (6)$$

Principal component analysis and shape encoding. To capture the anatomical variability of the cohort, we perform principal component analysis (PCA) on the spatial coordinates of the corresponding vertices across the cohort [25,28]. PCA decomposes the geometry variations into a mean geometry and a set of principal components, allowing each patient's geometry to be represented compactly as:

$$\Gamma_i = \bar{\Gamma} + \sum_{m=1}^{N_M} \alpha_{i,m} \phi_m \lambda_m, \quad i \in \{1, 2, \dots, N_{\mathcal{O}}\}, \quad (7)$$

where $\bar{\Gamma}$ is the mean geometry across the cohort, ϕ_m and λ_m are respectively the eigenvector and eigenvalue of the m th principal component, and $\alpha_{i,m}$ is the patient specific weight, or shape score, of that component in the i th patient. This PCA-based representation allows efficient encoding of anatomical variability using only the shape scores $\alpha_{i,m}$. We found that using 12 principal component-derived z-scores sufficiently captures the anatomical variability within the patient cohort. For the sampling and training of the surrogate model (Section 2.4), we normalized the shape coefficients as z-scores by their mean and standard deviation. Furthermore, we can generate new synthetic patient geometries by sampling novel shapes scores $\beta_{i,n}$ and plugging them into Eq. (7) [29,45].

2.2.3. Synthetic cohort generation

To augment our original cohort of $N_{\mathcal{O}}$ anatomies with additional synthetic yet realistic data, we introduce a novel three-step pipeline based on geometrical encoding and probabilistic sampling. This synthetic cohort generation procedure includes: (i) sampling new candidate shape scores and filtering them based on the probability density function of the original cohort, (ii) identifying candidate shape scores that are the closest to the original patients in the multivariate space, (iii) reconstructing and meshing synthetic anatomies from selected shape scores

Probabilistic sampling and filtering. We employ Latin Hypercube Sampling (LHS) [47] to generate synthetic candidate shape scores within a 12-dimensional multivariate space, with values ranging from $[-2.5, 2.5]$ standard deviations in each mode of anatomical variation. We sample a large set of samples to ensure broad coverage of anatomical variations. To ensure compatibility of synthetic anatomies with our original cohort, we filter these candidate shape scores based on their probability relative to the original cohort of $N_{\mathcal{O}}$ anatomies. Specifically, we estimate the underlying probability density function using a non-parametric Gaussian Kernel Density Estimation (KDE) [48] applied to the original z-scores:

$$\hat{f}(\mathbf{z}) = \frac{1}{nh_1 h_2 \dots h_d} \sum_{i=1}^n K\left(\frac{\mathbf{z} - \mathbf{Z}_i}{h}\right), \quad (8)$$

where $\hat{f}(\mathbf{z})$ is the estimated density function at point \mathbf{z} , n is the number of data points in the original cohort, \mathbf{Z}_i are the sample points from the original data, h_1, h_2, \dots, h_d are the bandwidths for each dimension, determining the width of the kernel, and $K(u)$ is the Gaussian kernel function.

This method allows us to model the distribution of the original z-scores without assuming a specific parametric form. We apply this KDE to the z-scores derived from the original cohort of $N_{\mathcal{O}}$ anatomies, to derive the probability distribution across the multivariate space defined by the modes of variation. Using this estimated density, we define a threshold probability corresponding to the least probable sample in the original cohort. We filter out any LHS-generated candidate scores falling below this threshold, thus ensuring that our generated samples are at least as probable to occur as the samples we observe in the original cohort and that our synthetic cohort accurately reflects realistic anatomical variations observed in our original population.

Selecting closest synthetic samples. Following probabilistic filtering, we identify synthetic generated samples that closely resemble actual patient geometries. We first transform both the original and synthetic normalized z-scores back to their original scales using the respective means and standard deviations of each mode. This approach ensures that the distances between shape scores reflect actual anatomical variations appropriately.

We then compute the Euclidean distances in the multivariate space between each pair of original shape scores $\alpha_{i,n}^{\text{orig}}$ and the filtered synthetic shape scores $\beta_{j,n}^{\text{gen}}$:

$$\ell_{ij} = \sqrt{\sum_{n=1}^N (\alpha_{i,n}^{\text{orig}} - \beta_{j,n}^{\text{gen}})^2} \quad (9)$$

where N is the total number of modes considered. For each original patient, we rank these synthetic candidates by their proximity, selecting the four closest matches. This step enables us to focus further on the samples that fall not only within the statistically probable range but are also geometrically close to the original data points, thus effectively enhancing the diversity yet maintaining anatomical realism in the synthetic cohort.

Reconstruction and mesh generation. We reconstruct synthetic surface meshes from the selected shape scores by applying Eq. (7), resulting in a cohort of $N_S = 52$ synthetic surface meshes. These surface meshes are further processed using Meshmixer [49] to fix small topological defects that may arise during reconstruction. We subsequently generate volumetric meshes for each synthetic patient using Gmsh [50]. Finally, to ensure full compatibility with our original cohort and integration with our simulation pipeline, we remesh all synthetic volumetric geometries using SimVascular.

2.3. Cardiac electrophysiology

We introduce the mathematical model and the numerical scheme that we adopt to generate activation sequences of the biventricular geometries. We simulate cardiac electrophysiology for both the synthetic and original patient cohorts by generating a total of $N_T = 65$ activation maps.

2.3.1. Mathematical model

We model cardiac electrophysiology in multiple 3D domains $\Omega^i = \Omega_{\text{BiV}}^i \cup \Omega_{\text{purk}}^i \subset \mathbb{R}^3$ represented by biventricular-Purkinje models of original and synthetic patients affected by ToF using the monodomain equation [14,51,52] coupled with the ten Tusscher-Panfilov ionic model [53], that is:

$$\begin{cases} \frac{\partial u}{\partial t} + I_{\text{ion}}(u, \mathbf{y}) - \nabla \cdot (\mathbf{D}_M \nabla u) = I_{\text{app}}(\mathbf{x}, t) & \text{in } \Omega^i \times (0, T], \\ (\mathbf{D}_M \nabla u) \cdot \mathbf{n} = 0 & \text{on } \partial\Omega^i \times (0, T], \\ \frac{d\mathbf{y}}{dt} = \mathbf{F}(u, \mathbf{y}) & \text{in } \Omega^i \times (0, T], \\ u(\mathbf{x}, 0) = u_0(\mathbf{x}), \mathbf{y}(\mathbf{x}, 0) = \mathbf{y}_0(\mathbf{x}) & \text{in } \Omega^i, \end{cases} \quad (10)$$

where $i = 1, \dots, N_T$, where $N_T = N_O + N_S$ represents the total count of biventricular geometries in the combined original and synthetic cohorts. The primal variable u defines the transmembrane potential, that is the electric signal propagating throughout the cardiac tissue, whereas vector $\mathbf{y} = (y_1, \dots, y_{M+P})$ introduces both the probability density functions of $M = 12$ gating variables, which represent the fraction of open channels across the membrane of a single cardiomyocyte, and the concentrations of $P = 6$ relevant ionic species, such as intracellular calcium Ca^{2+} , sodium Na^+ and potassium K^+ .

The right hand side $\mathbf{F}(u, \mathbf{y})$ defines a system of ordinary differential equations for the time evolution of gating and concentration variables. The ionic current $I_{\text{ion}}(u, \mathbf{y})$ enables cell-to-organ scale coupling between the ionic and monodomain models. The analytical expressions of both $\mathbf{F}(u, \mathbf{y})$ and $I_{\text{ion}}(u, \mathbf{y})$ derive from the mathematical formulation of the ten Tusscher-Panfilov ionic model [53]. The diffusion tensor is expressed as $\mathbf{D}_M = D_{\text{iso}}\mathbf{I} + D_{\text{ani}}\mathbf{f}_0 \otimes \mathbf{f}_0$ in Ω_{BiV} and $\mathbf{D}_M = D_{\text{purk}}\mathbf{I}$ in Ω_{purk} , where \mathbf{f}_0 defines the biventricular fiber field [14,54,55]. $D_{\text{ani}}, D_{\text{iso}}, D_{\text{purk}} \in \mathbb{R}^+$ represent the anisotropic, isotropic and Purkinje conductivities, respectively [4]. We impose homogeneous Neumann boundary conditions on $\partial\Omega^i$, where \mathbf{n} is the outward unit normal vector to the boundary, to prescribe an electrically isolated domain.

The action potential is triggered by a current $I_{\text{app}}(\mathbf{x}, t)$ applied at time $t = 0$ at the first nodes of both the left and right bundle branches. These nodes initiate the depolarization of the right and left ventricle Purkinje networks, which we generate using a fractal tree algorithm. This algorithm effectively covers the complex and irregular surfaces inside the ventricles through a two-step projection process [56]. We construct a fractal network by defining branch parameters, including initial length, generation count, median length, length standard deviation, angle, and a repulsion parameter which regulates curvature. Currently, no in vivo imaging technique can accurately capture Purkinje networks [57], which exhibit both significant inter- and inpatient variability in size and shape. We utilize a coverage-based approach to achieve a consistent and reproducible method to automatically create Purkinje networks across geometries.

As a first step in this approach, we methodically select two points in each ventricle to determine the direction of the first branch of each Purkinje network, as shown in Fig. 4. This selection is guided by heat simulations previously conducted to activate different regions of the biventricular geometries. From the apex-to-base heat simulation, we extract all points where the universal ventricular coordinate ϕ_{ab} reaches values of 0.99 and 0.97 for each ventricle. Next, we calculate the centroids of all the $\phi_{ab} = 0.99$ points for the LV and RV separately and draw a line connecting these two centroids. The starting points for the left and right Purkinje networks are then identified as the $\phi_{ab} = 0.99$ points closest to this line in the LV and RV respectively. For each network, the second point is designated as the $\phi_{ab} = 0.97$ point closest to its respective initial point. This second point does not necessarily correspond to the second node in the network but rather serves to determine the direction of the first segment. This systemic approach, ensures that the selected points are geometrically aligned and standardizes the process of determining the direction of the initial branch for each electrophysiology simulation.

For all geometries, we set median branch length $l = 1.8$ mm, standard deviation $\sigma = \sqrt{0.2} * l$, branch angle $\alpha = 0.2$ radians, and branch repulsion $w = 0.15$. We iterate the initial length and branch generation count, starting at 6.5 mm and 22 generations, respectively, until our algorithm produces a Purkinje network where at least 82% of all endocardial points in the ventricle are no more than 0.5 mm away from a node in the created network. Thus, chamber size and morphology directly influence branch placement and coverage.

In the absence of ECG data for our geometries, we adopt parameter values from a previous study that offered representative ECG profiles for CHD cases [36]. We apply a stimulus with amplitude $-52.0 \mu\text{A}/\text{cm}^2$ to the initial nodes of each of the left and right Purkinje networks for 5.0 ms and a cycle length of 800.0 ms. We conduct all electrophysiology simulations using a 3D monodomain

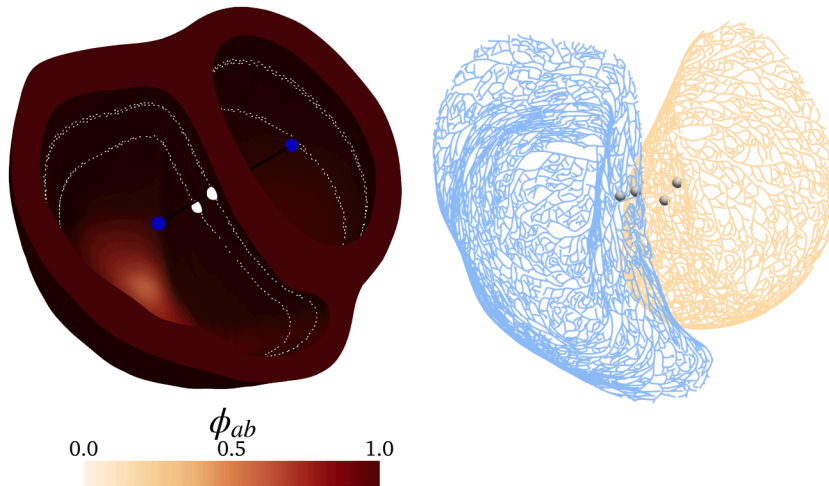


Fig. 4. Selection of the initial points for the Purkinje networks. On the left, the white contour points represent universal ventricular coordinates, $\phi_{ab} = 0.97$ and $\phi_{ab} = 0.99$, derived from the apex-to-base simulation of the Laplace heat equation [38]. The centroids of the ϕ_{ab} points for the LV and RV are marked by blue spheres. The first and second points in each ventricle are represented by white spheres, where the second point is chosen to determine the direction of the first segment. On the right, the corresponding RV and LV Purkinje networks are visualized in light blue and orange, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1
Cardiac electrophysiology simulation parameters.

Parameter	Value	Units
Maximal current conductance of Na^+	13.19	nS pF ⁻¹
Maximal current conductance of Ca^{2+}	7.91e-05	cm ms ⁻¹ μF^{-1}
Maximal current conductance of the rapid delayed rectifier I_{Kr}	0.11	nS pF ⁻¹
Myocardial anisotropic conductivity	0.0177	mm ² ms ⁻¹
Myocardial isotropic conductivity	0.0059	mm ² ms ⁻¹
Purkinje fiber anisotropic conductivity	1.60	mm ² ms ⁻¹
Purkinje fiber isotropic conductivity	1.60	mm ² ms ⁻¹

model with biventricular and Purkinje fiber domains. Simulations are performed for each geometry in our original and synthetic patient cohorts using svMultiPhysics, a high-performance C++ solver in the SimVascular project designed for multi-physics and multi-scale finite element simulations of the cardiovascular system [58]. The workflow from running heat simulations to generating the Purkinje network is performed on a single CPU. For electrophysiology simulations, we utilize a single node from the Stanford Research Computing Center cluster per simulation, with each node featuring 24 cores. Table 1 summarizes the key parameters used in the simulations, including conductance values for ionic currents and domain-specific tissue conductivities.

2.3.2. Numerical discretization

We perform space discretization of the monodomain equation coupled with the ten Tusscher-Panfilov ionic model using the Finite Element Method (FEM) with \mathbb{P}_1 Finite Elements. The tetrahedral mesh of the biventricular-Purkinje system has an average mesh size $h = 1$ mm, for each patient. We apply non-Gaussian quadrature rules to recover convergent conduction velocities of the biventricular-Purkinje system [2].

For time discretization, we employ an Implicit-Explicit numerical scheme where we first update the variables of the ionic model and then the transmembrane potential [59–61]. In particular, the diffusion term in the monodomain equation is treated implicitly and the ionic term is treated explicitly. Furthermore, we use the ionic current discretization by means of the Ionic Current Interpolation scheme [62]. We employ the Forward Euler scheme with a time step size $\Delta t = 0.1$ ms and we simulate ventricular depolarization only by considering $T = 100$ ms, for each patient.

We apply a Laplace-Dirichlet Rule-Based algorithm [38] to generate the fiber, sheet and sheet-normal distributions on the myocardium using the following parameters: $\alpha_{\text{epi}} = -60^\circ$, $\alpha_{\text{endo}} = 60^\circ$, $\beta_{\text{epi}} = 20^\circ$ and $\beta_{\text{endo}} = -20^\circ$.

2.3.3. Activation maps

We generate activation maps from simulation outputs consisting of 1000 VTK files, each representing a 0.1 ms time step over a total duration of 100 ms. This fine temporal resolution is essential to accurately capture the rapid propagation of the cardiac action potential, which occurs over just a few milliseconds. By sampling at 0.1 ms intervals, we ensure sufficient temporal resolution to identify precise activation patterns. We process each file sequentially to extract the action potential at each node in the mesh. As soon as the action potential of a node exceeds the depolarization threshold of -10 mV, we record the corresponding simulation time as

the activation time for that node. This depolarization threshold indicates the physiological transition from resting to active. We then construct an activation map by applying these activation times as scalar values to the 3D cardiac model. An individual activation map is generated for each original and synthetic patient, i.e. 65 activation maps in total.

2.4. Branched latent neural maps

We construct a surrogate model of cardiac function that generalizes across different anatomies by building a feedforward partially-connected neural network with Tanh activation functions. The model is trained to learn activation maps post-processed from the numerical solution of our physics-based electrophysiology model while structurally separating the role of the space coordinates \mathbf{x} and the anatomical parameters θ_{BiV}^i . This recently proposed scientific machine learning tool [36] allows for different levels of disentanglement between inputs and outputs. The surrogate model reads:

$$\mathbf{z}(\mathbf{x}) = \mathcal{B}\mathcal{L}\mathcal{N}\mathcal{M}(\mathbf{x}, \theta_{\text{BiV}}^i; \mathbf{w}) \text{ for } \mathbf{x} \in \Omega^i, i \in \{1, \dots, N_T\}. \quad (11)$$

The set $\mathbf{w} \in \mathbb{R}^{N_w}$ encodes weights and biases of a feedforward partially-connected NN, which represents a map $\mathcal{B}\mathcal{L}\mathcal{N}\mathcal{M} : \mathbb{R}^{3+N_M} \rightarrow \mathbb{R}^{N_z}$ from space coordinates $\mathbf{x} \in \mathbb{R}^3$ and model parameters $\theta_{\text{BiV}}^i \in \Theta \subset \mathbb{R}^{N_M}$ to an output vector $\mathbf{z}(t) = [z_{\text{AT}}(t), \mathbf{z}_{\text{latent}}(t)]^T \in \mathbb{R}^{N_z}$. In this application, the vector θ_{BiV} represents the anatomical z-scores of a biventricular geometry. When the model utilizes all 12 extracted anatomical z-scores, $N_M = 12$ and the vector is defined as $\theta_{\text{BiV}} = [\theta_1, \theta_2, \dots, \theta_{12}]$.

$\mathbf{z}(t)$ contains activation times $z_{\text{AT}}(\mathbf{x}) \in \mathbb{R}$ on the different geometries Ω^i . In addition, the vector leverages some latent variables $\mathbf{z}_{\text{latent}}(t)$, which function as auxiliary internal variables that expand the feature space and enhance the learned space dynamics. Conditioning predictions on these latent variables enables the network to locally adapt to regions with sharp geometric features or steep activation gradients. This helps the network capture subtle yet physiologically relevant conduction patterns that would be difficult to approximate with the physical outputs alone. The inclusion of latent variables and the use of partial connections improves the network's ability to represent complex activation patterns while also producing a smoother landscape of local minima compared to a fully connected network. For this reason, we optimize neural network parameters \mathbf{w} by using the Adam optimizer rather than a second-order optimizer [36]. This allows for faster training while accounting for similar performance in accuracy with respect to more complex optimizers [36].

During nonlinear optimization, we track the mean squared error (MSE) produced by the surrogate model with respect to physics-based numerical simulations to find an optimal set of weights and biases \mathbf{w} , that is:

$$\mathcal{L}(\tilde{z}_{\text{AT}}(\tilde{\mathbf{x}}), \tilde{z}_{\text{obs}}(\tilde{\mathbf{x}}); \hat{\mathbf{w}}) = \arg \min_{\hat{\mathbf{w}}} \left[\|\tilde{z}_{\text{AT}}(\tilde{\mathbf{x}}) - \tilde{z}_{\text{obs}}(\tilde{\mathbf{x}})\|_{L^2(\Omega^i)}^2 \right] \text{ for } i \in \{1, \dots, N_T\}. \quad (12)$$

where $\tilde{z}_{\text{AT}}(t) \in [-1, 1]$ represents BLNM outputs and $\tilde{z}_{\text{obs}}(t) \in [-1, 1]$ defines the post-processed activation maps from physics-based numerical simulations, both in non-dimensional form. Space $\tilde{\mathbf{x}} \in [-1, 1]^3$ and z-scores modes $\tilde{\theta}_{\text{BiV}}^i \in [-1, 1]^{N_M}$ are also normalized during the training and inference. For further details on the BLNM method, we refer the reader to [36]. Furthermore, the Python implementation of our BLNM model is publicly available under MIT License at <https://github.com/StanfordCBCL/BLNM>.

2.4.1. Preparation of training and testing sets

We generated activation maps for 52 synthetic and 13 actual ToF patient geometries. Across our 65 geometries, sizes vary significantly, with the smallest biventricular mesh containing 166,026 points and the largest 372,117 points. The average number of points per geometry is approximately 265,394, with a standard deviation of 42,033 points. Prior to training, we normalize both spatial coordinates and z-scores for each geometry to ensure consistency across different scales. Our training set consists of synthetic geometries, while our testing set contains the original ToF patient cohort.

2.4.2. BLNM hyperparameter tuning

We use Ray Tune, a Python library designed to optimize the search for the best model hyperparameters. We explore 100 neural network configurations, where each configuration is generated by selecting hyperparameters within predefined ranges. Each hyperparameter, except for the learning rate, is sampled from a discrete set of values. These ranges are listed in the Potential Values column of Table 2. The learning rate is sampled from a logarithmic uniform distribution between 10^{-4} and 10^{-3} . We employ OptunaSearch to efficiently identify the hyperparameters that yield the lowest MSE. OptunaSearch is a hyperparameter optimization algorithm within Ray Tune that leverages Optuna, a flexible and efficient optimization library [63]. This sampling process is repeated throughout the 1,000-epoch training period, with model performance evaluated every ten epochs.

The optimal BLNM configuration consists of 16 layers with 42 neurons per layer, 5 states (3 spatial parameters, 12 anatomical variability modes, and 4 latent variables), and a disentanglement level of 2 (Table 2). Fig. 5 provides a visualization of this structure. During each epoch, we randomly select 1485 points from our training geometries as spatial coordinates. Model performance is evaluated every ten epochs, with the entire training process spanning a total of 1000 epochs.

3. Results

We present the mapping results, including the reconstruction accuracy achieved after our anatomical mapping. We also present the activation maps generated from our cardiac electrophysiology simulations for both the original and synthetic cohort. Finally, we showcase the results of our BLNM model trained on the synthetic cohort. Our findings demonstrate that BLNMs can accurately predict cardiac activation times across ToF geometries with substantial anatomical variability, achieving an average mean squared error of 0.0034 on the testing set.

Table 2
Ranges and Ray Tune-optimized parameter values for BLNM.

BLNM Parameters	Potential Values	Optimal Values
Layers	[5, 500]	16
Neurons	[5, 20]	42
Learning Rate	$\log U(10^{-4}, 10^{-3})$	0.0005
# States	[1, 5]	5
Disentanglement	[1, 4]	2
# of Training Points	[1500, 10000]	1485




Fig. 5. Sketch of the optimal BLNM architecture of 16 layers with 42 neurons per layer and 5 states as determined using Ray Tune with the OptunaSearch algorithm. The network starts with separate branches: the first branch processes spatial coordinates (x, y, z) , and the second processes the geometry-specific modes θ_1 through θ_{12} . These branches operate independently up to a specified disentanglement level of 2 before merging.

3.1. Anatomical mapping results and synthetic cohort

Our anatomical mapping process yields an average reconstruction error of 0.086 mm across the cohort, with individual patient reconstruction errors ranging from 0.055 to 0.11 mm. This low geometric reconstruction error indicates a high accuracy in capturing subject-specific anatomies and high fidelity of anatomical correspondence across subjects using our methodology. Additionally, since the reconstruction error is significantly below both imaging and mesh resolution, we consider our approximation trustworthy for geometrical encoding. Hence, we apply PCA to the mapped geometries, choosing a finite set of $N_{\mathcal{M}} = 12$ principal modes of variation. This choice effectively captures essential anatomical features while reducing dimensional complexity, which facilitates efficient sampling and subsequent analysis.

Using Latin Hypercube Sampling, we generate an extensive pool of 10,000 candidate samples. After filtering these samples according to the estimated probability density function from the original cohort, we retain 294 realistic synthetic samples. Fig. 6 visualizes these results using parallel plots. The solid colored lines represent the ensemble of z-scores from the original patient cohort, while the dashed lines illustrate the selected synthetic samples closest to each patient. For further clarity, we also provide overlays between each original anatomy and its four nearest synthetic counterparts, highlighting the anatomical similarities and variations captured by the synthetic data generation method.

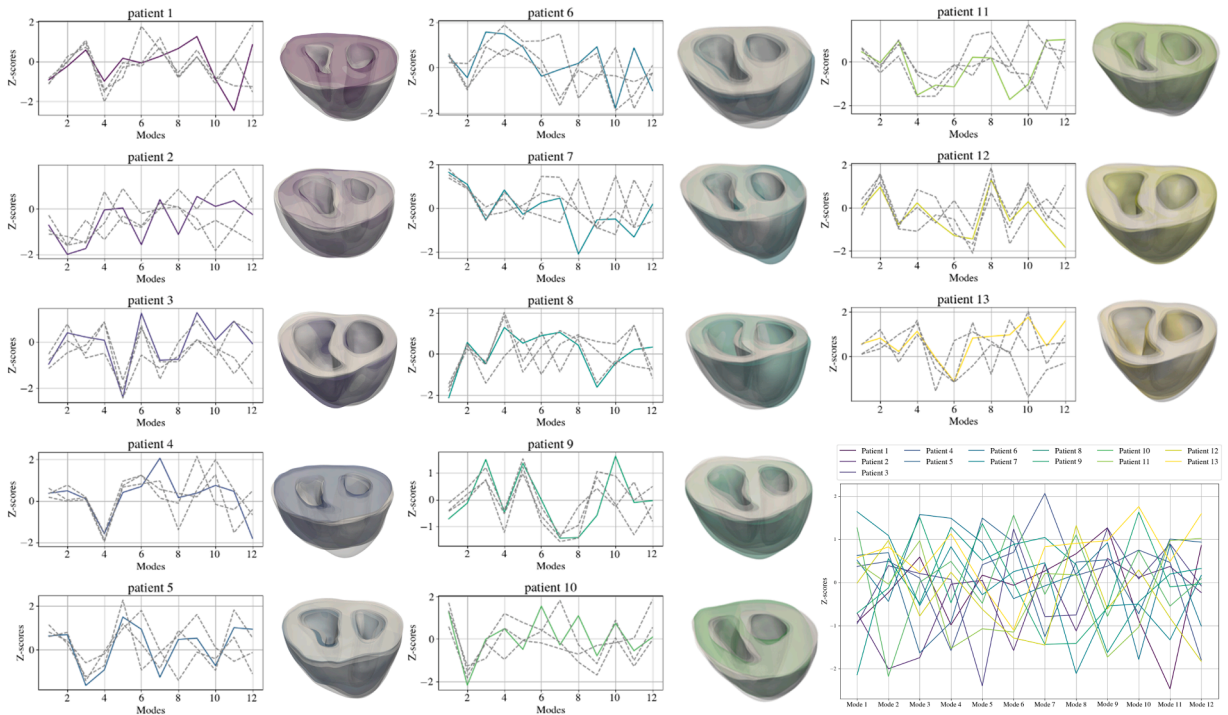


Fig. 6. Parallel plots and overlay meshes comparing original and generated patients. Solid colored lines represent original cohort z-scores, while dashed lines indicate generated synthetic samples closest to each subject.




Fig. 7. Action potential progression for synthetic patient 19 at sequential time points: $t = 0.1$ ms, $t = 16.1$ ms, $t = 32.3$ ms, $t = 64.5$ ms.

3.2. Electrophysiology simulations

From our electrophysiology simulations, we can measure the time required for complete depolarization across the entire biventricular geometry. A position (x,y,z) is considered depolarized when its action potential exceeds -10 mV. Activation time is the time a node takes to depolarize.

Figs. 7 and 8 illustrate the action potentials at different time steps for these patients. These figures depict the action potential at 0.1 ms, and approximately $1/4$, $1/2$, and the full duration required to achieve complete depolarization, highlighting the temporal progression of depolarization across the biventricular geometry. In the synthetic patient cohort, the shortest duration recorded for complete depolarization of the geometry was 64.5 ms (patient 19), while the longest was 92.7 ms (patient 7). In the original patient cohort, the durations for complete biventricular depolarization ranged from 67 ms (patient 3) to 88 ms (patient 10).

This variation in depolarization times highlights the differences in electrophysiological behavior across patient geometries, with each geometry exhibiting a unique depolarization rate. To focus on the impact of anatomical differences, we intentionally used the same electrophysiological parameters across all geometries, ensuring that variations in depolarization dynamics could be attributed solely to structural factors. However, the resulting activation maps have not been clinically validated against direct patient measurements. The physics-based model used in this study serves as a reference for assessing activation patterns.

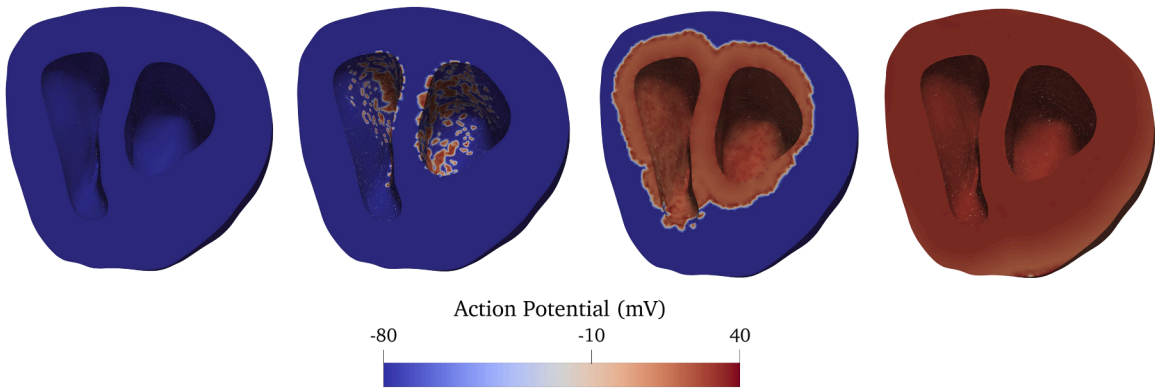


Fig. 8. Action potential progression for synthetic patient 7 at sequential time points: $t = 0.1$ ms, $t = 23.2$ ms, $t = 46.4$ ms, $t = 92.7$ ms.

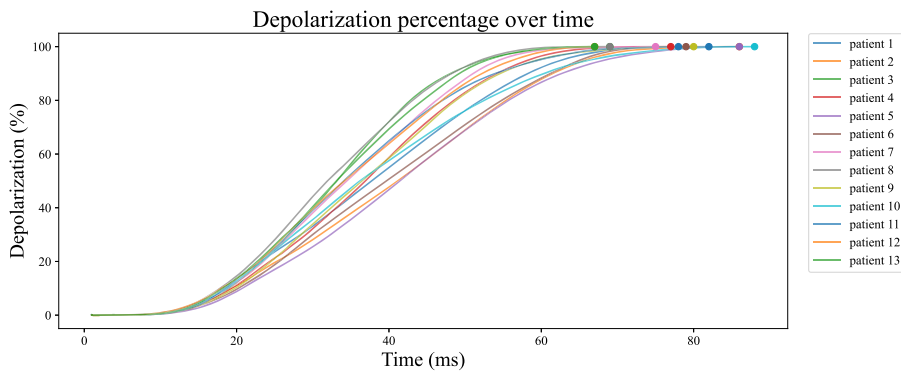


Fig. 9. Percentage of depolarized points within the biventricular geometry of each patient over time measured in ms. Dots on each line mark the point at which 100% depolarization is reached.

The impact of these geometric differences is further illustrated in Fig. 9, which charts the percentage of points that have been depolarized over time in the original patient cohort. Each patient geometry exhibited distinct depolarization patterns and rates. For example, patients 3 and 10 showed the greatest difference in depolarization times despite having mid-range values in overall size and shape metrics. Similarly, patients 5 and 2 reached complete depolarization at nearly the same time despite large differences in volume, whereas patients 1 and 4 exhibited similar depolarization durations despite varying in wall thickness and apicobasal distance. This indicates that the dynamics of depolarization cannot be explained by a single dominant anatomical parameter.

Fig. 10 presents the activation maps for the 13 original patient geometries, where each map visually encodes activation time using a color gradient — blue represents early activation, while red indicates later activation. These visualizations enable us to identify regions of rapid activation (e.g. Purkinje fibers) and areas of delayed activation (e.g the ventricular myocardium).

Fig. 11 provides the corresponding activation maps for the synthetic cohort. Both figures highlight the variability in electrical signal propagation across the biventricular geometries. These variations emphasize the individualized nature of cardiac electrophysiology and the inherent heterogeneity present in both cohorts.

3.3. Branched latent neural maps results

We train the BLNM by using an optimal configuration of 16 layers, 42 neurons per layer, 5 states, a disentanglement level of 2, and a learning rate of 0.0005. For each epoch, we randomly select 1485 points from our training geometries. The BLNM model runs for 1000 epochs with the objective of predicting the activation time for all nodes within the geometry, specifically aiming to minimize MSE.

Training our BLNM surrogate model on a single Apple M3 Pro-(12-core) CPU took approximately 12 minutes (711 seconds) in total. The average inference time per prediction is 0.245 ms, computed by running predictions on each geometry in our testing cohort.

We achieve the best overall training MSE of 0.0012 and the best overall testing MSE of 0.0034. The MSE across both training and testing cohorts maintain an order of magnitude of 10^{-3} , with a couple of training geometries producing an MSE of magnitude 10^{-4} . Each individual training geometry, from synthetic patient 1 to synthetic patient 52, records a loss ranging between 0.0009 and 0.0014. The range of losses for our testing cohort varies, with patient 12 demonstrating a loss of 0.0017, while patient 6 exhibits the highest testing loss of 0.0051. Table 3 provides an error analysis for each geometry in the testing cohort.

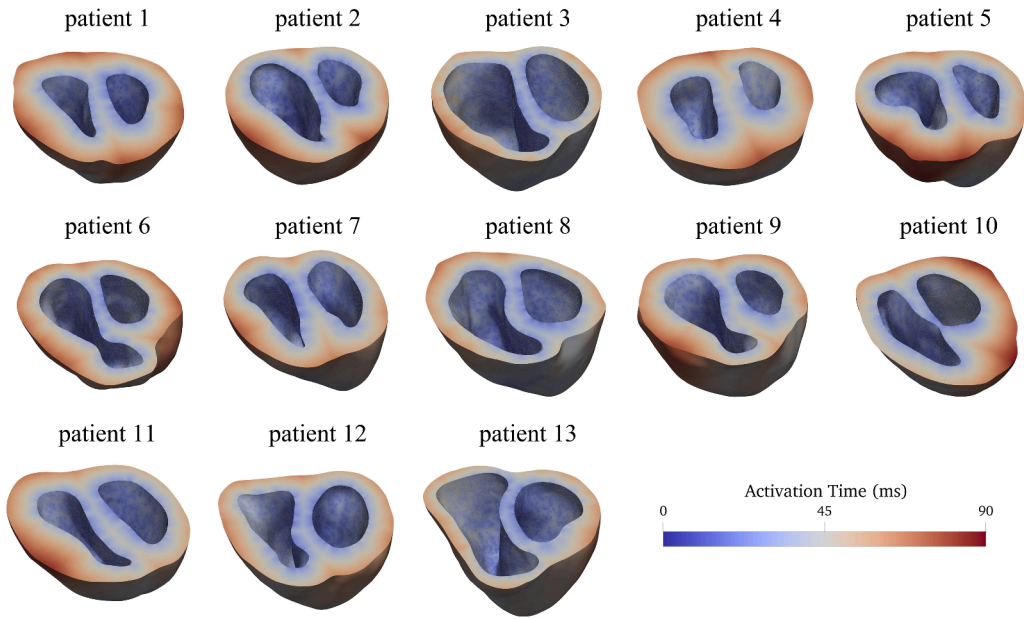


Fig. 10. Simulated activation maps of the original cohort.

Table 3

BLNM MSE for each patient of the original cohort.

Patient Label	Test Error	Patient Label	Test Error
patient 1	0.0021	patient 2	0.0046
patient 3	0.0028	patient 4	0.0049
patient 5	0.0050	patient 6	0.0051
patient 7	0.0045	patient 8	0.0028
patient 9	0.0036	patient 10	0.0031
patient 11	0.0031	patient 12	0.0017
patient 13	0.0023		

Figs. 12 and 13 illustrate the actual and predicted activation times for our worst-performing testing geometry (patient 6) and our best-performing testing geometry (patient 12), respectively. Each figure presents activation times from four distinct perspectives, offering a comprehensive view of the spatial distribution of electrophysiological activity in these 3D activation maps. In both images, the top row displays the activation times obtained from our electrophysiology simulation, while the bottom row shows the activation times predicted by our BLNM model. The color gradient represents the full range of activation times across the biventricular geometry, illustrating regional variations in conduction and the complex spatial distribution of electrophysiological activity.

This pictorial comparison reveals a closer alignment between the predicted and simulated activation times for patient 12, while patient 6 exhibits larger discrepancies, reflecting the differences captured by our MSE values. Nonetheless, in both geometries we see that the BLNM prediction exhibits a smoother gradient transition in comparison to the numerical simulation. This indicates that while the BLNM captures well most of the activation pattern, it tends to average out the finer details. Latent dynamics networks prioritize stable global representations over fine-grained local variations.

We now focus on identifying the specific regions within each testing patient geometry that exhibit significant discrepancies between the simulated and predicted activation times. The prediction errors, defined as a difference in ms between the simulated and actual activation times, range from a maximum value of 18.8 ms to 27.4 ms across patients.

Fig. 14 illustrates the geometries in the original cohort, each color-coded to highlight the top 50% of prediction errors. This visualization helps identify regions where the BLNM struggles most in accurately predicting cardiac activation times. We observe that areas with higher errors often coincide with structurally complex regions of the heart. For example in patients 6,7, and 9, we observe higher errors in parts of the geometry where the right ventricle is particularly narrow. Additionally, each geometry demonstrates significant errors at the initial points of the Purkinje network. We also observe that our BLNM predicts significantly higher activation times for the nodes in the first branch — the main segment immediately following the first Purkinje node before splitting into other branches. This occurs in both the LV and RV Purkinje networks.

Fig. 15 further supports the conclusion that we obtain higher errors at the initial points of the Purkinje. In this figure we compare the average error for every node across a geometry that shares an activation time. For each of the 13 geometries, we notice significantly higher discrepancies between the simulated and predicted activation times at nodes with low activation times. Action potential is

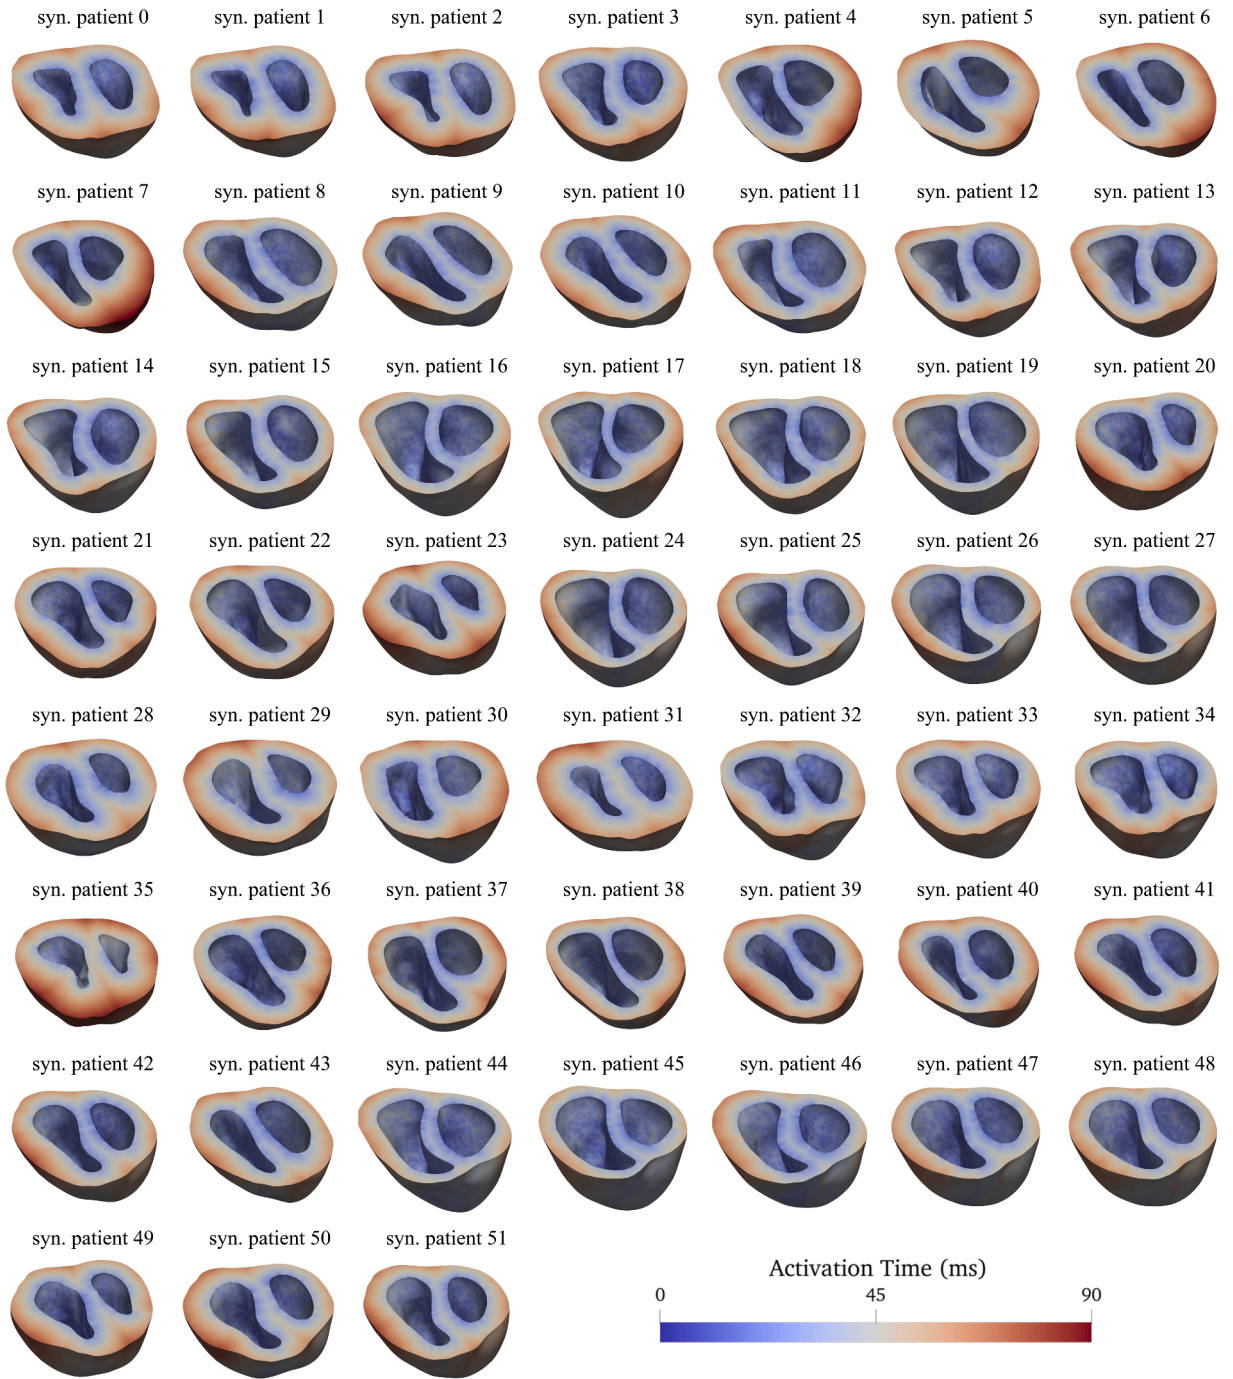


Fig. 11. Simulated activation maps of the synthetic cohort.

applied at time $t = 0$ to the nodes that initiate our Purkinje networks. These points have the lowest activation times and, as evidenced by the plot, exhibit the highest error. Error generally decreases as activation time increases but rises again for the last activating nodes, though still exhibiting smaller absolute errors than the initial activation times.

The geometrical variability of our cohort can be effectively explained using shape scores $\alpha_{i,m}$, which describe the contribution of each principal mode of variation to a given subject, S_i . These shape coefficients are normalized as z-scores. Investigating how the number of modes used in the surrogate model impacts its performance is essential for understanding the extent of variability that can be captured and determining the number of z-scores needed to effectively represent the data. We evaluate model accuracy while systematically varying the number of modes used in our surrogate model and maintaining the same optimal hyperparameters previously determined.

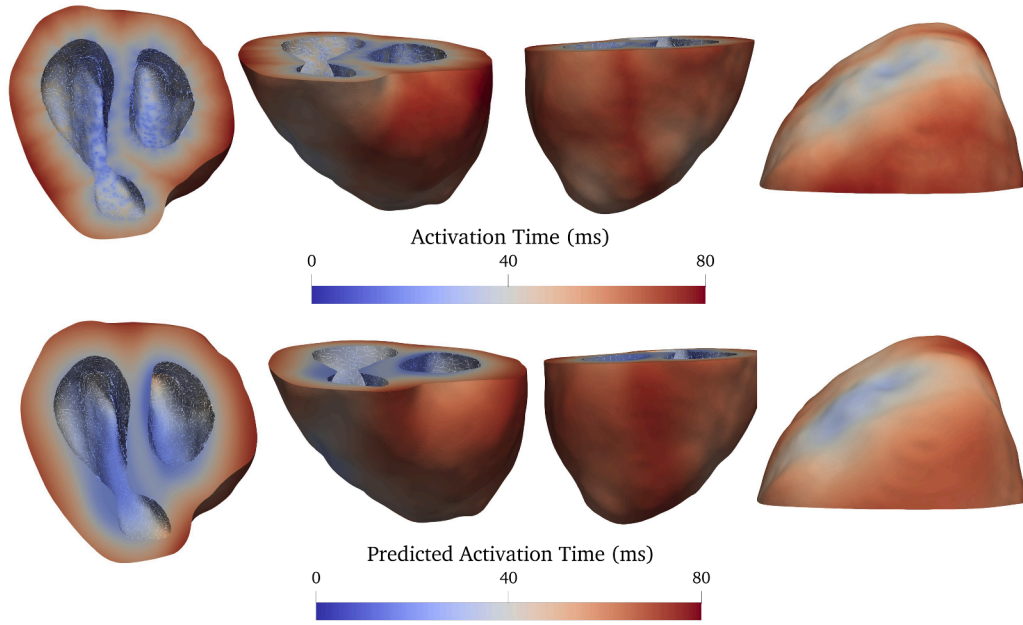


Fig. 12. Simulated activation map (top row) vs. BLNM prediction (bottom row) for patient 6.

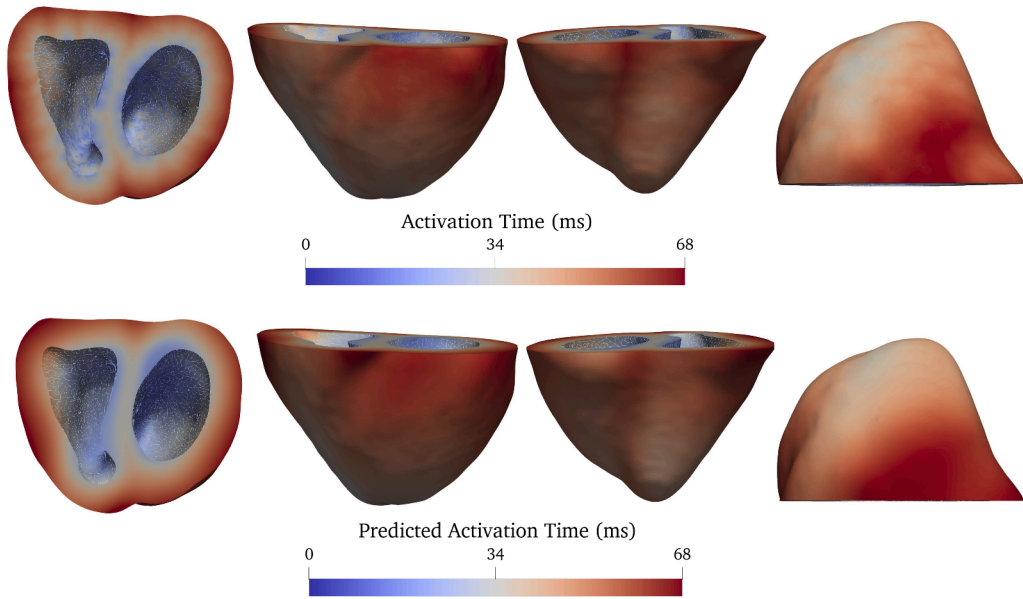


Fig. 13. Simulated activation map (top row) vs. BLNM prediction (bottom row) for patient 12.

From Table 4 we see that adding more modes generally reduces the error between predicted and simulated activation time, indicating improved model performance. However, beyond around 5 modes, the rate of improvement diminishes, and additional modes capture less meaningful variations. Nonetheless, the best losses across all testing geometries are ultimately achieved using all 12 modes. Through this convergence study, we evaluate how well our surrogate model encodes and generalizes the geometrical variability within the testing cohort.

4. Discussion

Accurate modeling of cardiac anatomy and function is essential for advancing the understanding and treatment of congenital heart disease. Neural networks, in particular, have emerged as powerful tools for capturing spatial dynamics and approximating key aspects of cardiac hemodynamics, including pressure, flow rates, cardiac mechanics, and electrodynamics [17,18,47,56,64–67]. By integrating

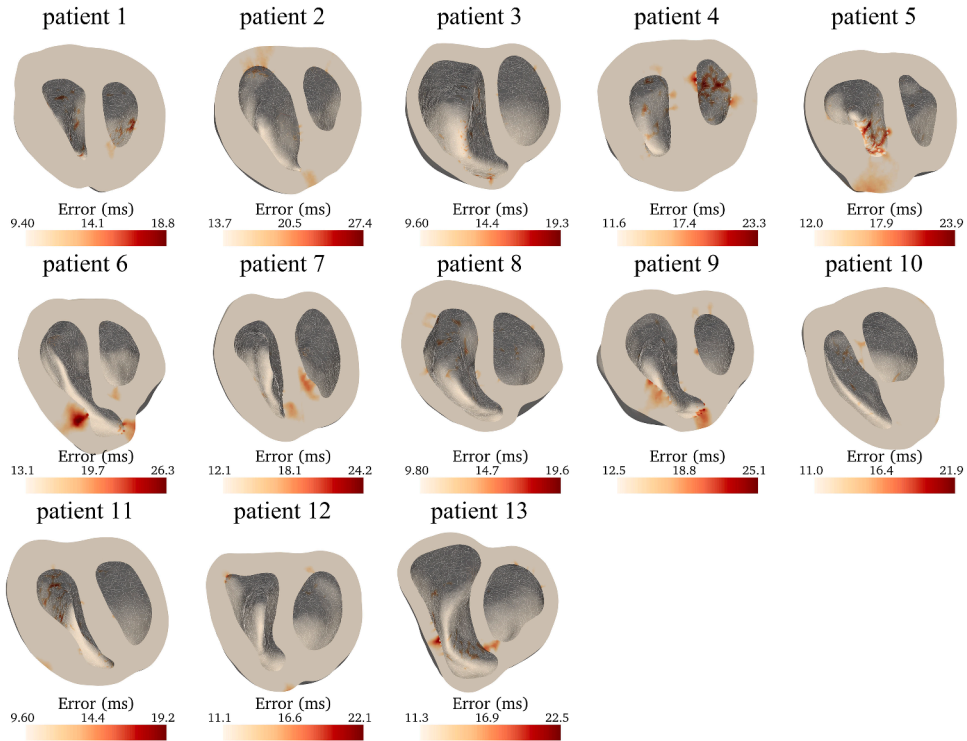


Fig. 14. Top 50% of absolute errors between predicted and simulated activation times for the original cohort. The color scale ranges from white, indicating areas where the prediction error is at or below the median, to red, where errors reach the maximum for that particular geometry. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

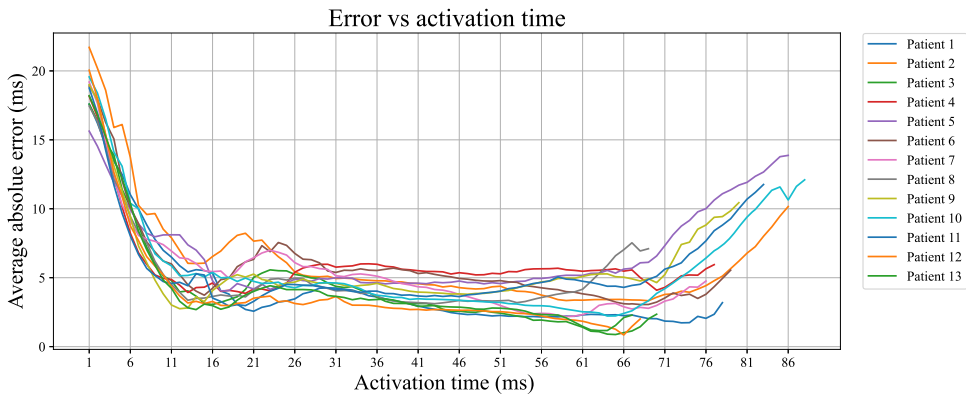


Fig. 15. Average error for all nodes sharing an activation time within a geometry, with each line representing a different patient from patients 1 through 13.

physics-based and data-driven methods, our surrogate model enhanced computational efficiency while preserving prediction accuracy, facilitating progress towards clinically applicable, patient-specific models for understanding and treating congenital heart disease.

Branched Latent Neural Maps (BLNMs) are a novel neural network architecture designed to model the spatial dynamics of physical processes by separating intrinsically different input spaces while leveraging latent variables to capture the underlying variability. These models enable the prediction of both spatial and temporal quantities by integrating diverse data modalities, including spatial coordinates, temporal dynamics, and structural and functional patient-specific characteristics. However, their application thus far has been limited to a single patient geometry for parameter calibration and in silico electrocardiogram generation [36]. To date, BLNMs have not been explored for cohort analysis or for predicting spatiotemporal dynamics with geometry-specific variables.

In this study, we extended the application of BLNMs to predict activation time using spatial coordinates and z-scores. To evaluate their ability to generalize across complex geometries, we utilized a cohort of ToF patients. Given its significant anatomical variability, ToF served as an ideal test case to assess the ability of the model to learn across diverse cardiac structures. We generated 13

Table 4

MSE for the BLNM surrogate model across varying numbers of modes. The table reports the overall best training and testing loss, as well as the testing loss of the worst-performing and best-performing geometries, when utilizing between 1 and 12 modes of variation.

#modes	Training Loss	Testing Loss	Worst Geom Loss	Best Geom Loss
1	0.0056	0.0079	0.0154	0.0047
2	0.0043	0.0062	0.0112	0.0042
3	0.0030	0.0053	0.0095	0.0038
4	0.0022	0.0048	0.0091	0.0035
5	0.0016	0.0039	0.0051	0.0023
6	0.0015	0.0038	0.0063	0.0023
7	0.0012	0.0038	0.0051	0.0022
8	0.0012	0.0036	0.0053	0.0022
9	0.0012	0.0036	0.0066	0.0020
10	0.0013	0.0038	0.0061	0.0020
11	0.0014	0.0039	0.0079	0.0024
12	0.0012	0.0035	0.0051	0.0017

biventricular meshes based on CT scans of pediatric ToF patients. These meshes exhibited significant differences in ventricular size, myocardial volume, wall thickness, and septal morphology — factors that strongly influence cardiac electrophysiology and function [68,69]. The diversity and complexity of cardiac geometries in CHD patients are key reasons why current surrogate models for predicting cardiac function require retraining to generalize across patient-specific anatomies. Although the study focused on ToF patients, the proposed framework is applicable to other cardiac pathologies because it directly incorporates geometric variability using statistical shape modeling and learns full-field activation patterns from physics-based electrophysiology simulations. To demonstrate this adaptability, future research will focus on other CHD cohorts. Validation using other pathologies will provide a rigorous test of the framework’s ability to generalize across diverse structural and electrophysiological patterns.

Nonetheless, for our BLNM surrogate to perform effectively, a larger dataset than the original cohort is beneficial. To address this limitation, statistical shape modeling played a crucial role in augmenting our cohort of original patients. We leveraged an LDDMM-based approach to conduct anatomical mapping of the biventricular geometries in the original population, and we proposed a novel sampling approach to augment the original population whilst retaining the geometrical variability of the original cohort. Several strategies have been proposed for the generation of in-silico cohorts of patients [70,71]. In this regard, data-driven generation methods, including bootstrapping and Gaussian sampling, are particularly suited when the goal of the sampling is to reproduce the existing population [29]. However, such methods perform poorly when the original population is small, and can result in unfeasible generated geometries [29]. Acceptance criteria based on clinical biomarkers are often used to filter out improbable shapes, but are very challenging to compute especially for complex anatomies such as congenital heart disease patients. In this work, we proposed a novel generation strategy that allowed us to constrain the sample selection to the regions of the shape parameter space that are more representative of the underlying anatomical variability in the population. By selecting the four closest samples of this restricted distribution, we effectively augmented the original cohort to improve the performance of our surrogate model. Using this synthetic cohort for training and the original cohort for testing, we were able to overcome the challenge of working with a small number of patient geometries. Future research could extend this framework by performing sensitivity analysis with respect to the sampling threshold.

We conducted 65 electrophysiology simulations, generating activation maps for both the 13 original and 52 synthetic geometries. One key challenge in such simulations is the accurate representation of the Purkinje network, whose intricate structure is difficult to capture through clinical imaging and tough to model directly [72]. To address this, we refined an existing algorithm by enforcing a standardized coverage requirement, ensuring that the network consistently spans a set proportion of each ventricle. We also introduced a systematic procedure for selecting the initial simulation points and determining the direction of the first branch in each network, while still preserving randomized branching to reflect natural physiological variability. These refinements are crucial because the Purkinje system plays a fundamental role in coordinating ventricular depolarization and strongly affects activation timing. Nonetheless, the lack of real patient-specific certainty in branch density, endpoint placement, and branching patterns remains a limitation for the physiological accuracy of depolarization dynamics. Future research could address this limitation by using a probabilistic, ECG-based reconstruction of the Purkinje system. In this approach, a population of plausible Purkinje networks is generated where each network reproduces the patient’s ECGs within a specified tolerance, providing an estimate of topological uncertainty [57]. Incorporating such methodology to our framework would improve physiological accuracy and enhance the robustness of our surrogate model.

In this study, we intentionally focused on isolating the effects of geometry on depolarization dynamics by maintaining uniform electrophysiological properties across all cases. Our simulation parameters, including anisotropy, conductivities, and the ten Tusscher-Panfilov model, maintained the same across all geometries. Our resulting electrophysiology simulations revealed significant variability in activation times and depolarization dynamics in a diverse cohort of biventricular geometries. Notably, each patient geometry exhibited distinct depolarization patterns even among geometries with comparable surface area, volume, wall thickness and apicobasal distance. This suggests that depolarization dynamics is influenced by a combination of geometric factors rather than one dominant anatomical parameter. Geometric variability affects the simulations in two distinct ways: in the definition of the computational

domain of the monodomain model and in the formation of the Purkinje network, whose construction depends on ventricular size and morphology. Due to this dual influence, surrogate models must capture both global anatomical factors and Purkinje network distribution to generalize properly across patients. However, we acknowledge that additional sources of variability such as fibrosis, tissue heterogeneity, or differences in ion channel distributions can impact conduction properties [73].

Beyond electrophysiology, this same framework can be generalized to cardiac mechanics and blood flow dynamics. In a similar way to how BLNMs learned full-field electrical activation patterns from physics-based electrophysiology simulations, a BLNM surrogate model can be created to reproduce spatiotemporal patterns of myocardial strain from finite-element models of cardiac mechanics or flow fields from computational fluid dynamics simulations. Incorporating mechanics and hemodynamics to create an integrated digital twin that captures electromechanical behavior and fluid-structure interaction represents an interesting avenue of future research.

We reported the results of the BLNM surrogate model using hyperparameters determined using OptunaSearch, a hyperparameter optimization algorithm within Ray Tune. Our results demonstrated the ability of BLNMs to predict activation times across a range of complex biventricular geometries with high accuracy relative to physics-based models, which are used as ground truth for evaluation, while also maintaining computational efficiency.

A visual comparison between the simulated and BLNM-predicted activation maps in Figs. 12 and 13 showed that our model effectively captured the general spatial distribution of activation times. However, the predicted results exhibited smoother gradients compared to the simulated results, where activation follows a more localized variation with sharper gradients. This smoothing effect causes activation values in the endocardium, particularly near Purkinje nodes, to be influenced by a gradient diffusion effect, reducing contrast. To mitigate this limitation, Fourier embeddings can be incorporated for input space coordinates [74]. By leveraging Fourier feature mappings, the model can better represent sharp transitions and localized activation variations, minimizing the averaging effect. These embeddings can enhance the model's ability to capture high-frequency functions rather than favoring smooth features during training. However, this improvement comes at the cost of an increased computational complexity [75].

Furthermore, by analyzing prediction errors, measured as the absolute difference between predicted and simulated activation times (in ms), we identified the regions where BLNMs struggled the most to accurately predict activation times. One possible explanation for the largest errors occurring at the earliest activation times, specifically at the first branch of the Purkinje network, is that the algorithm required varying initial branch lengths to ensure adequate coverage of the network. While these early activation regions exhibit the highest discrepancies, small inaccuracies in activation can propagate through the network, affecting downstream activation times. This highlights the importance of accurately modeling the initial activation points and propagation patterns of the Purkinje network to minimize cumulative errors.

Additionally, while prediction errors decrease and remain consistent after the initial Purkinje nodes, we observed a secondary spike in prediction error toward the end of activation. This is visually supported by high errors at the base of the geometries, particularly in regions with greater structural complexity, such as a narrow RV. This issue arises due to the 1D-3D coupling problem, where our fractal tree-based Purkinje network is projected onto a 3D surface, making it difficult for the BLNM to accurately capture activation propagation near the base of the geometry. Namely, the BLNM struggled to reproduce the abrupt termination of the fast conducting Purkinje network. These errors are not solely a limitation of the BLNM, but also highlight the critical role the Purkinje network plays in cardiac activation [57]. Despite this, even though the Purkinje network is generated independently for each geometry, BLNMs successfully capture the activation maps of the 1D-3D Purkinje-myocardium system with MSE of order 10^{-3} . Incorporating more detailed modeling of these networks, such as patient-specific variability in branch lengths and angles could significantly improve predictions. Given the complexity of Purkinje activation, the ability of BLNMs to generalize efficiently while maintaining computational efficiency is a significant achievement. Although challenges remain to accurately capture activation times in these complex regions, our results highlight the robustness and adaptability of the model, with potential for further refinement.

To further assess model performance, we varied the number of anatomical variability modes used in our surrogate model while maintaining the same BLNM architecture. The results, summarized in Table 4, indicate that increasing the number of anatomical modes generally enhances model performance, reducing both training and testing errors. With only a few modes, the model primarily captures large-scale or global variations, such as overall size or orientation, but does not account for finer structural details. As more modes are incorporated, the model becomes capable of encoding more nuanced and localized variations, such as curvature differences or small regional deformations. Using all 12 modes, which performed best across all metrics, allows the model to generalize effectively to unseen shapes by capturing the full spectrum of anatomical variability. In general, our analysis demonstrates the critical role of the number of modes in determining the performance of the surrogate model, with the 12 modes providing the most complete representation.

5. Conclusions

In this study, we evaluated the performance of a Branched Latent Neural Maps (BLNM) surrogate model in predicting activation times generated by physics-based models across a diverse and complex cohort of biventricular geometries derived from patients with Tetralogy of Fallot (ToF). By leveraging both patient-derived and synthetic datasets, our analysis highlights the importance of capturing anatomical variability to develop a robust and generalizable surrogate model. We found that BLNMs generate activation times given space coordinates and z-scores as inputs with high accuracy and computational efficiency, achieving low MSE values of order 10^{-3} across all geometries.

The innovations presented in this study span both methodology and application. The inclusion of synthetic geometries generated through statistical shape modeling (SSM) enabled us to address the limitations of working with a small and challenging cohort, thereby improving the model's ability to generalize across highly varying geometries. Furthermore, we applied BLNMs to predict geometry-

specific activation times, demonstrating that we can develop a surrogate model for electrophysiology that does not require retraining across different geometries. By optimizing hyperparameters, such as network depth, neurons per layer, and disentanglement levels we achieved an MSE of 0.0034, with errors ranging from 0.0017 to 0.0051 across testing geometries. The training process, which completed in approximately 12 minutes (711 seconds) on a single CPU, highlights the computational efficiency of our approach. This rapid training time, combined with the ability of the model to maintain a low MSE across both training and testing cohorts, underscores the potential of BLNMs for large-scale applications, including cohort studies and personalized cardiac simulations.

Despite these achievements, our analysis identified key areas for improvement, particularly in capturing the Purkinje network variability and refining activation time prediction in regions of high structural complexity. The variability in prediction errors further highlights the sensitivity of cardiac activation dynamics to geometric features, such as myocardial wall thickness, septal morphology, and ventricular size. To enhance the model's predictive capabilities future work should focus on improving patient-specific Purkinje network representation and refining the model's ability to capture sharp transitions in activation by incorporating Fourier embeddings.

Additionally, while our results demonstrate that the surrogate model can accurately predict activation patterns within the framework of our physics-based simulations, clinical accuracy depends on validation using patient-derived data. Comparing BLNM predictions to patient ECG signals, particularly in individuals with CHD, would help evaluate the clinical relevance of the model and its ability to capture key electrophysiological features. Recent work has already demonstrated the feasibility of this approach by aligning digital twin electrophysiology simulations with clinical ECGs in patients with hypoplastic left heart syndrome, a form of single ventricle physiology [19]. Building on this foundation, our ongoing efforts aim to extend this framework to a broader cohort of CHD patients, where we will assess the validity of our BLNM surrogate model using patient-specific ECGs. Nevertheless, our findings establish that the model can successfully learn activation dynamics, reinforcing its potential for patient-specific cardiac modeling.

In conclusion, this study advances the application of BLNMs for interpatient analysis, demonstrating their potential for accurately predicting cardiac activation times across anatomically diverse biventricular geometries. Future work should prioritize integrating more detailed patient-specific anatomical data, particularly regarding the Purkinje network, and exploring larger, more diverse datasets to enhance the robustness and clinical applicability of surrogate modeling in cardiac electrophysiology. Additionally, we propose incorporating multiple pathological conditions, such as hypoplastic left heart syndrome and transposition of the great arteries, into the same network. This would enable the model to generalize across multiple independent patient cohorts, making it a more versatile tool for electrophysiological modeling and potential clinical applications.

CRediT authorship contribution statement

Elena Sabdy Martinez: Writing – original draft, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization; **Beatrice Moscoloni:** Writing – original draft, Methodology, Investigation, Formal analysis, Data curation; **Matteo Salvador:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Conceptualization; **Fanwei Kong:** Writing – review & editing, Software, Data curation; **Mathias Peirlinck:** Writing – review & editing, Project administration, Funding acquisition, Data curation, Conceptualization; **Alison Lesley Marsden:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition, Conceptualization.

Data availability

Code is publicly available publicly available under MIT License at <https://github.com/StanfordCBCL/BLNM>.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the [National Science Foundation](#) (DGE-2146755 to E.M and 2310909 to M.S, F.K, A.L.M); the [National Institutes of Health](#) (Grant no. R01HL173845 to M.S, F.K, A.L.M); the Research Foundation – Flanders, Fonds voor Wetenschappelijk Onderzoek – Vlaanderen (Grant no. 11P5524N, to B.M.) and European Union's Horizon Europe Research and Innovation Program (VITAL – Grant No. 101136728, to M.P.). We acknowledge the Stanford Cardiovascular Institute.

References

- [1] J.A. Leopold, J. Loscalzo, The emerging role of precision medicine in cardiovascular disease, *Circ. Res.* 122 (9) (2018) 1302–1315.
- [2] O.Z. Tikenogullari, M. Peirlinck, H. Chubb, A.M. Dubin, E. Kuhl, A.L. Marsden, Effects of cardiac growth on electrical dyssynchrony in the single ventricle patient, *Comput. Methods Biomech. Biomed. Eng.* 27 (8) (2024) 1011–1027.
- [3] O.J. Britton, A. Bueno-Orovio, K.V. Ammel, H.R. Lu, R. Towart, D.J. Gallacher, B. Rodriguez, Experimentally calibrated population of models predicts and explains intersubject variability in cardiac cellular electrophysiology, *Proc. Natl. Acad. Sci. USA* 110 (23) (2013) E2098–E2105.
- [4] M. Peirlinck, J. Yao, F. Sahli Costabal, E. Kuhl, How drugs modulate the performance of the human heart, *Comput. Mech.* 69 (6) (2022) 1397–1411, <https://doi.org/10.1007/s00466-022-02146-1>
- [5] F. Vadakkumpadan, V. Gurev, J. Constantino, H. Arvalo, N.A. Trayanova, *Modeling of Whole-Heart Electrophysiology and Mechanics: Toward Patient-Specific Simulations*, Springer, New York, NY, 2010. 145–165

- [6] G.R. Mirams, S.A. Niederer, R.H. Clayton, The fickle heart: uncertainty quantification in cardiac and cardiovascular modelling and simulation, *Philos. Trans. R. Soc. A* 378 (2173) (2020) 20200119.
- [7] M. Peirlinck, F.S. Costabal, J. Yao, J.M. Guccione, S. Tripathy, Y. Wang, D. Ozturk, P. Segars, T.M. Morrison, S. Levine, E. Kuhl, Precision medicine in human heart modeling: perspectives, challenges, and opportunities, *Biomech. Model. Mechanobiol.* 20 (3) (2021) 803–831, <https://doi.org/10.1007/s10237-021-01421-z>
- [8] S.A. Niederer, J. Lumens, N.A. Trayanova, Computational models in cardiology, *Nat. Rev. Cardiol.* 16 (2) (2019) 100–111.
- [9] M. Kadem, L. Garber, M. Abdelkhalik, B.K. Al-Khazraji, Z. Keshavarz-Motamed, Hemodynamic modeling, medical imaging, and machine learning and their applications to cardiovascular interventions, *IEEE Rev. Biomed. Eng.* 16 (2023) 403–423.
- [10] O.A. Jaffery, L. Melki, G. Slabaugh, W.W. Good, C.H. Roney, A review of personalised cardiac computational modelling using electroanatomical mapping data, *Arrhythm. Electrophysiol. Rev.* 13 (2024) e08. <https://doi.org/10.15420/aer.2023.25>, <https://doi.org/10.15420/aer.2023.25>
- [11] S. Giffard-Roisin, T. Jackson, L. Fovargue, J. Lee, H. Delingette, R. Razavi, N. Ayache, M. Sermesant, Noninvasive personalization of a cardiac arrhythmology model from body surface potential mapping, *IEEE Trans. Biomed. Eng.* 64 (9) (2017) 2206–2218.
- [12] R. Doste, M. Lozano, J.F. Gomez, A. Alcaine, L. Mont, A. Berrueto, O. Camara, R. Sebastian, Predicting the origin of outflow tract ventricular arrhythmias using machine learning techniques trained with patient-specific electrophysiological simulations, in: *2019 Computing in Cardiology (CinC)*, 2019, 1–4.
- [13] Stefania Fresca, Andrea Manzoni, Luca Dedè, Alfio Quarteroni, Deep learning-based reduced order models in cardiac electrophysiology, *PLoS One* 15 (2020) 1–32.
- [14] M. Peirlinck, F. Sahli Costabal, E. Kuhl, Sex differences in drug-induced arrhythmogenesis, *Front. Physiol.* 12 (2021), <https://doi.org/10.3389/fphys.2021.708435>
- [15] C.H. Roney, I. Sim, J. Yu, M. Beach, A. Mehta, J.A. Solis-Lemus, I. Kotadia, J. Whitaker, C. Corrado, O. Razeghi, E. Vigmond, S.M. Narayan, M. O'Neill, S.E. Williams, S.A. Niederer, Predicting atrial fibrillation recurrence by combining population data and virtual cohorts of patient-specific left atrial models, *Circulation* 15 (2) (2022) e010253.
- [16] F. Regazzoni, M. Salvador, L. Dedè, A. Quarteroni, A machine learning method for real-time numerical simulations of cardiac electromechanics, *Comput. Methods Appl. Mech. Eng.* 393 (2022) 114825.
- [17] L. Cicci, S. Fresca, A. Manzoni, A. Quarteroni, Efficient approximation of cardiac mechanics through reduced-order modeling with deep learning-based operator approximation, *Int. J. Numer. Method Biomed. Eng.* 40 (1) (2023) e3783 <https://doi.org/10.1098/rsif.2023.0729>.
- [18] L. Pegolotti, M.R. Pfaller, N.L. Rubio, K. Ding, R. Brugarolas Brufau, E. Darve, A.L. Marsden, Learning reduced-order models for cardiovascular simulations with graph neural networks, *Comput. Biol. Med.* 168 (2024) 107676.
- [19] M. Salvador, F. Kong, M. Peirlinck, D.W. Parker, H. Chubb, A.M. Dubin, A.L. Marsden, Digital twinning of cardiac electrophysiology for congenital heart disease, *J. R. Soc. Interface* 21 (215) (2024), <https://doi.org/10.1098/rsif.2023.0729>
- [20] F. Regazzoni, L. Dedè, A. Quarteroni, Machine learning of multiscale active force generation models for the efficient simulation of cardiac electromechanics, *Comput. Methods Appl. Mech. Eng.* 370 (2020) 113268, <https://doi.org/https://doi.org/10.1016/j.cma.2020.113268>
- [21] F.S. Costabal, Y. Yang, P. Perdikaris, D.E. Hurtago, E. Kuhl, Physics-informed neural networks for cardiac activation mapping, *Front. Physiol.* 8 (2020).
- [22] E.M. Cherry, F.H. Fenton, Effects of boundaries and geometry on the spatial distribution of action potential duration in cardiac tissue, *J. Theor. Biol.* 285 (1) (2011) 164–176.
- [23] C. Rodero, M. Strocchi, M. Marciniak, S. Longobardi, J. Whitaker, M.D. O'Neill, K. Gillette, C.M. Augustin, G. Plank, E.J. Vigmond, P. Lamata, S.A. Niederer, Linking statistical shape models and simulated function in the healthy adult human heart 17 (4) (2021) e1008851–e1008851, <https://doi.org/https://doi.org/10.1371/journal.pcbi.1008851>
- [24] C. Nagel, S. Schuler, O. Dössel, A. Loewe, A bi-atrial statistical shape model for large-scale in silico studies of human atria: model development and application to ECG simulations, *Med. Image Anal.* 74 (2021) 102210, <https://doi.org/https://doi.org/10.1016/j.media.2021.102210>
- [25] S. Kollar, E. Balaras, L.J. Oliveri, Y.-H. Loke, F. Capuano, Statistical shape modeling reveals the link between right ventricular shape, hemodynamic force, and myocardial function in patients with repaired tetralogy of Fallot, *AJP Heart Circ. Physiol.* 323 (3) (2022) H449–H460, <https://doi.org/https://doi.org/10.1152/ajpheart.00228.2022>
- [26] F. Sophocleous, L. Standen, G. Doolub, R. Laymouna, C. Bucciarelli-Ducci, M. Caputo, N. Manghat, M. Hamilton, S. Curtis, G. Biglino, Analysing functional implications of differences in left ventricular morphology using statistical shape modelling, *Sci. Rep.* 12 (2022) 19163, <https://doi.org/https://doi.org/10.1038/s41598-022-15888-y>
- [27] S. Cutugno, T. Ingrassia, V. Nigrelli, S. Pasta, On the left ventricular remodeling of patients with stenotic aortic valve: a statistical shape analysis, *Bioengineering* 8 (5) (2021) 66, <https://doi.org/https://doi.org/10.3390/bioengineering8050066>
- [28] B. Moscoloni, C. Beeche, J.A. Chirinos, P. Segers, M. Peirlinck, Unveiling sex dimorphism in the healthy cardiac anatomy: fundamental differences between male and female heart shapes (2025), <https://doi.org/10.48550/ARXIV.2503.00197>
- [29] P. Romero, M. Lozano, F. Martínez-Gil, D. Serra, R. Sebastián, P. Lamata, I. García-Fernández, Clinically-driven virtual patient cohorts generation: an application to aorta, *Front. Physiol.* 12 (2021), <https://doi.org/https://doi.org/10.3389/fphys.2021.713118>
- [30] B. Thamsen, P. Yevtushenko, L. Gundelwein, A.A.A. Setio, H. Lamecker, M. Kelm, M. Schafstedde, T. Heimann, T. Kuehne, L. Goubergrits, Synthetic database of aortic morphometry and hemodynamics: overcoming medical imaging data availability, *IEEE Trans. Med. Imaging* 40 (5) (2021) 1438–1449, <https://doi.org/https://doi.org/10.1109/tmi.2021.3057496>
- [31] S. Bridio, G. Luraghi, A. Ramella, J.F. Rodriguez Matas, G. Dubini, C.A. Luisi, M. Neidlin, P. Konduri, N. Arrarte Terreros, H.A. Marquering, C.B.L.M. Majoie, F. Migliavacca, Generation of a virtual cohort of patients for in silico trials of acute ischemic stroke treatments, *Appl. Sci.* 13 (18) (2023) 10074, <https://doi.org/https://doi.org/10.3390/app131810074>
- [32] K. Ostendorf, K. Bäuml, D. Mastrodicasa, D. Fleischmann, B. Preim, G. Mistelbauer, Synthetic surface mesh generation of aortic dissections using statistical shape modeling, *Comput. Graph.* 124 (2024) 104070, <https://doi.org/https://doi.org/10.1016/j.cag.2024.104070>
- [33] S. Saitta, L. Maga, C. Armour, E. Votta, D.P. O'Regan, M.Y. Salmasi, T. Athanasiou, J.W. Weinsaft, X.Y. Xu, S. Pirola, A. Redaelli, Data-driven generation of 4D velocity profiles in the aneurysmal ascending aorta, *Comput. Methods Programs Biomed.* 233 (2023) 107468, <https://doi.org/https://doi.org/10.1016/j.cmpb.2023.107468>
- [34] R. Scuoippo, S. Castelbuono, S. Cannata, G. Gentile, V. Agnese, D. Bellavia, C. Gandolfo, S. Pasta, Generation of a virtual cohort of TAVI patients for in silico trials: a statistical shape and machine learning analysis, *Med. Biol. Eng. Comput.* 63 (2025) 467–482, <https://doi.org/https://doi.org/10.1007/s11517-024-03215-8>
- [35] S. Niederer, Y. Aboelkassem, C.D. Cantwell, C. Corrado, S. Coveney, E.M. Cherry, T. Delhaas, F.H. Fenton, A.V. Panfilov, P. Pathmanathan, G. Plank, M. Riabiz, C. Roney, R. Weber, L. Wang, Creation and application of virtual patient cohorts of heart models, *Philos. Trans. R. Soc. A* 378 (2173) (2020) 20190558, <https://doi.org/https://doi.org/10.1098/rsta.2019.0558>
- [36] M. Salvador, A.L. Marsden, Branched latent neural maps, *Comput. Methods Appl. Mech. Eng.* 418 (2024) 116499.
- [37] A. Karoui, M. Bendahmane, N. Zemzemi, Cardiac activation maps reconstruction: a comparative study between data-driven and physics-based methods, *Front. Physiol.* 12 (2021) 686136.
- [38] J.D. Bayer, R.C. Blake, G. Plank, N. Trayanova, A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models, *Ann. Biomed. Eng.* 40 (2012) 2243–2254.
- [39] Sadd M. Khan, Nigel E. Drury, John Stickley, David J. Barron, William J. Brawn, Timothy J. Jones, Robert H. Anderson, Adrian Crucean, Tetralogy of fallot: morphological variations and implications for surgical repair, *Eur. J. Cardio-Thoracic Surg.* 56 (2018) 101–109.
- [40] X. Xu, T. Wang, H. Yuan, Q. Jia, ImageCHD: a 3D computed tomography image dataset for classification of congenital heart disease, *Proc. Med. Image Comput. Comput. Assist. Interv. (MICCAI)* 12264 (2020) 77–87.
- [41] A. Updegrove, N.M. Wilson, J. Merkow, H. Lan, A.L. Marsden, S.C. Shadden, SimVascular: an open source pipeline for cardiovascular simulation, *Ann. Biomed. Eng.* 45 (2016) 525–541.
- [42] J. Glauès, A. Qiu, M.I. Miller, L. Younes, Large deformation diffeomorphic metric curve mapping, *Int. J. Comput. Vis.* 80 (3) (2008) 317–336, <https://doi.org/https://doi.org/10.1007/s11263-008-0141-9>
- [43] A. Bône, M. Louis, B. Martin, S. Durrleman, Deformetrica 4: an open-Source software for statistical shape analysis, *Shape Med. Imaging* 11167 (2018) 3–13.

- [44] S. Durrleman, M. Prastawa, N. Charon, J.R. Korenberg, S. Joshi, G. Gerig, A. Trounev, Morphometry of anatomical shape complexes with dense deformations and sparse parameters, *NeuroImage* 101 (2014-11) 35–49, <https://doi.org/10.1016/j.neuroimage.2014.06.043>
- [45] J.L. Bruse, K. McLeod, G. Biglino, H.N. Ntsinjana, C. Capelli, T.-Y. Hsia, M. Serresant, X. Pennec, A.M. Taylor, S. Schievano, A statistical shape modelling framework to extract 3D shape biomarkers from medical imaging data: assessing arch morphology of repaired coarctation of the aorta, *BMC Med. Imaging* 16 (1) (2016) 40, <https://doi.org/10.1186/s12880-016-0142-z>
- [46] P.J. Besl, N.D. McKay, A method for registration of 3-D shapes, *IEEE Trans. Pattern Anal. Mach. Intell.* 14 (2) (1992) 239–256, <https://doi.org/https://doi.org/10.1109/34.121791>
- [47] M. Peirlinck, F. Sahli Costabal, K.L. Sack, J.S. Choy, G.S. Kassab, J.M. Guccione, M. De Beule, P. Segers, E. Kuhl, Using machine learning to characterize heart failure across the scales, *Biomech. Model. Mechanobiol.* 18 (6) (2019) 1987–2001, <https://doi.org/10.1007/s10237-019-01190-w>
- [48] Y.-C. Chen, A tutorial on kernel density estimation and recent advances, *Biostat. Epidemiol.* 1 (1) (2017) 161–187, <https://doi.org/https://doi.org/10.1080/24709360.2017.1396742>
- [49] R. Schmidt, K. Singh, Meshmixer: an interface for rapid mesh composition, in: *ACM SIGGRAPH 2010 Talks, SIGGRAPH '10*, Association for Computing Machinery, New York, NY, USA, 2010, <https://doi.org/10.1145/1837026.1837034>
- [50] C. Geuzaine, J.-F. Remacle, Gmsh: a 3-D finite element mesh generator with built-in pre- and post-processing facilities, *Int. J. Numer. Methods Eng.* 79 (11) (2009) 1309–1331, <https://doi.org/https://doi.org/10.1002/nme.2579>
- [51] P. Colli Franzone, L.F. Pavarino, S. Scacchi, *Mathematical Cardiac Electrophysiology*, 13, Springer, 2014.
- [52] A. Quarteroni, L. Dede, A. Manzoni, C. Vergara, *Mathematical Modelling of the Human Cardiovascular System: Data, Numerical Approximation, Clinical Applications*, Cambridge University Press, 2019.
- [53] K.H. ten Tusscher, A.V. Panfilov, Alternans and spiral breakup in a human ventricular tissue model, *Am. J. Physiol. Heart Circ. Physiol.* 291 (2006) 1088–1100.
- [54] F. Sahli Costabal, J. Yao, E. Kuhl, Predicting drug-induced arrhythmias by multiscale modeling, *Int. J. Numer. Method Biomed. Eng.* 34 (5) (2018) e2964.
- [55] R. Piersanti, P.C. Africa, M. Fedele, C. Vergara, L. Dede, A.F. Corno, A. Quarteroni, Modeling cardiac muscle fibers in ventricular and atrial electrophysiology simulations, *Comput. Methods Appl. Mech. Eng.* 373 (2021) 113468.
- [56] F.S. Costabal, D.-E. Hurtado, E. Kuhl, Generating Purkinje networks in the human heart, *J. Biomech.* 49 (2016) 2455–2465.
- [57] F. Álvarez-Barrientos, M. Salinas-Camus, S. Pezzuto, F. Sahli Costabal, Probabilistic learning of the Purkinje network from the electrocardiogram, *Med. Image Anal.* 101 (2025) 103460.
- [58] C. Zhu, V. Vedula, D. Parker, N. Wilson, S. Shadden, A. Marsden, svFSI: a multiphysics package for integrated cardiac modeling, *J. Open Source Softw.* 7 (2022) 4118.
- [59] F. Regazzoni, M. Salvador, P.C. Africa, M. Fedele, L. Dede, A. Quarteroni, A cardiac electromechanical model coupled with a lumped-parameter model for closed-loop blood circulation, *J. Comput. Phys.* 457 (2022) 111083.
- [60] R. Piersanti, F. Regazzoni, M. Salvador, A.F. Corno, L. Dede, C. Vergara, A. Quarteroni, 3D-0D closed-loop model for the simulation of cardiac biventricular electromechanics, *Comput. Methods Appl. Mech. Eng.* 391 (2022) 114607.
- [61] M. Fedele, R. Piersanti, F. Regazzoni, M. Salvador, P.C. Africa, M. Bucelli, A. Zingaro, L. Dede, A. Quarteroni, A comprehensive and biophysically detailed computational model of the whole human heart electromechanics, *Comput. Methods Appl. Mech. Eng.* 410 (2023) 115983.
- [62] S. Krishnamoorthi, M. Sarkar, W.S. Klug, Numerical quadrature and operator splitting in finite element methods for cardiac electrophysiology, *Int. J. Numer. Method Biomed. Eng.* 29 (2013) 1243–1266.
- [63] T. Akiba, S. Sano, T. Yanase, T. Ohta, M. Koyama, Optuna: a next-generation hyperparameter optimization framework, in: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19*, Association for Computing Machinery, New York, NY, USA, 2019, p. 2623–2631, <https://doi.org/10.1145/3292500.3330701>
- [64] F. Regazzoni, S. Pagani, A. Quarteroni, Universal solution manifold networks (USM-Nets) non-Intrusive mesh-free surrogate models for problems in variable domains, *J. Biomech. Eng.* 144 (12) (2022) 121004.
- [65] F.S. Costabal, K. Matsuno, J. Yao, P. Perdikaris, E. Kuhl, Machine learning in drug development: characterizing the effect of 30 drugs on the QT interval using Gaussian process regression, sensitivity analysis, and uncertainty quantification, *Comput. Methods Appl. Mech. Eng.* 348 (2019) 313–333.
- [66] J. Xie, B. Yao, Physics-constrained deep active learning for spatiotemporal modeling of cardiac electrodynamics, *Comput. Biol. Med.* 146 (2022) 105586.
- [67] J. Xie, B. Yao, Physics-Constrained deep learning for robust inverse ECG modeling, *IEEE Trans. Autom. Sci. Eng.* 20 (1) (2023) 151–166.
- [68] N.A. Trayanova, A. Lyon, J. Shade, J. Heijman, Computational modeling of cardiac electrophysiology and arrhythmogenesis: toward clinical translation, *Physiol. Rev.* 104 (3) (2024).
- [69] D.D. Chan, K.C. Wu, Z. Loring, L. Galeotti, G. Gerstenblith, G. Tomaselli, R.G. Weiss, G.S. Wagner, D.G. Strauss, Comparison of the relation between left ventricular anatomy and QRS duration in patients with cardiomyopathy with versus without left bundle branch block, *Am. J. Cardiol.* 113 (10) (2014) 1717–1722.
- [70] Sabine Verstraeten, Martijn Hoeijmakers, Pim Tonino, Jan Brüning, Claudio Capelli, Wouter van Huberts, Generation of synthetic aortic valve stenosis geometries for in silico trials, *Int. J. Numer. Method Biomed. Eng.* 40 (1) (2023) e3778, <https://doi.org/https://doi.org/10.1002/cnm.3778>
- [71] F. Kong, S. Stocker, P.S. Choi, M. Ma, D.B. Ennis, A. Marsden, SDF4CHD: generative modeling of cardiac anatomies with congenital heart defects, *Med. Image Anal.* 97 (2024) 103293.
- [72] S. Krishnamoorthi, M. Sarkar, W.S. Klug, Numerical quadrature and operator splitting in finite element methods for cardiac electrophysiology, *Int. J. Numer. Method Biomed. Eng.* 29 (11) (2013) 1243–1266.
- [73] S. Verheule, U. Schotten, Electrophysiological consequences of cardiac fibrosis, *Cells* 10 (11) (2021) 3220.
- [74] M. Tancik, P.P. Srinivasan, B. Mildenhall, N.R. Sara Fridovich-Keil, U. Singhal, R. Ramamoorthi, J.T. Barron, R. Ng, Fourier features let networks learn high frequency functions in low dimensional domains, *ACM Digit. Lib.* 632 (2020) 7537–7547.
- [75] M. Salvador, A.L. Marsden, Liquid Fourier latent dynamics networks for fast GPU-based numerical simulations in computational cardiology, *ResearchGate* (2024), <https://doi.org/10.48550/arXiv.2408.09818>.