



EU Horizon 2020 Research & Innovation Program  
Digital transformation in Health and Care  
SC1-DTH-06-2020  
Grant Agreement No. 101016496

## **SimCardioTest - Simulation of Cardiac Devices & Drugs for in-silico Testing and Certification**



### **Technical Report**

D 5.5: Report on advanced data science processing of in-silico trials

### **Work Package 5 (WP 5)**

In-silico Trials & Data Science

Task Lead: Inria, France  
WP Lead: Inria, France

PUBLIC



## DELIVERABLE INFORMATION

<b>Deliverable number</b>	5.5
<b>Deliverable title</b>	Report on advanced data science processing of in-silico trials
<b>Description</b>	Data science for in-silico trials and in-silico data
<b>Lead authors</b>	Irene Balelli
<b>Contributors</b>	Irene Balelli, Marta Saiz, Michael Leguebe, Yves Coudière, Delphine Deshors, Beatriz Trenor
<b>Due date</b>	30/06/2025
<b>Submission date</b>	7 July 2025
<b>Comments</b>	

Document history			
Date	Version	Author(s)	Comments
22/05/2025	Template	I. Balelli	
23/06/2025	1st finalized version	I. Balelli, M. Saiz, M. Leguebe, Y. Coudière, D. Deshors, B. Trenor	
03/07/2025	Quality Check	O. Camara	
04/07/2025	Updated version	I. Balelli, Y. Coudière	
07/07/2025	Final editing	M. Barbier	



## Table of Contents

<b>EXECUTIVE SUMMARY</b>	<b>4</b>
<b>1. INTRODUCTION</b>	<b>5</b>
<b>2. SENSITIVITY ANALYSIS FOR ELECTROMECHANICAL MODELS</b>	<b>5</b>
<b>2.1 GOAL-ORIENTED SENSITIVITY ANALYSIS OF A MODEL OF AN ARTIFICIAL PACEMAKER</b>	<b>6</b>
<b>2.1.1 OBJECTIVES</b>	<b>6</b>
<b>2.1.2 METHODOLOGY</b>	<b>6</b>
<b>2.1.3 RESULTS AND CONCLUSIONS</b>	<b>10</b>
<b>2.2 SENSITIVITY ANALYSIS THROUGH CAUSAL DISCOVERY</b>	<b>14</b>
<b>2.2.1 OBJECTIVES</b>	<b>14</b>
<b>2.2.2 METHODOLOGY</b>	<b>14</b>
<b>2.2.3 RESULTS AND CONCLUSIONS</b>	<b>18</b>
<b>3. A DIGITAL TWIN APPROACH INTEGRATING CLINICAL, STRUCTURAL, AND IN-SILICO HEMODYNAMIC DATA TO UNCOVER STROKE RISK FACTORS</b>	<b>22</b>
<b>3.1 OBJECTIVES</b>	<b>22</b>
<b>3.2 METHODOLOGY</b>	<b>22</b>
<b>3.3 RESULTS</b>	<b>26</b>
<b>3.4 DISCUSSION AND CONCLUSIONS</b>	<b>30</b>
<b>4. CAUSAL LEARNING FOR IN-SILICO BIOMARKERS ANALYSIS AND TdP RISK ASSESSMENT</b>	<b>31</b>
<b>4.1 LINEAR NON-GAUSSIAN ACYCLIC MODEL FOR ION CHANNEL BLOCKADE</b>	<b>32</b>
<b>4.1.1 OBJECTIVES</b>	<b>32</b>
<b>4.1.2 METHODOLOGY</b>	<b>32</b>
<b>4.1.3 RESULTS AND CONCLUSIONS</b>	<b>33</b>
<b>4.2 MULTIMODAL CAUSAL VAE WITH IN-SILICO ELECTROPHYSIOLOGICAL BIOMARKERS</b>	<b>35</b>
<b>4.2.1 OBJECTIVES</b>	<b>35</b>
<b>4.2.2 METHODOLOGY</b>	<b>36</b>
<b>4.2.3 RESULTS AND CONCLUSIONS</b>	<b>38</b>
<b>5. CONCLUSIONS</b>	<b>40</b>
<b>6. BIBLIOGRAPHY</b>	<b>42</b>

## EXECUTIVE SUMMARY

This report constitutes the SimCardioTest WP5 Deliverable D5.5 due in June 2025 (M54) on advanced data science for in-silico trials. This report is organized in three main sections, each one related to a specific use case and addressing a specific scientific question, from sensitivity analysis and model reduction to biomarker extraction, risk stratification and phenotyping. A general *Introduction* and *Conclusions* sections are included to provide a unified overview on the challenges tackled and the remaining open questions and perspectives in the field. All works presented here, realized under Task 5.6 (M24-M54), concerns the development and application of statistical and machine learning-based approaches designed to enhance the use of in-silico models and in-silico generated data, thus providing scientific evidence for their usefulness and relevance in supporting clinical-decision making.

## 1. Introduction

Work Package 5 (WP5) has two main objectives: (i) to conduct in-silico trials across the selected use cases and evaluate the simulation results against acquired data; and (ii) to develop and apply advanced statistical and machine learning (ML) methodologies to extract deeper insights from the simulations, with the potential to discover novel informative biomarkers. Results related to objective (i) have been presented in deliverables D5.1–D5.4. The present report focuses on objective (ii), and details the research outcomes developed under Task 5.6 (T5.6), *Advanced Data Science for Biomarker Discovery*. T5.6 consolidates the expertise of the consortium partners (INRIA, University of Bordeaux -UBx, Universidad Pompeu Fabra -UPF, Universitat Politècnica de Valencia - UPV) to explore modern data science methods for more robust and insightful analysis of in-silico data.

This includes, for instance, the use unsupervised machine learning (ML) techniques—such as multiple kernel learning and variational autoencoders—to identify latent representations that support phenotyping and risk assessment, while enabling the integration of in-silico features and providing evidence on the utility of simulations for decision support and personalized medicine. In addition, we investigate sensitivity analysis approaches tailored to quantify the impact of parameters in highly complex models, and apply causal discovery techniques for robust feature selection.

Overall, the report is organized as follows. Section 2 presents two different approaches developed within the consortium (UBx and INRIA being the main contributors for these two studies, respectively) for performing sensitivity analysis of highly complex mathematical models, related to Use Case 1 (UC1). Section 3 describes a study (realized by UPF) in which multi-domain data—including clinical, morphological and in-silico hemodynamic features—are jointly analyzed using unsupervised ML to identify phenotypes associated with cardioembolic thrombus, relevant to UC2. Finally, Section 4 reports on two works (led by INRIA in collaboration with UPV) which address the characterization and assessment of drug-induced Torsade-de-Pointes (UC3). These studies employ causal learning-based techniques, either for performing feature selection, or for the construction of an informative latent space using unsupervised multimodal ML. Section 5 summarizes the main results presented here, and concludes the report.

## 2. Sensitivity analysis for electromechanical models

For the last couple of decades, cardiac modeling has been shown to be a useful and reliable tool to study cardiac function in healthy and under arrhythmia conditions [1], in the context of cardiac resynchronization therapy [2], to investigate mechanoelectrical feedback in healthy and left bundle branch block [3] as well as to elucidate the effect of fiber organization in cardiac output [4]. The emergence of advanced imaging technologies (e.g., high resolution Computed Tomography (CT), photon counting CT, high resolution cardiac Magnetic Resonance Imaging (MRI)) has considerably increased the availability of data to better characterize cardiac functions. From one side, this can allow models to reach a higher level of realism and details hence improved accuracy. From the other side, this is done at the cost of an increasing model complexity due to the need to integrate a greater number of parameters. Consequently, investigating the interplay between the model parameters, and understanding their impact on the main model outputs of interest, represent a relevant but complex and challenging task: this motivates the need for reliable approaches to sensitivity analysis [5].

We distinguish between local sensitivity analysis (LSA) and global sensitivity analysis (GSA). LSA's objective is to assess the effects on model outputs upon small perturbations of one parameter at a time around a nominal parameter value. Besides the advantage of being computationally efficient and easy to interpret, LSA only partially answers our initial question, and is unable to explain the effect of the simultaneous variation of several input parameters, especially in highly nonlinear models. On the other hand, GSA seeks to explore the input parameters space in a more comprehensive way, and quantify the input parameter importance based on a characterization of the resulting output response surface. For this reason, global approaches for quantitative and model-free sensitivity analysis may be preferred [6], in order to manage the effect of each parameter and their interaction.

## 2.1 Goal-oriented sensitivity analysis of a model of an artificial pacemaker

### 2.1.1 Objectives

The computational model used to perform in-silico trials of a pacemaker lead solves complex partial differential equations on a three-dimensional (3D) mesh. Its high computational cost prevents from performing a global sensitivity analysis to determine which of its input parameters may produce high uncertainty on the result, and can be calibrated from experimental data acquired during the project. To this aim, we developed a surrogate model without spatial dimension (0D), which can produce results at a significantly reduced cost, and can be used in statistical studies. In order to evaluate the sensitivity of the 3D model to its parameters, we have to proceed in two steps:

1. Perform the sensitivity analysis on the 0D model with respect to its own input parameters, some of which being different from the 3D model.
2. Build a mapping between the 3D and 0D parameters to add to the analysis the transfer of uncertainty from one set of parameters to the other.

In this document, we report the results of the first step, which were obtained during the M42-M54 period. These results were obtained in two phases. First, we performed a GSA using the theory of Sobol indices, for a large set of input and output quantities related to the question of interest determined in the V&V process (deliverables D6.1 and D6.2). It was used to determine the major output of interest, and most influential inputs. Second, we refined the GSA following a techniques of design optimization and uncertainty, from [7]. The results from the Sobol analysis have been published in [8].

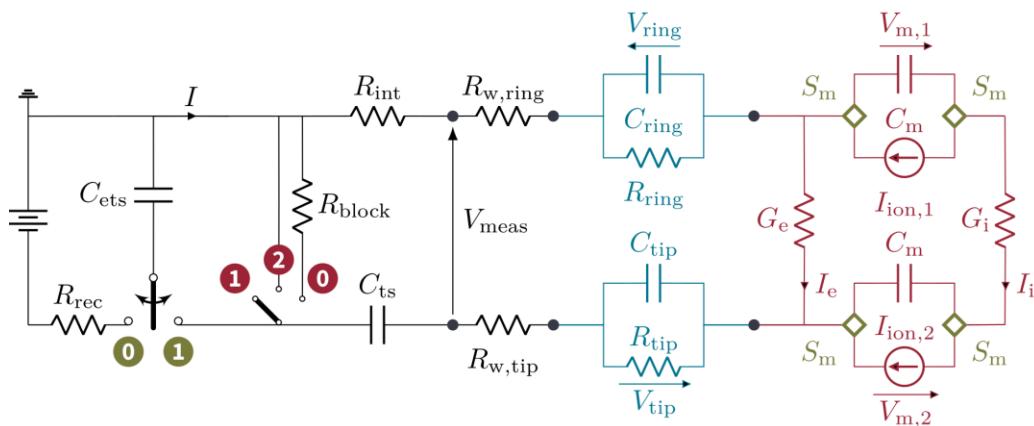
### 2.1.2 Methodology

#### The 0D model

The 0D model, developed during the PhD thesis of V. Pannetier and presented at the conference FIMH2025 [9], computes the transmembrane voltage (TMV) of cell membranes excited by a pacemaker lead. This quantity is used to determine whether or not an action potential was triggered by the pacemaker, which then allows to generate the so-called Lapicque curves that are the output of the in-silico trial.

The 0D model represents the pacemaker device, the cardiac tissue and the contact between them on the electrodes as an electric circuit (Figure 1). Cell membranes are represented by non-linear current generators that depend on the TMVs  $V_{mi}$ ,  $i=1,2$ .

The parameters of the stimulation device ( $C_{ets}$ ,  $C_{ts}$ ,  $R_{rec}$ ,  $R_{block}$ ,  $R_{int}$ ,  $R_{w,ring}$ ,  $R_{w,tip}$ ) are considered uncertain in the first step. This is due to the fact that two types of pacing devices were considered, the BOREA pacemaker system for which the values were known, and a commercial PSA (Pulse System Analyzer) for which the values were uncertain, and obtained by experimental characterization. Afterwards, they were fixed, since their influence on the output were proved less important than the other modeling parameters.



**Figure 1.** Electric circuit of the 0D bipolar pacemaker model. The part corresponding to the device is drawn in black, the cardiac tissue in red, and the contacts between them in blue. Diamonds indicate where the current are scaled by the surface factor  $S_m$ . The parameters of the black part and  $G_e$ ,  $G_i$ ,  $S_m$  (in red) are defined in the text and considered uncertain. The contact parameters  $R_{tip,ring}$ ,  $C_{tip,ring}$  (in blue) have been calibrated (see D2.3), and the ionic parameters  $C_m$ ,  $I_{ion,1,2}$  (in red) are taken from the literature.

The membrane parameters  $C_m$  and those included in the ionic currents  $I_{ion,i}$  are taken from the literature.

Only seven parameters remain, related to the bio-electrode contact, and the biological tissue:

- $C_{ring}$ ,  $R_{ring}$ ,  $C_{tip}$  and  $R_{tip}$  are part of the model of the bio-electric contact impedance between the tissue and the pacemaker electrodes (called ring and tip). They were first calibrated with data from bench experiments where the electrodes were placed in a saline bath [10]. However, we suppose that the contact with a different medium can change significantly their value.
- $G_e$ ,  $G_i$  are surrogate parameters for the conductivity of the tissue, which is represented in 3D by a tensorial conductivity oriented along the myocardial fibers.  $G_e$  and  $G_i$  correspond to the equivalent conductance of the medium through which current flows.
- $S_m$  is a scaling parameter that makes the units compatible between the ionic models (expressed in units per  $\text{cm}^2$ ) and the currents in ampere delivered by the device. It corresponds to the total surface ( $\text{in cm}^2$ ) of cell membrane through which the current passes.

We used a first guess for  $S_m$ ,  $G_i$  and  $G_e$  by considering a stimulated volume of  $2 \text{ mm}^3$  and a cylindrical shape of cells with a  $100 \mu\text{m}$  length and  $10 \mu\text{m}$  diameter. However, these estimations are highly uncertain. For this reason, we will rather conduct the sensitivity analysis on the logarithm of all parameters, rather than their absolute value, as their order of magnitude is uncertain.

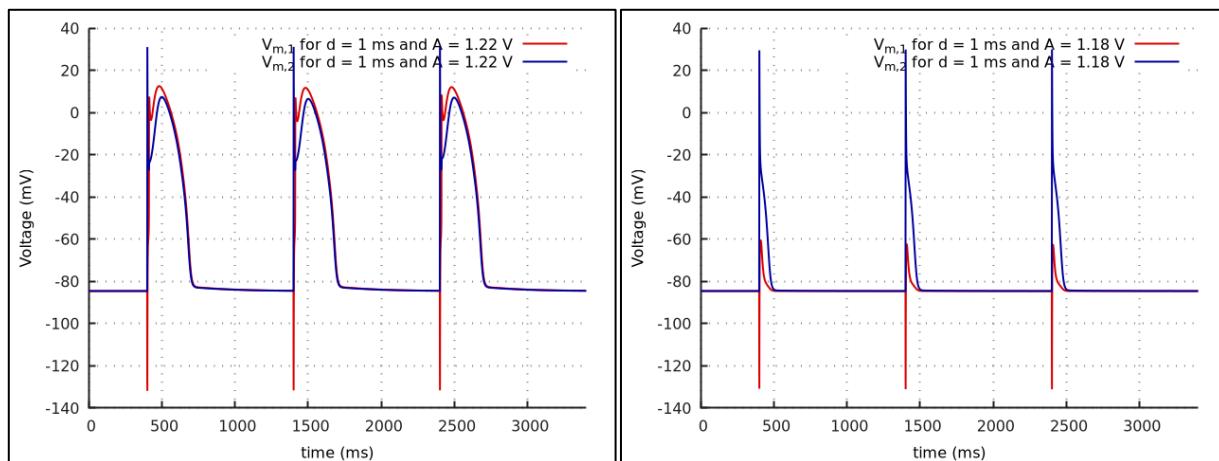
For a fixed amplitude  $A$  (in volt), and duration  $d$  (in ms) of stimulation, we use our in-house [CirCE software](#) to compute the TMVs  $V_{m1}$  and  $V_{m2}$  as illustrated in Figure 2. These voltages depend on time, and can have either the shape of an action potential, when there is capture, or close to their baseline when there is no capture, i.e. if the amplitude and duration of the stimulation were large enough to trigger the action potential.

In the first step of the analysis [8], we computed the total order Sobol indices of 7 output features for each of the paced TMVs, namely:

- the average of  $(V_{mi} - V_0)$   $i=1,2$  over time,  $V_0$  being the baseline of  $V_{mi}$ ;
- the min and max values of  $V_{mi}$ ;
- the last value of  $V_{mi}$  in the time frame;
- the min and the max of the time derivative  $dV_{mi}/dt$ ;
- the APD50 value.

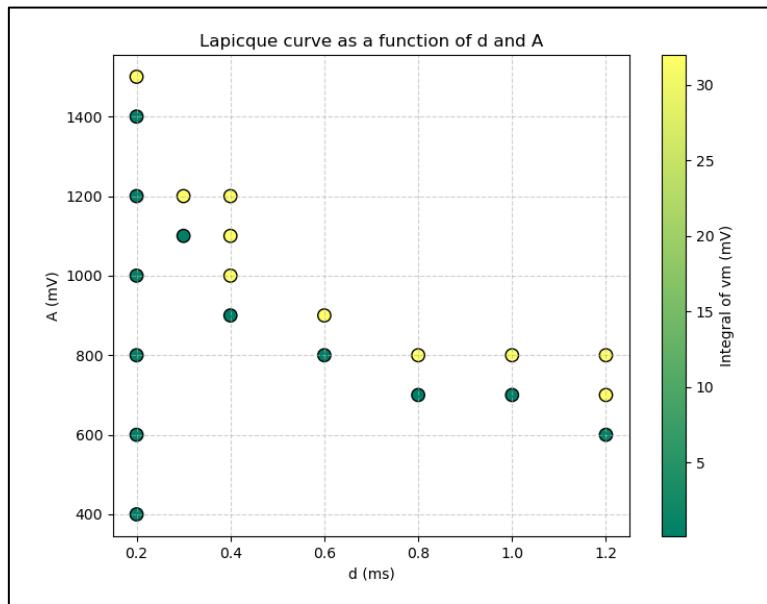
We also computed the same 7 features on the voltage measured between the tip and ring electrodes (the possible electric measure during ex-vivo experiments,  $V_{meas}$  on Figure 1). The Sobol indices were computed with the [SALib Python library](#).

Based on results of this first analysis, we selected the average of  $(V_{m1} - V_0)$  over the last stimulation period as the best quantity to discriminate between capture and no capture.



**Figure 2.** Transmembrane voltage  $V_{m1}$  (red) and  $V_{m2}$  (blue) from the 0D model, for a capturing stimulation (left), and a non-capturing one (right).

When running computations for various amplitude-duration couples, we can perform a threshold search which is similar to the experimental procedure, as shown in Figure 3. This search was conducted while using the reference value of the parameters to study and calibrate (Table 1).



**Figure 3.** Threshold search performed with the 0D model with parameters from Table 1. Green circles represent absence of capture, while yellow circles correspond to a regular action potential. The Lapicque curve should be located between the green and yellow circles. Voltages and durations are those that the device can deliver.

**Table 1.** Reference value for the 7 parameters of the 0D model (see Figure 1), with  $t = RC$ .

$S_m(\text{cm}^2)$	$G_e(\text{mS}/\text{cm}^2)$	$G_i(\text{mS}/\text{cm}^2)$	$R_{\text{ring}}(\text{kOhm})$	$t_{\text{ring}}(\text{s})$	$R_{\text{tip}}(\text{kOhm})$	$t_{\text{tip}}(\text{s})$
40	0.5	500	2	37.48	0.03	0.1665
$C_{\text{ets}}$ (microF)	$C_{\text{ts}}$ (microF)	$R_{\text{int}}$ (kOhm)	$R_{\text{pulse}}$ (kOhm)			
9.36	10.62	0.018	0.007			

### Local Sensitivity Analysis

In this part, we consider a fixed stimulation duration  $d$ , and study the influence of the 0D model parameters  $X = (S_m, G_e, G_i, R_{\text{ring}}, t_{\text{ring}}, R_{\text{tip}}, t_{\text{tip}})$  on the minimum amplitude  $A$  which triggers an action potential. We then modify the value of a parameter  $X_i$  by  $\pm 20\%$ . For each modified parameter  $X_i^\pm$ , we define the local sensitivity index  $SI_i^\pm$  and  $SI_i$  as

$$SI_i^\pm = \frac{A_i^\pm - A_{\text{mean}}}{X_i^\pm - X_{\text{mean}}}, \quad SI_i = \max(SI_i^-, SI_i^+),$$

where  $A_{\text{mean}}$  and  $X_{\text{mean}}$  are the mean values of the amplitudes and parameter  $X_i$ , respectively.

### Global Sensitivity Analysis

Instead of centering the sensitivity analysis around a single reference set of parameters, we now sample the whole parameter space, which is determined by choosing a search interval for each parameter. We consider each parameter as a random variable, which takes values in the intervals given in Table 2. In the first step, 786 432 evaluations of the model were used for the analysis (Sobol

indices), and in the second step, 1 024 samples were computed (the parameter space is of dimension 7 instead of 11).

**Table 2.** Search interval for each parameter. We actually sample the logarithm of the parameter in these intervals.

$S_m$ (cm <sup>2</sup> )	$G_e$ (mS/cm <sup>2</sup> )	$G_i$ (mS/cm <sup>2</sup> )	$R_{ring}$ (kOhm)	$t_{ring}$ (s)	$R_{tip}$ (kOhm)	$t_{tip}$ (s)
[0.1, 10 <sup>3</sup> ]	[10 <sup>-4</sup> , 10 <sup>2</sup> ]	[10 <sup>-4</sup> , 10 <sup>2</sup> ]	[10 <sup>-4</sup> , 10 <sup>3</sup> ]	[10 <sup>-3</sup> , 10 <sup>2</sup> ]	[10 <sup>-4</sup> , 10 <sup>3</sup> ]	[10 <sup>-3</sup> , 10 <sup>2</sup> ]
$C_{ets}$ (microF)	$C_{ts}$ (microF)	$R_{int}$ (kOhm)	$R_{pulse}$ (kOhm)			
[10 <sup>-2</sup> , 10 <sup>3</sup> ]	[10 <sup>-2</sup> , 10 <sup>3</sup> ]	[10 <sup>-4</sup> , 10]	[10 <sup>-4</sup> , 10]			

We noticed that the 0D model produces results for the average of  $V_m - V_0$  that are bimodal: on the 1024 model evaluations at step 2, 83% of the simulations resulted in no capture with a value of the average of  $V_m - V_0$  that is close to 0, the remaining 17% lead to TMV averages which are distributed around a positive value near 30 mV. Using variance-based techniques such as Sobol indices is not appropriate for such a model, since the average and variance around the average have no real meaning for bimodal distributions. This is why, in a second step, we rather used goal-oriented sensitivity analysis techniques [7].

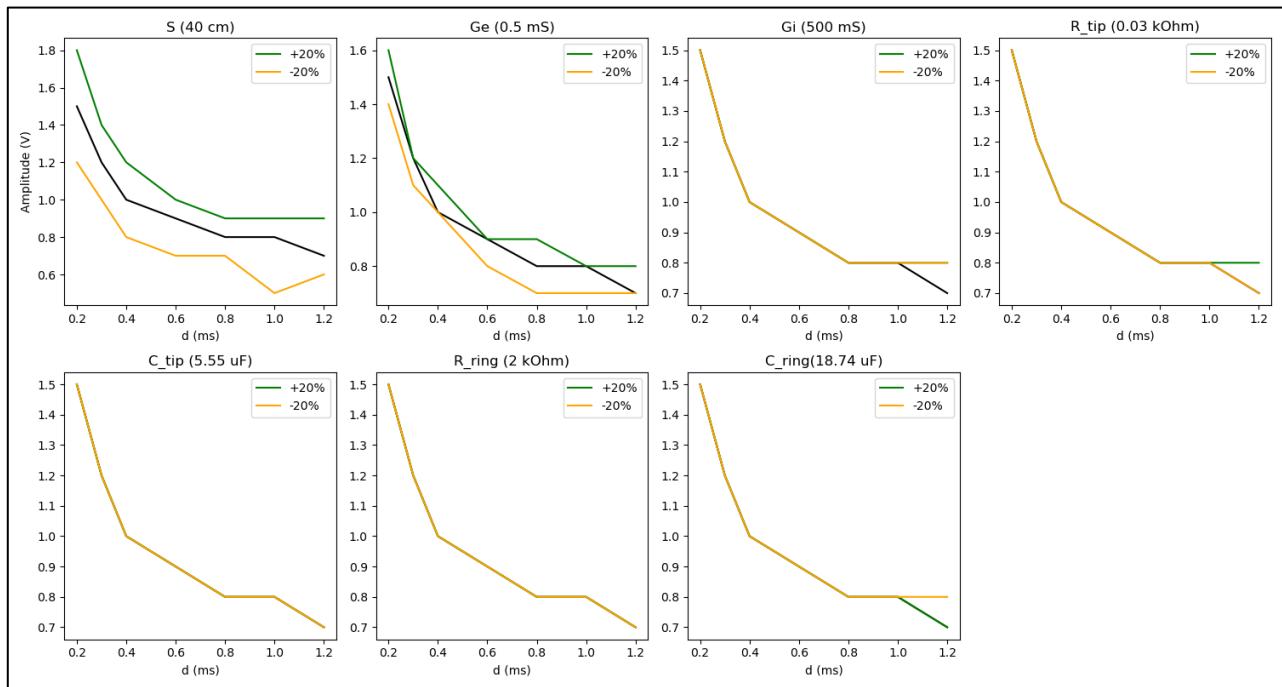
To this aim, we reformulated the question of interest as: *given an amplitude A and duration d of stimulation, what is the probability of capture with respect to the distributed parameters?*

We used the Kolmogorov-Smirnov test (KST) to quantify the influence of the parameters on the output. This statistical test allows to compare two distributions by measuring the maximum Kolmogorov-Smirnov distance between them. In our case, for each couple (A,d), the initial distribution is the sampling of the parameters itself, which is close to uniform in log space. We compare it to the distribution of the parameters in which we have excluded the sets that did not produce a capture. The KST is then applied to evaluate the distance between the two distributions. A large distance means that the parameters have a significant effect on the output, hence on the probability of capture.

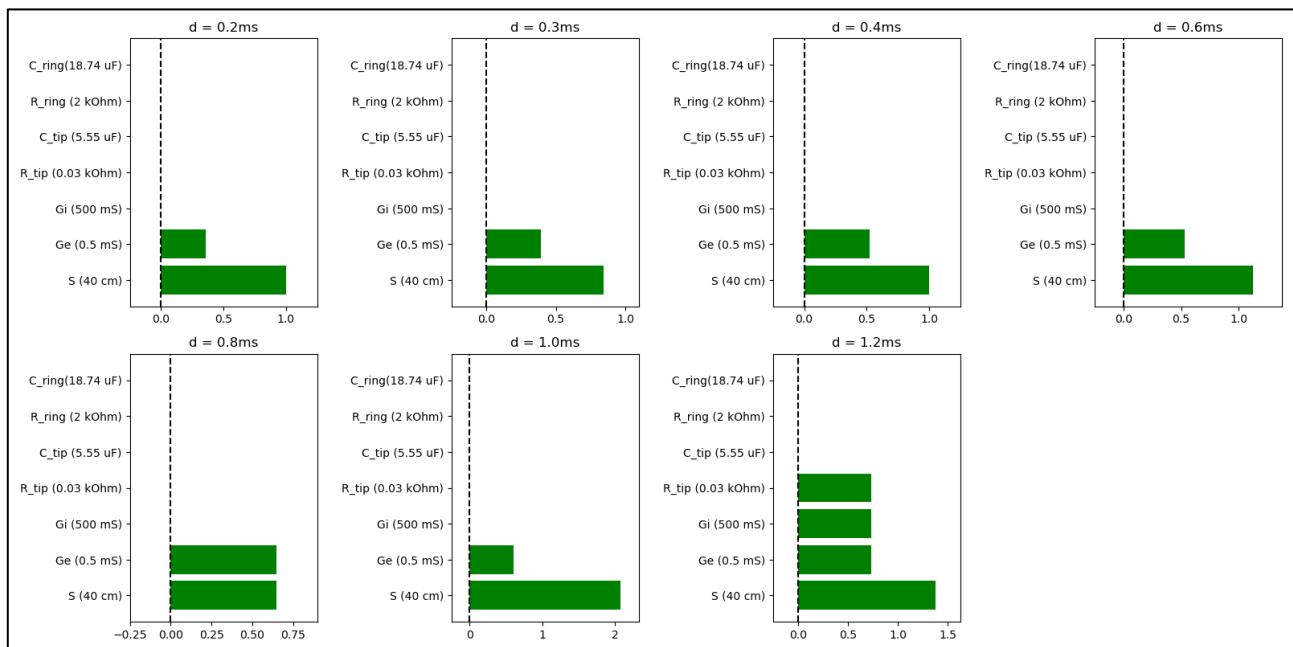
### 2.1.3 Results and conclusions

#### Local Sensitivity Analysis

On Figure 4, we show the resulting Lapicque curves when varying a single parameter of the 0D model by  $\pm 20\%$ . We can already see qualitatively that the parameters  $S_m$  and  $G_e$  have the largest influence. This is confirmed by the indices SI which were computed for several durations d, and are reported in Figure 5.



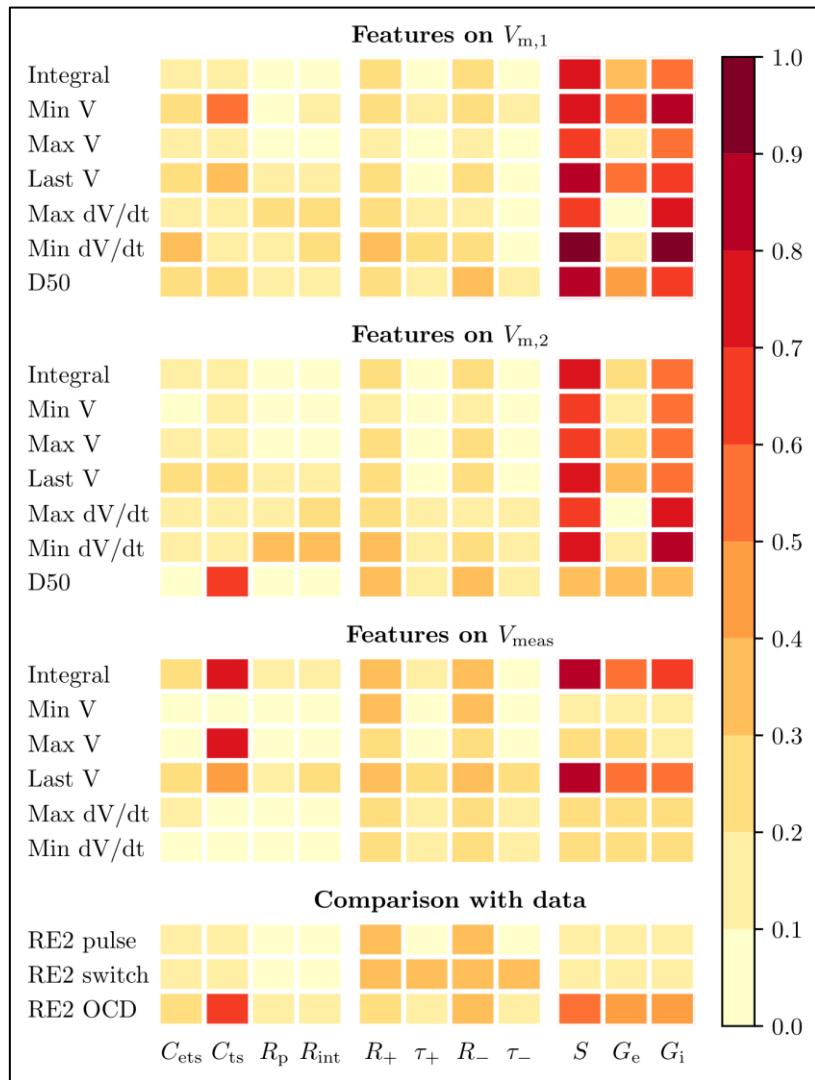
**Figure 4.** Comparison between Lapicque curves computed with the 0D model, while modifying each parameter by  $\pm 20\%$ , one at a time. Values in parentheses are the reference value of each parameter.



**Figure 5.** Local sensitivity indices  $SI_i$  for different stimulation durations ( $d$ ).

### Global Sensitivity Analysis

Over the 786 432 model evaluations from the first step, the simulation failed for 2% of them due to physically incompatible values. The evaluation of these failed samples is replaced by the average evaluation of those remaining during the analysis of the Sobol indices. The results are shown on Figure 6.



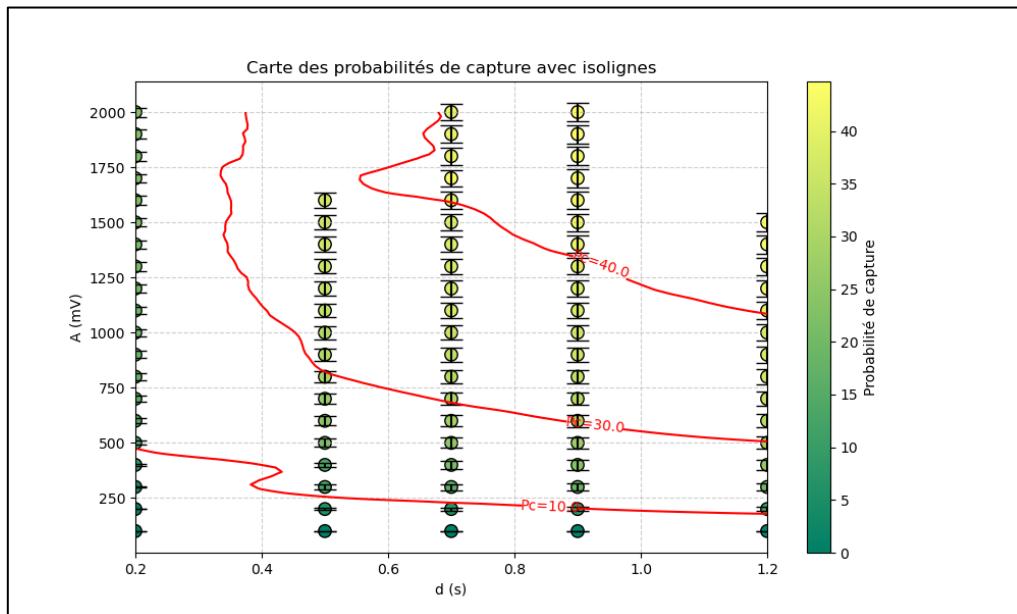
**Figure 6.** Results of total order indices of Sobol analysis on several outputs features on 786 432 samples. The three top blocks report the indices for observable quantities derived from  $V_{m1}$ ,  $V_{m2}$  and  $V_{meas}$  respectively. The lower block is dedicated to  $L^2$  relative difference between the model and experimental data.

The analysis shows a strong influence of the tissue parameters ( $S$ ,  $G_i$ ,  $G_e$ ) on the various characteristics of the TMVs, particularly for the first membrane  $V_{m1}$ . The capacitance value  $C_{ets}$  has little influence on these outputs. Additionally, we can also note the parameters that are not identifiable from the experimental data, such as the capacitance  $C_{ets}$  and the two resistances of the pacemaker circuit  $R_{pulse}$  and  $R_{int}$ . In addition, the average of  $(V_{m1} - V_0)$  appeared to be the best determinant of capture, and not influenced either by the second capacitance  $C_{ts}$ . For these reasons, the second step of the GSA was performed with only 7 uncertain parameters, after discarding ( $C_{ts}$ ,  $C_{ets}$ ,  $R_{pulse}$ ,  $R_{int}$ ).

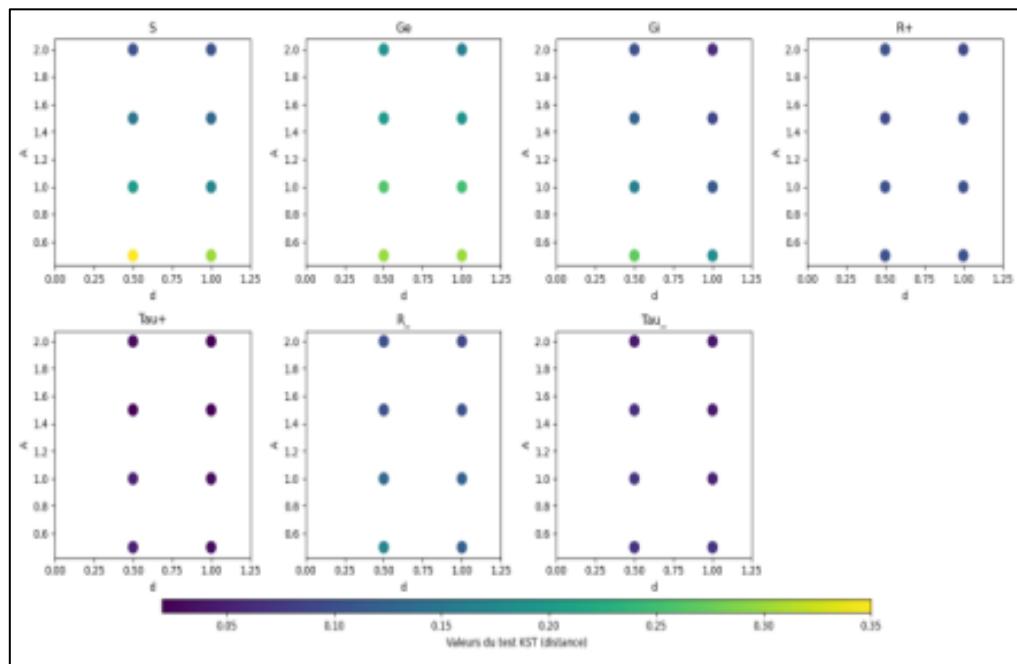
In the second step, computations of the 1024 outputs of the model resulted in a numerical error in 7% of the cases, due to a physical incompatibility of the parameters with the model, or lack of robustness of the numerical solver. Specific solvers for ionic solvers would be needed to address this issue.

Figure 6 displays the capture probability map in the amplitude-duration plane. In Figure 7, we plot the KST results at different amplitudes and durations. This analysis confirms that the contact

parameters  $R_{ring}$ ,  $t_{ring}$ ,  $R_{tip}$  and  $t_{tip}$  have a lesser influence on the capture results than the 3 parameters  $S$ ,  $G_e$  and  $G_i$ , that are specific to the 0D model. We can also observe that the influence is larger in a region around the position of the reference Lapicque curve. This shows that there is a need to accurately determine the 0D parameters, but also confirms that a calibration is possible, since the parameters seem actually identifiable. Additionally, this result reinforces the need to characterize the relationship between parameters of the 3D and 0D models as much as possible.



**Figure 7.** Capture probability in the amplitude-duration space. Red lines are isolines of probability 10%, 30% and 40%.



**Figure 8.** Kolmogorov-Smirnov distance between sampled parameter distributions and parameters of capturing stimulations, for several values of stimulation amplitude and duration.

## 2.2 Sensitivity analysis through causal discovery

### 2.2.1 Objectives

We consider a fast electromechanical model [11] and two main outputs of interest, namely the ejection fraction (EF), defined as the ratio of stroke volume to the end diastolic volume, and the maximum change of pressure in the left ventricular cavity ( $\max(dP/dt)$ ), which is used in the context of cardiac resynchronization therapy to measure acute hemodynamic response to varying pacing parameters and lead locations.

The high complexity of the considered electromechanical model requires the application of advanced sensitivity analysis methods to deal with potential non-linearities, interactions between factors and discontinuities. This in turn requires a large number of model simulations, resulting in high computational costs and time. To address this challenge, we propose an alternative approach based on causality to perform global sensitivity analysis, and investigate the interrelationships among input parameters and the main outputs of interest. This work has been published and presented at FIMH 2025 [12].

### 2.2.2 Methodology

#### The electromechanical model

We consider a fast electromechanical model based on the work of Desrues et. al. [11]. The action potential propagation is computed using a fast-marching method and by solving the anisotropic Eikonal-Diffusion equation (ED) for each mesh vertex. The stress profile (SP) in the mechanical model is defined using an analytical formulation, consisting of two Kumaraswamy cumulative distributions, one for the contraction and a second one accounting for the active relaxation:

$$SP = \begin{cases} 1 - \left(1 - \left(\frac{x-t_d}{t_{APD}}\right)^a\right)^b, & x \in [t_d, t_{APD}] \\ \left(1 - \left(\frac{x-(t_d+t_{APD})}{t_{APD}}\right)^g\right)^h, & x \in [t_d + t_{APD}, t_r + t_{APD}] \end{cases}$$

Further, the heart is described as a passive isotropic Mooney-Rivlin material with corresponding strain energy function given by:

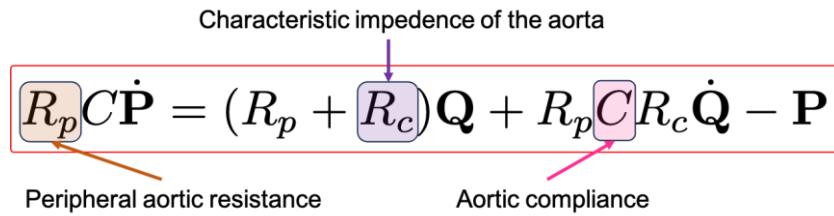
$$W = C_1(\bar{I}_1 - 3) + C_2(\bar{I}_2 - 3) + \frac{\kappa}{2}(J - 1)^2$$

Finally, we couple a hemodynamic model implementing the four phases of the cardiac cycle (filling, isovolumetric contraction, ejection, isovolumetric relaxation) and use the three-element Windkessel model to compute the arterial pressures:

$$R_p C \dot{P} = (R_p + R_c) Q + R_p C R_c \dot{Q} - P$$

Characteristic impedance of the aorta

Peripheral aortic resistance      Aortic compliance

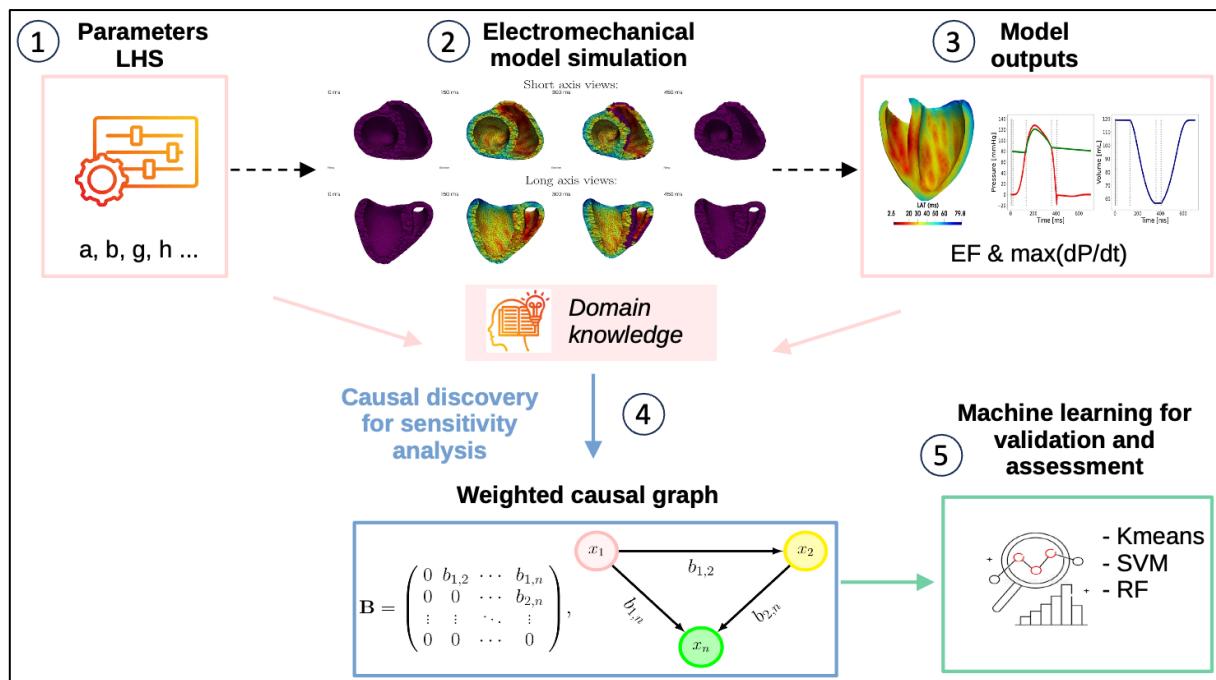


### Global sensitivity analysis and our causality-based approach

Global sensitivity analysis (GSA) methods can be broadly classified into four main categories: variance-based methods, derivative-based methods, density-based methods, and feature-additive methods [5]. Briefly, variance-based methods quantify the contribution of each input variable to the overall output variance; derivative-based methods assess sensitivity by measuring output changes using gradients; density-based methods analyze the variations in the probability distribution of outputs; while feature-additive methods decompose the model output into additive components to evaluate the impact of each input separately. For benchmarking purposes, we will compare our approach against two classical GSA methods: Pawn [13], a density-based GSA method, and Sobol [14], a widely used variance-based GSA method. For both methods we use the Python implementation available in the [SALib](#) package.

We propose to tackle the GSA problem from a causal perspective, and specifically using causal discovery. The aim of causal discovery [15] is to uncover causal relationships among a set of observed variables, which can be graphically represented through Directed Acyclic Graphs (DAGs). In a DAG, variables are represented by nodes, and the causal directions are encoded through directed edges, whose associated weights provide the strength of the influence of one variable on another one following the causal paths; acyclicity means that there is no cycle, i.e., there is no directed path starting and ending in the same node. This graph-based representation offers a holistic overview of the system under study [16], and provide straightforward interpretation of the relationships between the considered variables.

Figure 9 shows schematically our pipeline to perform causality-based GSA. Firstly, the input model parameters are sampled in their respective ranges through Latin Hypercube Sampling (LHS) and the model outputs of interest are computed for each sampled parameters' combination by running the considered electromechanical model (more details about the simulations will follow below). Causal discovery is then deployed on the dataset of model parameters and outputs: the learned causal graph and the information contained therein quantify the effect of the parameters' perturbations over the selected outputs. In particular, the causal weights define our causality-based sensitivity indices. It is worth noticing that prior knowledge about the relationships between model parameters and outputs can eventually be incorporated in the causal discovery phase, if available, to facilitate the estimation of the corresponding weighted causal graph. In order to independently validate the causal graph and its sensitivity interpretation, we additionally performed classification using classical machine learning approaches.



**Figure 9.** Graphical pipeline of causality-based global sensitivity analysis.

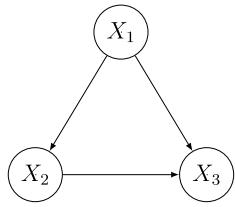
Of note, our causality-based sensitivity indices are not forced to be positive, as opposed to other classical approaches. Besides being an added value of our approach, since the sign of the causal weights can provide additional relevant information to further interpret the mutual relationships between model parameters and outputs, when comparing our results to the benchmark we will use the absolute values of the causality-based sensitivity indices, which denotes the absolute causal strength. In order to investigate the stability of the obtained results, for both the causal weights and the Pawn sensitivity indices we compute the mean values over 5-fold cross-validation. Concerning Sobol, to adhere to the sampling scheme assumption proposed by Slatelli et al. [17] (and deal with the limited number of available model simulations), we choose to repeat the Sobol analysis 10 times by randomly removing each time 8 simulations among the 200 available, and reported the average values, rather than performing the 5-fold cross-validation scheme.

### Causal Discovery through DierctLiNGAM

We choose to rely on Linear non-Gaussian Acyclic models [18] to solve the task of discovering the causal graph relating model parameters and outputs (our observed variables, and the nodes of the searched DAG). LiNGAM assumes that the functional dependencies of each cause on its effects are linear, with associated additive independent and centered non-Gaussian error terms. Formally, let  $X = \{x_n\}_n$  our dataset, where  $x_i, i = 1, \dots, n$ , is a vector containing the observed values of the  $i$ -th variable of interest (here, either a model parameter or a model output). The generating process for each  $x_i \in X$  writes:

$$x_i = \sum_{j \in \text{Pa}_i} b_{ji} x_j + e_i, \quad (1)$$

where  $\text{Pa}_i$  is the set of parents of  $x_i$  (i.e., the nodes who precede  $x_i$  following the sense of directed edges),  $b_{ji}$  defines the linear causal effect of  $x_j \rightarrow x_i$ , and  $e_i$  denotes the noise of  $x_i$ , supposed independent from any other  $e_j$ , when  $j \neq i$ .



In order to fix the ideas, let us suppose we observe 3 variables,  $X_1, X_2, X_3$ , which are causally related following the DAG in Figure 10. In this case we can say that  $X_1$  has no parent;  $X_2, X_3$  are both children of  $X_1$ ; and  $X_1, X_2$  are both parents of  $X_3$ .

**Figure 10.** A simple DAG with 3 variables.

In a matrix form, Eq. (1) can be written as:

$$X = (I - B)^{-1}E,$$

where  $I$  is the identity matrix,  $B$  is a matrix whose  $ij$ -th element is  $b_{ij}$  and  $E$  contains the non-Gaussian error terms. The acyclicity assumption implies that it exists a permutation of the rows and columns of matrix  $B$  so that the permuted matrix is strictly upper triangular.  $A := (I - B)^{-1}$  is called the mixing matrix.

Under the LiNGAM assumptions, i.e., the non-Gaussianity of the independent additive noises with the linear causal relationships, LiNGAM parameters can be proved to be identifiable, which ensures the possibility to recover the true causal structure, at least from a theoretical viewpoint. The ICA-LiNGAM estimation method then proposes to estimate  $A$  through Independent Component Analysis (ICA). However, the ICA optimization, typically based on iterative search, may get trapped in a local optimum, leading to potential computational instability. For this reason, we decided to use DirectLiNGAM [19] which implements an alternative non iterative optimization scheme, hence providing formal guarantees for convergence, while avoiding the specification of additional hyperparameters which would necessitate fine tuning. Briefly, in order to estimate the mixing matrix, DirectLiNGAM identifies exogenous variables through independence tests from the residuals of pairwise regressions. Then, the impact of the identified variables is removed from the others through least squares regression.

### Setup of simulations

We perform a total of 200 simulations of the considered electromechanical model after simultaneously sampling 10 model parameters of interest in their respective ranges (see Table 3), using a LHS scheme. Each simulation of the considered electromechanical model generates two output curves: left ventricular volume and pressure over time. Using these outputs, we compute the ejection fraction (EF), defined as the ratio of stroke volume to end-diastolic volume, and  $\max(dP/dt)$ , the maximum rate of change in left ventricular pressure. We are particularly interested in these two biomarkers to achieve a better model calibration and refine our knowledge on the input parameter ranges.

The biventricular geometry used in this study comprises 7 279 nodes and 31 330 tetrahedra. To hold the model in space and to represent the fibrous tissue around the valves, a spring force is added to all nodes surrounding the valves, so that these nodes are attached to their initial position but can still have a displacement. The value of 9 kPa for the spring stiffness associated to these nodes has

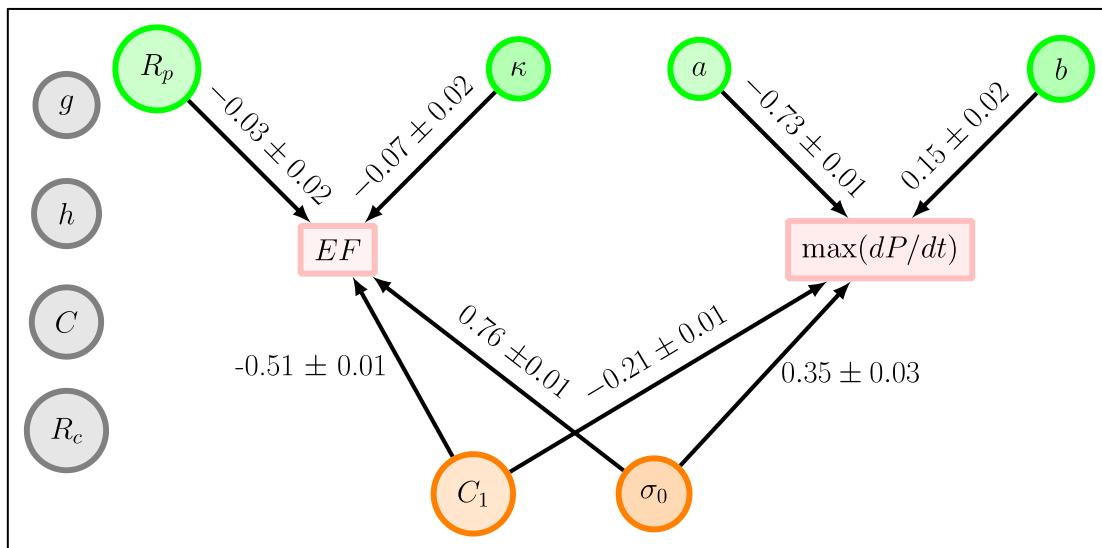
been used and seems to be a good trade-off between the mesh deformation and the numerical stability by keeping the mesh attached in space. The implementation of the whole dynamical system has been done in the open-source framework SOFA [20]. The simulated cases where the mathematical model didn't converge were excluded. Before entering each GSA pipeline, the final dataset of parameters and outputs have been standardized.

**Table 3.** Parameter ranges for LHS, and corresponding units.  $C_1, k$  govern the heart tissue properties;  $a, b, g, h$  define the stress profile in the mechanical model;  $R_p, R_c, C$  describe the arterial pressure;  $\sigma_0$  gives the peak contractility. nd: non dimensional.

$C_1$ (kPa)	$k$ (kPa)	$a$ (nd)	$b$ (nd)	$g$ (nd)	$h$ (nd)	$R_p$ (kPa·s/m <sup>3</sup> )
[1e4, 15e4]	[3e5, 3e6]	[0.5,4]	[1,8]	[0.5,3]	[1,4]	[1e8,1e9]
$R_c$ (kPa·s/m <sup>3</sup> )	$C$ (m <sup>3</sup> /kPa)	$\sigma_0$ (kPa)				
[1e5,4.5e6]	[9e-9,1e-8]	[1e5,1.5e6]				

### 2.2.3 Results and conclusions

Thanks to our causality-based GSA approach we have identified a subset of 6 parameters as the most impacting for the considered model's outputs, over the 10 initially investigated: results are depicted in Figure 11 through a DAG representation. As one can extrapolate from the graph, causal relationships are observed from  $R_p$  and  $k$  towards EF and from  $a$  and  $b$  towards  $\max(dP/dt)$ . Moreover,  $C_1$  and  $\sigma_0$  (the peak contractility) simultaneously influence both outputs.

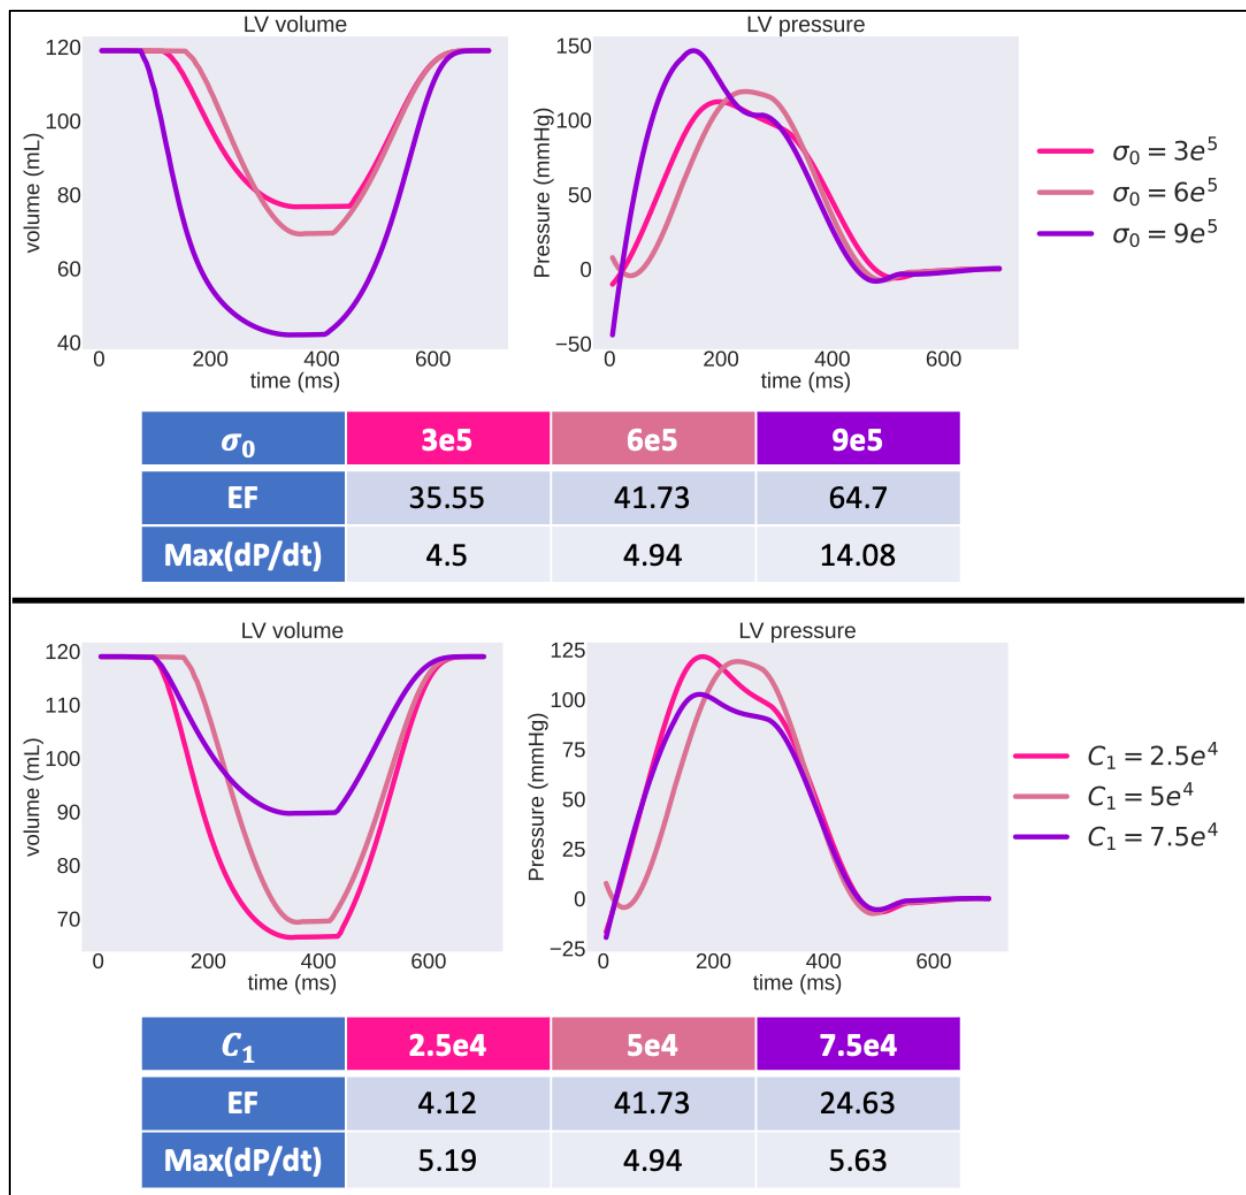


**Figure 11.** Causal graph obtained using DirectLiNGAM for the 10 parameters described in Table 10 and the two output of interest, EF and  $\max(dP/dt)$ .

The numerical values associated to the directed arrows in Figure 11 provide the mean causal weights across the 5 folds and the corresponding standard deviations: we interpret them as the causality-based sensitivity indices. To ease the lecture of the graph, nodes are colored depending on their role: green nodes represent the parameters which have been identified as influencing a

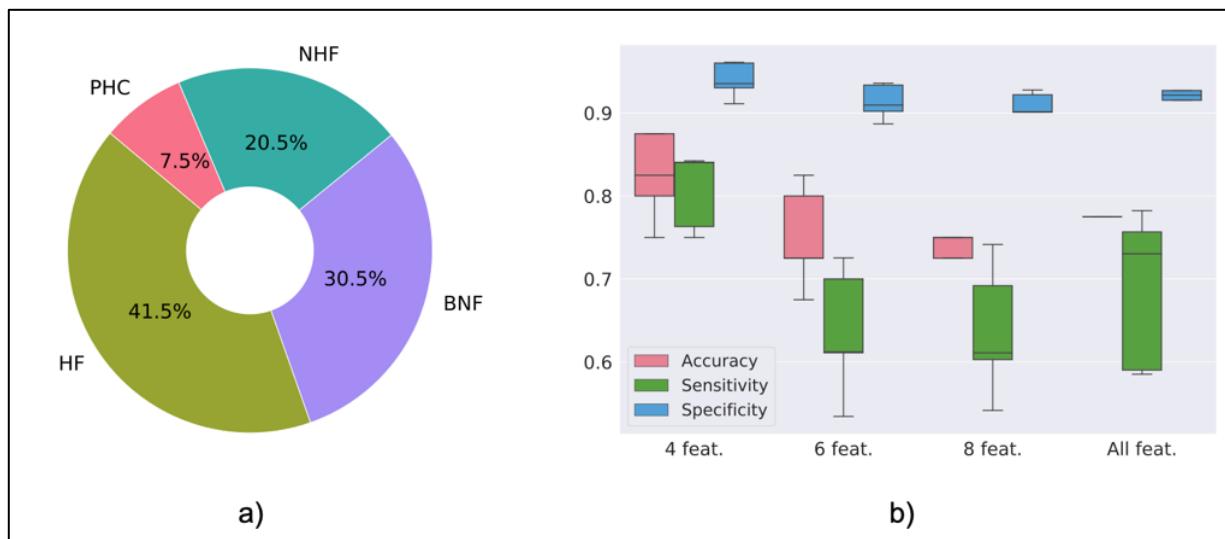
single outcome of interest (pink rectangles), while orange nodes indicate parameters that affect both outcomes; the gray nodes correspond to parameters that have no causal impact over the outcomes of interest.

The signs of the weights provide further insight on the expected behavior of parameter-output relationship. For instance,  $\sigma_0$  is found to positively influence both outcomes (i.e., increasing  $\sigma_0$  will result in higher EF and  $\text{max}(dP/dt)$ ). An example of this behavior is presented in Figure 12 (top panel), where we show results for three different values of  $\sigma_0$ . On the opposite, the material parameter  $C_1$ , that describes the stiffness of the tissue, negatively affect the two outputs of interest: as  $C_1$  increases, EF and  $\text{max}(dP/dt)$  decrease. Examples are shown in Figure 12 (bottom panel). This is an expected behavior, as stiffer tissue is less prone to contract, thus implying a decrease in both biomarkers.

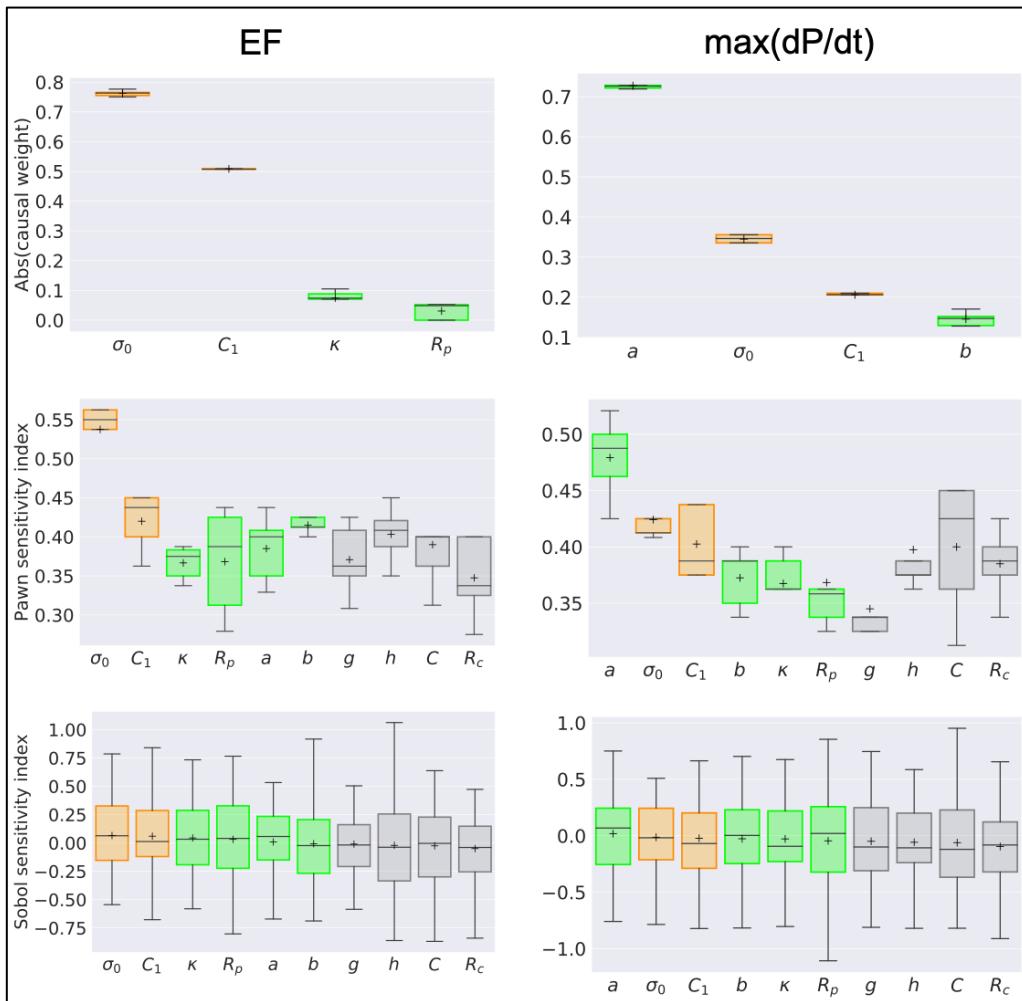


**Figure 12.** Effects of the variation of the peak contractility  $\sigma_0$  (upper panel), and the tissue parameter  $C_1$  (bottom panel) over the volume and pressure curves, and the extracted EF and  $\text{max}(dP/dt)$ .

To further validate our results in terms of parameters' selection, we perform multi-class classification considering the four features that have been identified as causally affecting EF ( $R_p$ ,  $k$ ,  $C_1$ ,  $\sigma_0$ ), and four classes of cardiac activity based on EF: normal heart function (NHF), where EF ranges from 55% to 70%; below-normal heart function (BNF), characterized by an EF between 40% and 54%; possible heart failure (HF), where EF is less than 40%, indicating impaired heart function; and potential hypertrophic cardiomyopathy (PHC), associated with an EF exceeding 70%. In Figure 13, we show the statistical distribution of EF values across our dataset and the classification performance obtained when accounting for several combinations of parameters, starting from the four model's inputs parameters identified as being directly related to EF, up to the inclusion of all available ones. The best performance is achieved using the set revealed in the causal graph, which validate our parameter selection.



**Figure 13.** a) EF classes proportions. b) Classification performances.



**Figure 14.** Sensitivity indices obtained with causal discovery (upper panel), Pawn (middle panel) and Sobol (bottom panel) for all input parameters included in this study (see Table 3). Left column: sensitivity of the EF outcome. Right column: sensitivity of the  $\max(dP/dt)$  outcome. Box plots are colored following the same code as in Figure 11. In the first row, parameters  $g, h, C, R_c$  have been excluded (no discovered causal effect).

We finally compare our results to two widely used GSA approaches: Pawn and Sobol. The obtained sensitivity indices with all methods are illustrated in Figure 14. One can see that the Pawn method provides equivalent sensitivity indices for most parameters, especially for  $\max(dP/dt)$ . For instance, it identifies parameters  $b, h$  and  $C$  as having an important effect on EF: this has been invalidated by independently performing single sensitivity analysis, further stressing the relevance of causal discovery to identify the most influential model parameters. Moreover, in this specific case where we deal with a large parameter space, a highly non-linear model, and a limited number of simulations, Sobol appears to be completely unreliable to identify the key parameters and their effect over our outputs of interest. Indeed, the mean value of all first-order Sobol's sensitivity indices is found to be zero. Finally, one can appreciate the stability of the results obtained with our method compared to both Pawn and Sobol.

To conclude, we have proposed to leverage causal discovery to perform global sensitivity analysis for a fast electromechanical model and considering two clinically significant biomarkers: ejection fraction (EF) and maximum rate of pressure change ( $\max(dP/dt)$ ). The obtained causal graph

reveals and quantifies the impact of the most relevant model parameters on the considered outputs, yielding stable results compared to classical global sensitivity analysis methods, despite a limited number of available simulations. Besides the electromechanical model considered here to illustrate our method, the proposed approach is general and could be relevant on others complex models, where parameters have a very indirect link with simulation outputs.

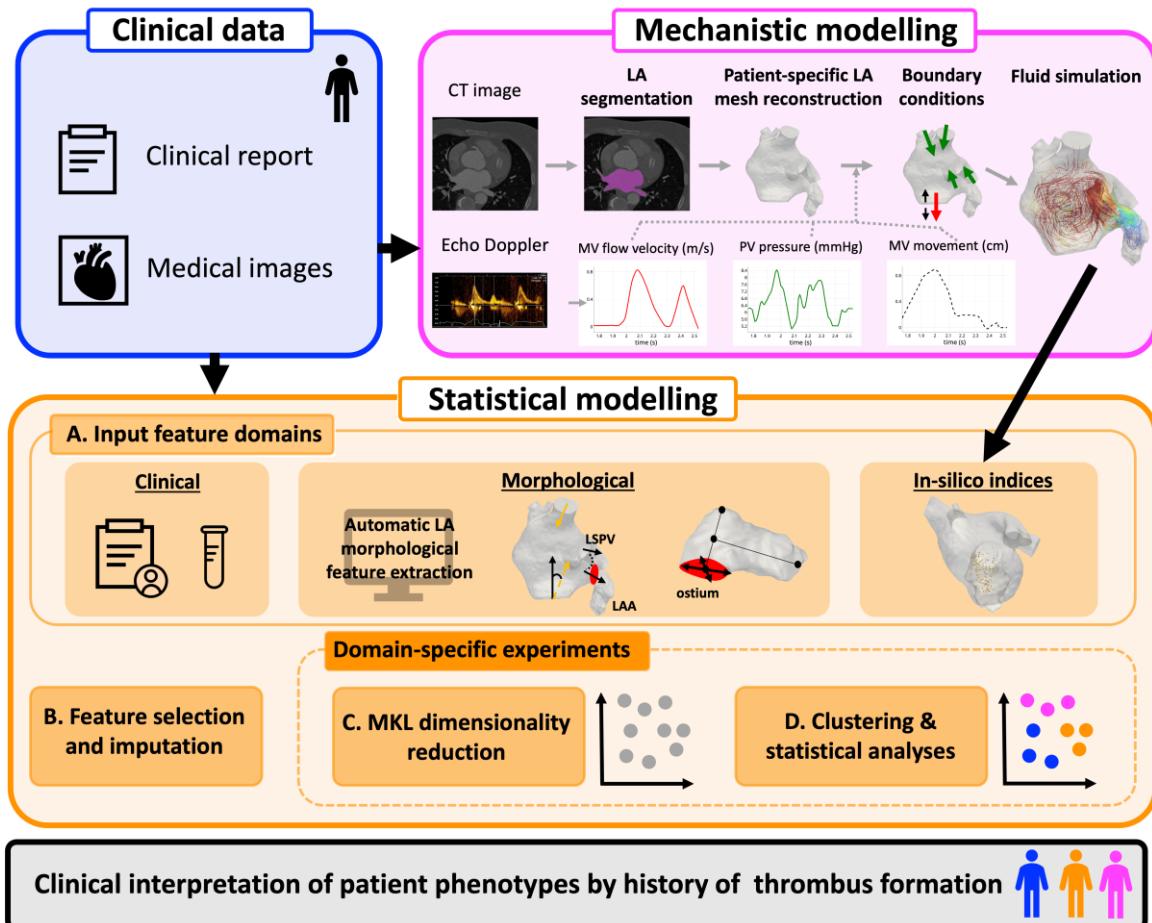
### **3. A digital twin approach integrating clinical, structural, and in-silico hemodynamic data to uncover stroke risk factors**

#### **3.1 Objectives**

In the context of stroke, digital twins can aid clinical decision-making across multiple medical specialties (e.g., neurology, cardiology) by enhancing the identification of risk factors linked to thrombus formation. While comprehensive stroke digital twins are highly complex, those focused specifically on the left atrium (LA) have already demonstrated their value in treatment planning and in improving our understanding of left atrial blood flow patterns in atrial fibrillation patients at risk of ischemic stroke [21,22,23]. The advent of mechanistic models using computational fluid dynamics (CFD) enables the simulation of blood flow in patient-specific atrial geometries [24,25]. From these models in-silico hemodynamic indices can be extracted, offering valuable insights into atrial blood flow behavior. In the following section, we present a Digital Twin that integrates multi-domain data to identify stroke risk factors in patients with non-valvular atrial fibrillation (AF). Clinical information, morphological features, and in-silico hemodynamic indices were jointly analyzed in a cohort of 130 patients using unsupervised machine learning to cluster them into distinct phenogroups based on their thrombus history. We hypothesized that integrating all available patient data within a digital twin framework would support the identification of phenogroups associated with cardioembolic thrombus. The results that will be presented here have been recently published in npj Digital Medicine [26].

#### **3.2 Methodology**

Figure 15 illustrates the proposed methodological pipeline. A Digital Twin of the left atrium was constructed from computed tomography (CT) images by segmenting the anatomy and extracting morphological parameters. A 3D mesh of the atrium was then used for patient-specific fluid simulations, from which hemodynamic indices—such as the number of fluid particles remaining inside the left atrial appendage (LAA)—were derived. Clinical data was integrated with morphological and hemodynamic features through the unsupervised multiple kernel learning (MKL) technique [27]. After applying feature selection, MKL was implemented to reduce the dimensionality of the feature space. From the low-dimensional output space, patients were clustered into phenogroups based on clinical, morphological, and hemodynamic similarities. These phenogroups were clinically interpreted and labeled according to increasing thrombus history (Phenogroups 0–2).



**Figure 15.** Scheme of the methodological pipeline. AF atrial fibrillation, LA: left atrium, LAA: left atrial appendage, PV: pulmonary vein, MKL: multiple kernel learning algorithm, CT: computed tomography [26].

### Study population

Clinical data was provided by Hospital Haut-Lévêque (Bordeaux, France), comprising medical records and pre-procedural CT scans from 130 patients with non-valvular atrial fibrillation (AF) scheduled for a left atrial appendage occlusion (LAAO) procedure. The study population was severely diseased, with high risk of thromboembolic events (mean CHA2DS2-VASc score of 4). A history of stroke/TIA or a detected LAA thrombus was observed in 44.6% of patients. The anticoagulation therapy was short-termed due to history of bleeding among patients (90.2%), with stroke/TIA generally occurring after its removal.

### Input features

Table 4 shows the clinical, morphological and in-silico hemodynamic features extracted. The features are divided by input features to the unsupervised machine learning algorithm ('Input features') and features used as an additional characterization of patients once they have already been stratified ('descriptive features').

### **- Clinical features**

History of thrombus—defined as a prior ischemic stroke or TIA, peripheral embolism, and/or LAA thrombus detected on preprocedural CT—was used as the label to stratify the patient population. Baseline characteristics, including age, body mass index (BMI), gender, hypertension, type 2 diabetes, vascular disease, dyslipidemia, type of atrial fibrillation (paroxysmal vs. non-paroxysmal), and HAS-BLED score, were extracted from electronic health records. Echocardiographic data such as left ventricular ejection fraction (LVEF) were also included. Heart failure was defined either by a clinical diagnosis or by reduced left ventricular systolic function—classified as severely reduced ( $LVEF \leq 40\%$ ) or moderately reduced ( $LVEF 41-49\%$ ). From pre-procedural medical records, laboratory parameters were obtained, including C-reactive protein (CRP), B-type natriuretic peptide (BNP), hematocrit, creatinine, glomerular filtration rate (GFR), urea, neutrophil-to-lymphocyte ratio, mean platelet volume (MPV), platelet count (PC), troponin, and fibrinogen levels.

### **- Morphological features**

An in-house computational pipeline was used for the automatic extraction of left atrium (LA) and left atrial appendage (LAA) geometric parameters. This algorithm derives patient-specific LA and LAA morphological measurements from chest CT scans.

### **- In-silico hemodynamic features**

Construction of in-silico hemodynamic models. The LA cavity was reconstructed from binary segmentations of CT images using semi-automatic region-growing methods available in Slicer 4.10.1. Computational fluid dynamics (CFD) simulations were performed using the ANSYS Fluent Solver 19.2 (ANSYS Inc., United States) to solve the Navier-Stokes equations via the finite volume method. Blood was modeled as an incompressible Newtonian fluid with a density of  $1060 \text{ kg/m}^3$  and a dynamic viscosity of  $0.0035 \text{ Pa}\cdot\text{s}$ . The pipeline described by Mill et al. [28] was followed. All simulation models incorporated two distinct boundary conditions (BCs): (1) a pressure inlet at the pulmonary veins (PVs), using a pressure waveform obtained from catheterization data of an AF patient in sinus rhythm; and (2) a velocity outlet at the mitral valve (MV), derived from Doppler echocardiography measurements collected from the study cohort. Electrocardiography recordings were used to determine the duration of each cardiac cycle. Furthermore, LA wall motion was simulated as a passive response to the longitudinal movement of the mitral valve annulus, employing a spring-based dynamic mesh approach.

In-silico hemodynamic parameters. Particle-based in-silico indices were derived from the CFD simulations to assess blood stasis in the left atrial appendage (LAA), following the methodology described by Mill et al. [28]. Massless tracer particles were released from the pulmonary veins, and their motion was governed by the simulated velocity vector fields originating from their respective seed locations. These particles were tracked continuously throughout the cardiac cycle. At each time step, 100 particles were injected—50 from each side of the left atrium (LA), uniformly distributed across the pulmonary veins. Two key hemodynamic indices were calculated: (1) the total number of particles remaining in the LAA at the end of the cardiac cycle, and (2) the particle age, defined as the duration a particle remains within the LA before exiting.

**Table 4.** Extracted features and their associated domains, organized into two categories: input features used for machine learning (ML) analysis and descriptive features used for post-ML interpretation.

Domain	Input ML variables	Descriptive variables
Clinical	Age, BMI*, BNP levels*, creatinine levels*, hematocrit levels*, ratio MPV/PC*, LVEF*	AF type, gender, HF, HT, diabetes type 2, vascular disease, CRP, ratio neutrophyl/lymphocite dyslipidemia, HAS-BLED, troponin, fibrinogen, urea, GFR
Morphological	Ostium area, ostium irregularity, LA volume, LAA centerline length, LAA tortuosity, LA sphericity, LAA bending angle, LAA/LSPV alignment angle (2D)	LAA volume/area/height/depth, ostium diameter/eccentricity/perimeter, number of LAA lobes, LA area/centreline length, LAA shape/alignment label, ratio ostium area/LAA volume, ratio LAA/LA volume, LAA/LSPV alignment angle (3D) PVs diameter and orientation
In-silico hemodynamic	Total LAA particles	Particle age

### Multi-domain data integration

#### - Feature pre-processing and dimensionality reduction

A total of 47 variables were extracted, encompassing clinical, morphological, and in-silico blood flow data (Table 4). Of these, 12 clinical variables contained missing values. To address this, missing data were imputed using the Random Forest method for mixed-type data, implemented via the 'missForest' package (version 1.4) in R [29]. An initial feature selection was performed, by not including features with more than 30% missing values to reduce potential bias, also, features with bivariate Pearson correlations  $> 0.7$  were not included. The in-silico index *particle age* was excluded from the machine learning inputs, as previous studies found no significant associations with thrombus formation. An unsupervised Multiple Kernel Learning (MKL) algorithm was employed [27]—a validated machine learning approach previously used in cardiac disease research. MKL integrates diverse feature types by assigning each one a dedicated Gaussian kernel, combining them into a unified low-dimensional representation. In this space, patients are positioned based on the similarity of their input features. The bandwidth of each Gaussian kernel was set to the square root of the number of patients ( $\sqrt{N_p}$ ), averaging distances over the  $\sqrt{N_p}$  nearest neighbors. As all input features were continuous, Euclidean distance was used across all kernels.

#### - Clustering

As a result, patients with similar characteristics were positioned closer together in the MKL-transformed space, improving interpretability and enabling effective clustering. K-means++ was chosen for clustering due to its simplicity and efficient centroid initialization. To enhance robustness, each clustering run included 20 iterations with different initializations. Six separate clustering analyses were performed using various combinations of clinical, morphological, and hemodynamic features to assess the contribution of each feature set. For each analysis, the optimal number of clusters and dimensions was determined using a grid search based on the average silhouette score, calculated across both K-means++ and hierarchical clustering. The number of clusters tested ranged from 3 to 4, balancing granularity and sample size, while dimensionality was varied from 2 to 5. Each analysis explored different combinations of input domains for MKL-based dimensionality reduction and clustering (see Table 5).

### - Statistics

Continuous variables were assessed for normality using the Shapiro-Wilk test and for homoscedasticity with the Levene test. When the null hypothesis was rejected at a significance level of  $\alpha = 0.05$ —indicating non-normality or unequal variances—the non-parametric Wilcoxon Mann-Whitney test was used for univariate comparisons, and results were reported as median with interquartile range. For variables meeting parametric assumptions, the Student's t-test was applied, with values expressed as mean  $\pm$  standard deviation. Differences across clusters were evaluated using ANOVA for normally distributed data or the Kruskal-Wallis test for non-parametric data. Post-hoc pairwise comparisons were performed using Tukey's HSD (parametric) or Dunn's test with Bonferroni correction (non-parametric). Categorical variables were analyzed using the chi-square test, with false discovery rate (FDR) correction applied for multiple testing. Statistical significance was set at  $p < 0.05$ . All analyses were conducted in R.

**Table 5.** Domain-specific experiments performed with unsupervised multiple kernel learning (MKL), with different combinations of morphological, hemodynamic and clinical domains.

Analysis	ML domain combinations
1	Clinical
2	Morphological
3	Morphological + clinical
4	In-silico hemodynamic + clinical
5	Morphological + in-silico hemodynamic
6	Morphological + in-silico hemodynamic + clinical

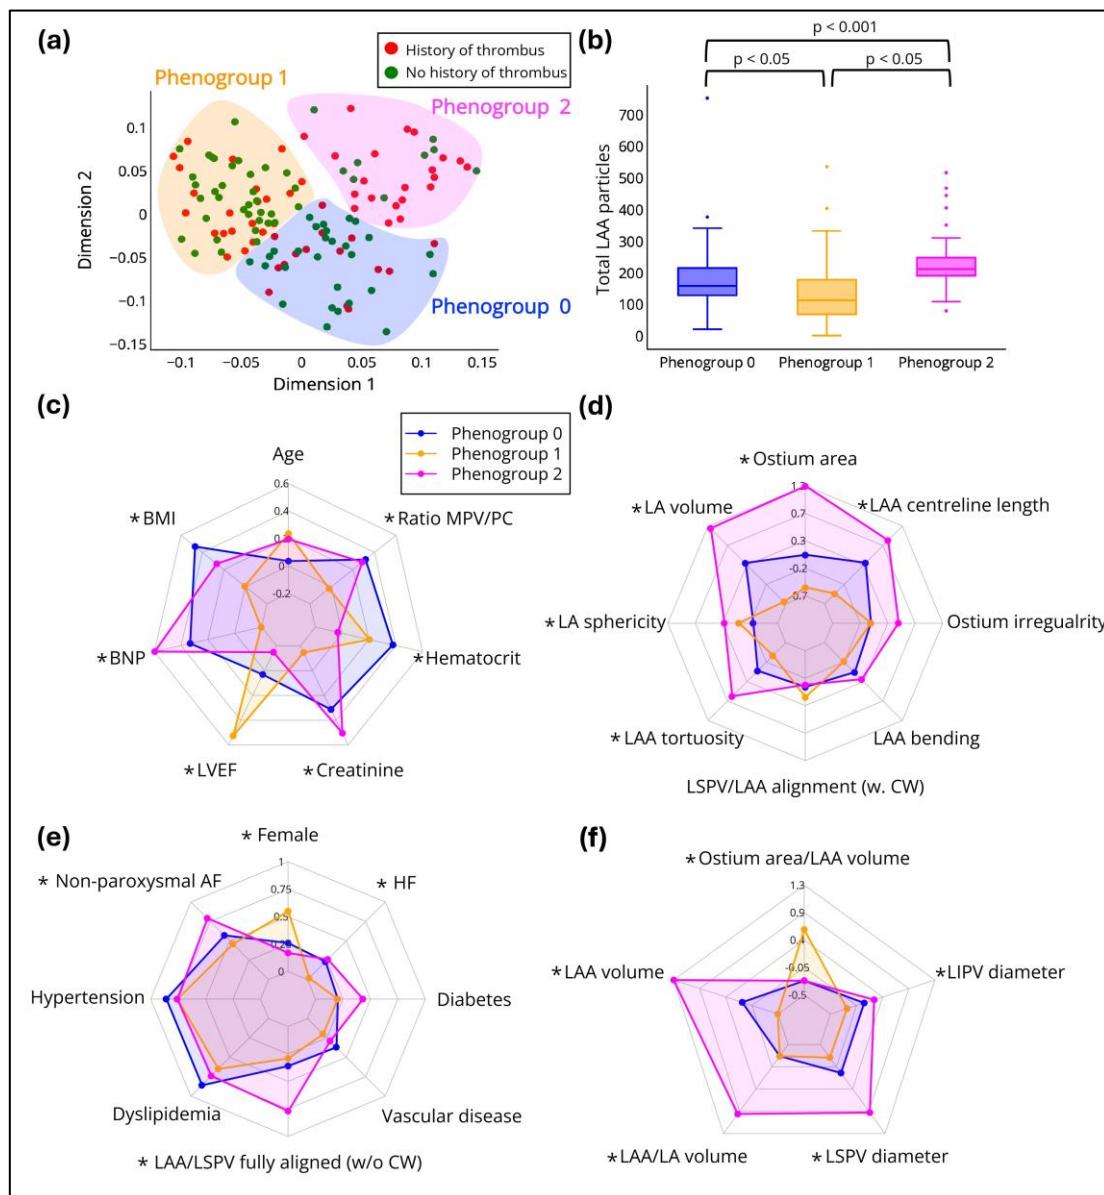
## 3.3 Results

The study included 130 atrial fibrillation (AF) patients with severely diseased left atria and high thromboembolic risk, incorporating both clinical data and medical imaging. A univariate statistical analysis was conducted to identify significant differences between patients with and without a history of thrombus. Among clinical variables, only dyslipidemia was significantly associated with thrombus history. Anatomically, patients with prior thrombus had larger LAA ostium dimensions and were more likely to have multi-lobed appendages ( $\geq 2$  lobes). In terms of in-silico hemodynamics, the total number of LAA particles was significantly higher in the thrombus group. Different combinations of input data domains were integrated using Multiple Kernel Learning (MKL) and analyzed through clustering techniques (see Table 5). Significant associations with thrombus history were identified in the clusters (phenogroups) derived from combining morphological and in-silico hemodynamic features (Analysis 5), as well as from the integration of clinical, morphological, and hemodynamic data (Analysis 6). Clusters based solely on morphological features (Analysis 2) were also further examined to explore their relationship with modeled hemodynamic parameters. Across all analyses, the configuration with three clusters in a two-dimensional space yielded the highest silhouette score.

**Digital twin clustering based on morphological factors.** Stratifying patients based solely on left atrial anatomical features revealed phenogroups with significant differences in the number of in-silico particles retained within the LAA, reflecting varying degrees of blood flow stasis. The phenogroup with the highest particle retention was characterized by larger ostium dimensions, increased LA and LAA volumes, and a strong alignment between the left superior pulmonary vein (LSPV) and the LAA.

**Digital twin clustering based on morphological and hemodynamic factors.** Including the number of in-silico LAA particles as an input feature led to the identification of phenogroups with significant differences in thrombus history. The phenogroup with the highest proportion of prior thrombus cases (64%) showed the greatest number of fluid particles inside the LAA, larger LA, LAA, and ostium areas, and stronger alignment between the LAA and the left superior pulmonary vein (LSPV). This group also exhibited less LAA bending—suggestive of non-chicken-wing morphologies—and elevated B-type natriuretic peptide (BNP) levels. Among all features, ostium area was the most predictive of phenogroup classification. Interestingly, thrombus cases within the phenogroup with the lowest overall thrombus prevalence were associated with a higher incidence of dyslipidemia, increased number of LAA fluid particles, and a steeper right superior pulmonary vein (RSPV) inflow angle.

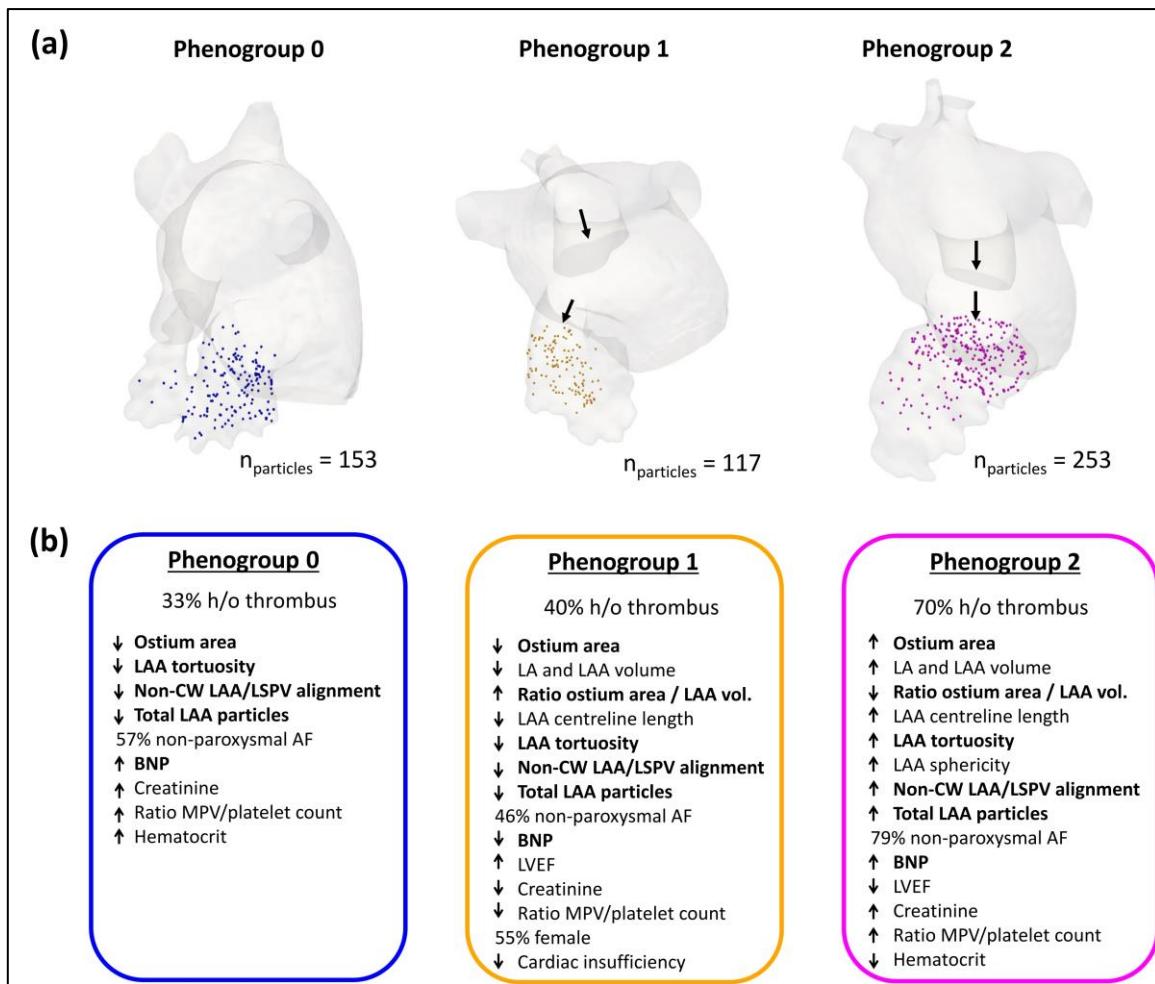
**Digital twin clustering based on clinical, morphological and hemodynamic factors.** Integrating clinical data with in-silico hemodynamic and morphological features enhanced patient stratification by thrombus history, identifying a phenogroup with a 70% prevalence of prior thrombus formation. The clustering analysis resulted in three distinct phenogroups, each with unique characteristics (Figure 16). Phenogroup 2 (n = 30) had a significantly higher proportion of patients with a history of thrombus compared to Phenogroup 0 (70% vs. 33%, p = 0.016) and Phenogroup 1 (70% vs. 40%, p = 0.03). As shown in Figure 16, this group was characterized by significantly larger left atrial (LA) volumes and ostium areas (though not a higher ostium area-to-LAA volume ratio), increased LAA anatomical complexity, greater LA sphericity, and stronger alignment between the LAA and the LSPV in non-chicken wing morphologies (Figures 16d and 16f). Phenogroup 2 also had the highest number of fluid particles retained within the LAA (Figure 16b). Figure 17a illustrates the residual particle distribution in the LAA for representative patients from each phenogroup, while Figure 17b summarizes the key clinical, anatomical, and hemodynamic characteristics defining each group.



**Figure 16.** Clustering results based on clinical, morphological, and hemodynamic features. (a) Cluster distribution showing patients with (red circles) and without (green circles) a history of thrombus. (b) Distribution of number of in-silico LAA particles by cluster. (c-f) Radar plots displaying median normalized values per cluster: (c) demographic and laboratory features, (d) anatomical (shape-related) markers, (e) categorical descriptive variables (shown as %), and (f) continuous descriptive variables. The blue, orange, and magenta lines represent Phenogroups 0, 1, and 2, respectively. Categorical variables are displayed as percentages (%), and statistically significant differences ( $p < 0.05$ ) are marked with an asterisk (\*). AF: atrial fibrillation; BMI: body mass index; BNP: B-type natriuretic peptide; CW: chicken wing; MPV: mean platelet volume; PC: platelet count; LVEF: left ventricular ejection fraction; LA: left atrium; LAA: left atrial appendage; LIPV: left inferior pulmonary vein; LSPV: left superior pulmonary vein. [26].

Clinically, Phenogroup 2 showed the highest BNP levels, MPV/PC ratio, creatinine (gender-adjusted), and non-paroxysmal AF prevalence, along with lower LVEF and hematocrit (although both within physiological ranges). Phenogroups 0 ( $n = 42$ ) and 1 ( $n = 58$ ) had fewer thrombus cases and distinct clinical profiles (Figure 16c, 16e). Phenogroup 1 had significantly lower BNP and MPV/PC ratios compared to Phenogroup 2, while Phenogroup 0 differed mainly in anatomical features—smaller ostium area, LA sphericity, LAA volume, and reduced LAA/LSPV alignment. Within Phenogroup 2,

patients without prior thrombus ( $n = 9$ ) had lower LA volume, BMI, hematocrit, and hypertension rates, and less steep LIPV angles. Conversely, thrombus cases in Phenogroup 0 had significantly higher MPV.

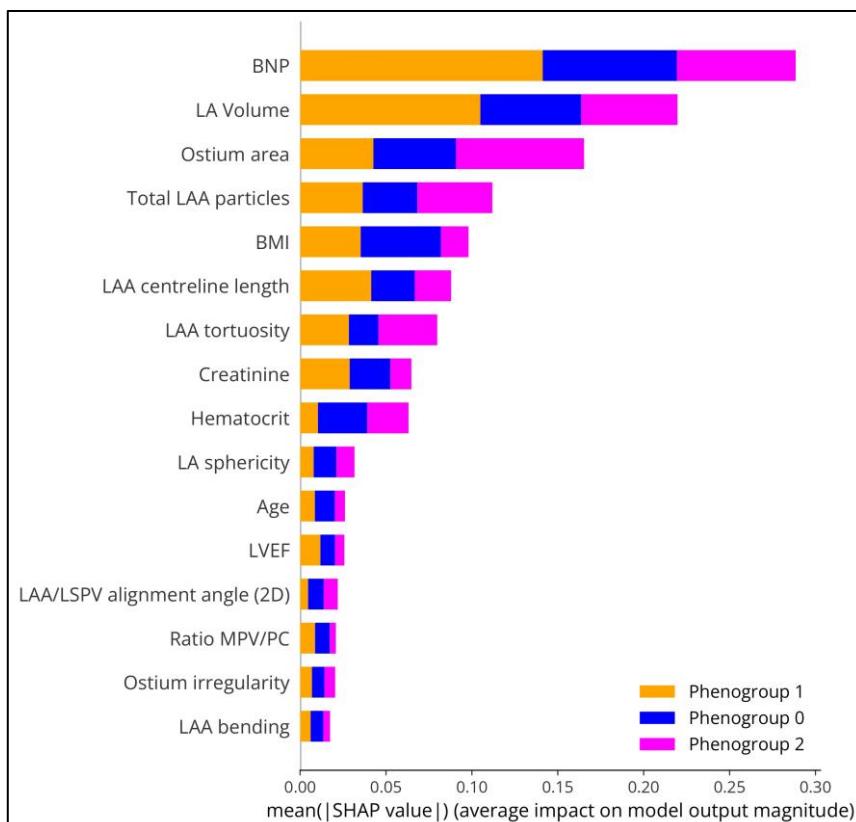


**Figure 17.** Phenogroup comparison based on clinical, morphological, and hemodynamic features. (a) Visualization of residual computational fluid dynamics (CFD) particles within the left atrial appendage (LAA) at the final time step (end-systolic phase) for representative patients from each phenogroup, illustrating typical particle counts ( $n_{particles}$ ). Black arrows indicate the degree of alignment between the LAA and the left superior pulmonary vein (LSPV). (b) Summary of the key clinical, anatomical, and hemodynamic characteristics defining the three identified phenogroups. [26].

To improve interpretability, a Random Forest classifier was trained using the cluster assignments as labels. Feature importance was then evaluated using Shapley Additive exPlanations (SHAP), which highlighted the most influential variables for phenogroup differentiation. As shown in Figure 18, BNP, LA volume, and ostium area emerged as the top contributors, with the total number of LAA particles also playing a relevant role. BNP was particularly important in distinguishing Phenogroups 0 and 1, whereas ostium area was the main driver for identifying Phenogroup 2.

### 3.4 Discussion and conclusions

The development of left atrial digital twins is essential for advancing stroke treatment and prevention in patients with atrial fibrillation (AF). In this study, we integrated multi-modal data—including demographics, blood biomarkers, and LA morphological features—with indices derived from computational fluid dynamics (CFD) simulations of left atrial blood flow to construct digital twins for a cohort of severely diseased AF patients, many of whom had a history of thrombus. We analyzed patient phenogroups based on various combinations of clinical, anatomical, and hemodynamic characteristics to uncover underlying relationships. Each phenogroup was then clinically interpreted, with a particular focus on the proportion of patients who had experienced thrombus formation. Our findings show that combining clinical, morphological, and circulatory data improves patient stratification based on prior thrombus formation. In conclusion, this descriptive study emphasizes the value of integrating clinical, morphological, and hemodynamic data to better understand the multifactorial drivers of thrombus formation in patients with advanced atrial fibrillation. By leveraging unsupervised machine learning, we demonstrate how multi-domain data integration can uncover distinct phenogroups and support thrombus risk stratification. Early identification of high-risk phenotypes—prior to left atrial appendage occlusion (LAAO)—can inform clinical decisions, such as intensifying monitoring or tailoring anticoagulation strategies, especially in patients with contraindications to oral anticoagulants. Our preliminary findings suggest that specific anatomical features, such as smaller ostium dimensions, atrial remodeling, LAA tortuosity, and certain pulmonary vein configurations (e.g., low LAA–LSPV alignment), may improve local flow dynamics by reducing stasis in the appendage. These morphological characteristics could offer protective effects even in clinically pro-thrombotic individuals. Future work should focus on refining this integrative framework to develop novel risk scores or predictive tools for thrombus formation.



**Figure 18.** Feature importance derived from SHAP analysis, incorporating clinical, morphological, and hemodynamic variables. The bar chart presents the average absolute SHAP value for each input feature, with taller bars representing a greater impact on the classification of phenogroups. Features are listed in descending order of importance. Abbreviations: BNP: B-type natriuretic peptide, BMI: body mass index, LA: left atrium, LAA: left atrial appendage, LSPV: left superior pulmonary vein, LVEF: left ventricular ejection fraction; MPV – mean platelet volume; PC – platelet count.

## 4. Causal learning for in-silico biomarkers analysis and TdP risk assessment

Drug-induced Torsade-de-Pointes (TdP) is one of the most frightening drugs' side effects, which can provoke ventricular fibrillation and lead to sudden death: consequently, pro-arrhythmia assessment is needed in the preclinical stages of drug development. TdP is closely related to abnormal repolarizations in single cardiac cells, hence by a prolongation of the QT interval. The human Ether-à-go-go-Related Gene (hERG) is responsible for the rapid component of the delayed rectifier current ( $I_{Kr}$ ) which is one of the major repolarizing currents in the heart [30]: proarrhythmic risk assessment of drugs is traditionally based on the evaluation of the hERG channel and the measure of the delayed ventricular repolarization on the electrocardiogram (ECG), i.e. the QT interval prolongation [31]. This method can accurately classify high TdP-risk drugs through a single analytical assessment considering a unique ionic channel and exclusively focusing on ventricular repolarization [32]. However, it has been observed (e.g. [33]) that  $I_{Kr}$  blocking may be not sufficient to assess drug-induced TdP risk and is prone to produce false positives.

The limited accuracy of strategies focusing on hERG block and QT prolongation has been at the basis of the [Comprehensive in Vitro Proarrhythmia Assay \(CiPA\)](#) initiative, promoting in-silico simulations to improve predictions by providing a mechanistic classification [34]. Electrophysiological models considering the blockage of the seven ionic currents identified by CiPA ( $I_{Kr}$ ,  $I_{Na}$ ,  $I_{NaL}$ ,  $I_{CaL}$ ,  $I_{K1}$ ,  $I_{Ks}$  and  $I_{to}$ ), instead of only the main repolarizing current  $I_{Kr}$ , are used to compute biomarkers able to better identify torsadogenic drugs. Consequently, recent studies have been proposed to incorporate additional ion channels other than  $I_{Kr}$  for a more reliable drug safety assessment, and several ion channels' combinations have been proposed as the most effective to assess drug-induced TdP risk [35,36,37].

Besides ionic currents blockade, in-silico electrophysiological and mechanical biomarkers can provide mechanistic proarrhythmic information for TdP-risk assessment. Indeed, detailed and complex biophysical models have been developed to better represent myocyte activity, including excitation-contraction coupling, the process by which changes in membrane potential trigger calcium release to activate myofilaments. Therefore, drug effects on ionic currents can be translated from excitation to contraction. So far, multiple indices have been proposed for TdP-risk assessment [38], most of them based on drug properties and action potential, and the latest increasing tendency is to combine several electrophysiological properties, including Ca-derived features [39,40]. Besides, instead of striving to find a single optimal predictor, the strategy of combining multiple features and analyzing them with machine learning tools seems promising [41].

Causal learning is recently attracting increasing interest in the e-medicine community as a powerful theoretical grounded tool to go beyond correlation-type relationships and promote interpretability and actionability. Indeed, the query of causality is of paramount importance in biomedical data analysis: assessing the causal relationships between the observed variables allows to improve our understanding of the tackled medical condition and better support decision-making. For instance, understanding the causal relationships between risk factors and disease can help develop effective prevention and treatment strategies and prioritize the

available information to improve risk assessment tools. In this section we present the work done on the characterization and assessment of drug-induced TdP risk, combining causal learning and in-silico derived biomarkers. The first study [42], described in Section 4.1, concerns the discovery of a causal structure relating drug-induced ionic channels' perturbations and TdP risk, with the objective of exploiting the obtained Bayesian causal network to infer the downstream impact of the ionic currents in the drug's safety. The second study [43], described in Section 4.2, lies in a multichannel context, with ionic currents blockades, torsadogenic indices, and electrophysiological biomarkers being three distinct channels, and aims at identifying between-channels causal relationships at a high level through the definition of a channel-specific latent representation. The objective is to effectively combine the different sources of information (channels) and achieve a more informative TdP-risk characterization.

## 4.1 Linear non-Gaussian acyclic model for ion channel blockade

### 4.1.1 Objectives

Discovering causal relationships between drug-induced ionic channels' perturbations could shed new light on the underlying mechanisms leading to TdP, and drive variable selection to improve TdP-risk assessment. This can be achieved through causal discovery, a branch of causal research whose aim is to learn the cause-effect relationships from data, which can be typically represented through Directed Acyclic Graphs (DAGs), a graph where nodes denote variables at hand, directed edges the causal relationships relating them, and which does not present any cycle (no pattern starting and ending in the same node).

Among classical existing methods for solving the task of learning such DAGs [15], we propose to apply the causal discovery method ICA-Linear Non-Gaussian Acyclic Model (ICA-LiNGAM) to uncover the relationships across the 7 ion channels identified by the Comprehensive in vitro Proarrhythmia Assay (CiPA) initiative as potentially related to the induction of TdP:  $I_{Kr}$ ,  $I_{Na}$ ,  $I_{NaL}$ ,  $I_{CaL}$ ,  $I_{K1}$ ,  $I_{Ks}$  and  $I_{to}$ . The obtained Bayesian causal network can be then explored to infer the downstream impact of the ionic currents in the drug's safety label. We further assess the relevance of our causality-based selection of ion channels by independently performing binary drug risk classification.

### 4.1.2 Methodology

We briefly recall here the assumptions and properties of a Linear non-Gaussian Acyclic Model (LiNGAM) [18] (already described in Section 2.2 of this deliverable). LiNGAM assumes that the causal relationships between variables can be represented by means of a directed acyclic graph (DAG) where the causal functions between each variable and its parents' nodes (i.e., the variables that causally precede it) are linear, and the error term associated to each variable follows a non-Gaussian distribution. Under these assumptions the theoretical identifiability of the model can be ensured. Denoting  $X$  the matrix of observations for each variable, then:

$$X = (I - B)^{-1}E,$$

where  $I$  is the identity matrix,  $B$  is a matrix whose  $ij$ -th element provides the strength of the linear causal relationship of the  $j$ -th variable over the  $i$ -th one, and  $E$  contains the non-Gaussian error terms. In order to estimate the mixing matrix  $A := (I - B)^{-1}$ , we use DirectLiNGAM [19] a well-established method for the estimation of LiNGAM models. DirectLiNGAM is based on Independent

Component Analysis (ICA) and implements a direct non-iterative optimization scheme, which provides formal convergence guarantees.

We consider a total of 109 drugs from [CredibleMeds](#) with known torsadogenic risk (37 with known risk, 14 with possible risk, 13 with conditional risk and 45 with no proven risk). For every drug, we consider two pharmacological data: the IC<sub>50</sub> for each of the seven ionic currents that have been proven to be potential for TdP induction ( $I_{Kr}$ ,  $I_{Na}$ ,  $I_{NaL}$ ,  $I_{CaL}$ ,  $I_{K1}$ ,  $I_{Ks}$ ,  $I_{to}$ ), and the effective free therapeutic plasma concentration (EFTPC), defined as the drug concentration in the plasma required to produce the desired therapeutic effect in the body. This information is combined to define the ion channels blocked fraction, here denoted by  $Bf_{Ion}$ :

$$Bf_{Ion} = \left[ 1 + \left( \frac{IC_{50}^{Ion}}{EFTPC} \right)^h \right]^{-1},$$

where  $h$  denotes the Hill coefficient, the number of drug molecules assumed to be sufficient to block an ion channel.

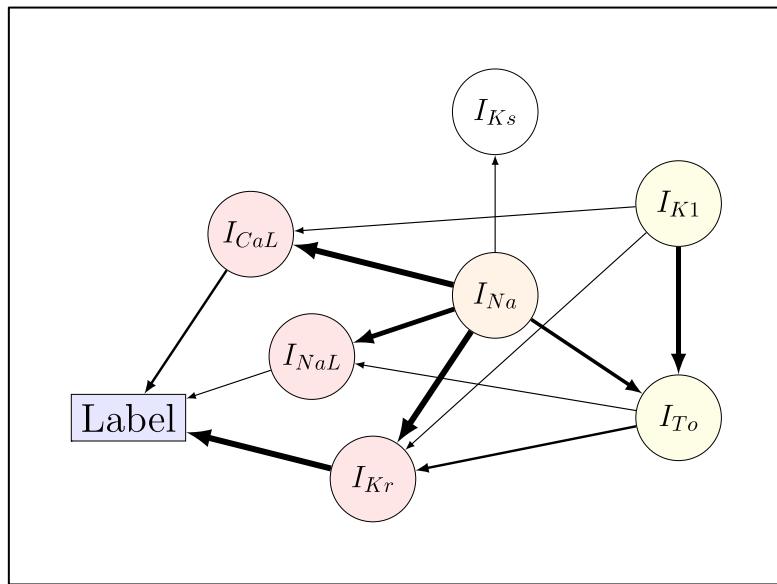
DirectLiNGAM is applied on the ion channels blocked fractions for the seven considered ionic currents; an extra binarized *Label* node (i.e., drugs of confirmed or possible TdP-risk are labelled as unsafe, and drugs with conditionally or no proven TdP risk as safe) is included as well to investigate the relationships of the blockade parameters with the known proarrhythmic risk. We implemented a 5-fold cross-validation over the 109 drugs and keep any directed arrow that was discovered at least on 2 folds.

We further use the obtained causal graph to inform about variable selection for TdP risk classification. Specifically, the nodes which appear as being consistently associated to the *Label* node were used as input of classical machine learning (ML) models to perform an independent binary classification of the considered 109 drugs. In particular, two classical ML classifiers were used: RandomForest and K-Nearest-Neighbors, whose prescriptions were finally combined through a majority voting classifier [44]. A 5-fold cross-validation was performed as well for the classification task.

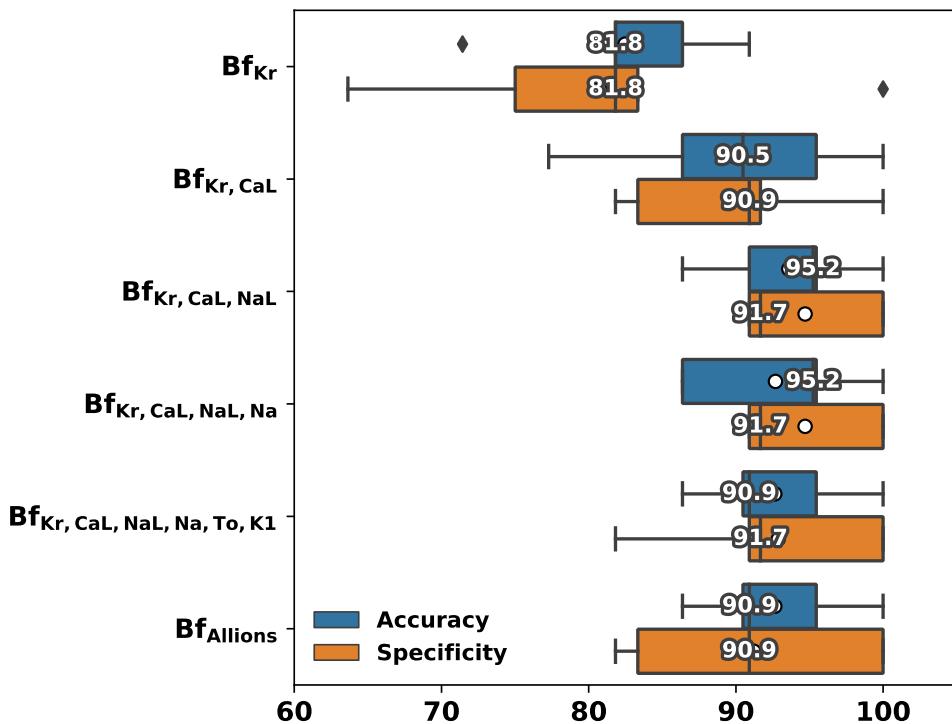
#### 4.1.3 Results and conclusions

In Figure 19 we show the causal graph obtained by applying DirectLiNGAM to the ion channels blocked fractions for the seven considered ionic currents: the arrow thickness represents the number of occurrences of the causal arrows over the 5 folds. One can see that the  $I_{Kr}$  channel is systematically identified as a direct cause leading to the drug label, in accordance with the established role it plays for repolarization. Also,  $I_{CaL}$  and  $I_{NaL}$  ionic currents appear both as directly affecting the drug label, even if with a milder relevance with respect to  $I_{Kr}$ . Despite the fact that fast Na channel is not affecting the node label directly, it still plays an important and central role since it appears to be causally related to  $I_{Kr}$ ,  $I_{CaL}$ , and  $I_{NaL}$  ion currents.

Our causality-based results were further confirmed by performing an independent binary classification of the 109 drugs used for this study. In Figure 20 we present the accuracy and specificity scores for the ion channels combinations revealed by the obtained causal graph (Figure 19), starting from the most stably identified parent of the label node,  $I_{Kr}$ , up to its farthest ancestors.



**Figure 19.** Causal graph obtained relating the blocked fractions for the seven considered ionic channels (round nodes) and the TdP risk label (square node), through DirectLiNGAM. We highlight the nodes hierarchy in relation to the target “Label” node, by coloring nodes from red (directed parents) to white (no directed path).



**Figure 20.** Accuracy and specificity for drug-induced TdP risk classification based on the combinations of ions as discovered by DirectLiNGAM

The combination of the 3 ions identified through DirectLiNGAM,  $I_{Kr}$ ,  $I_{CaL}$ ,  $I_{NaL}$ , maximizes the classification accuracy and specificity, while a progressive degradation of both metrics is observed when other channels are increasingly added, indicating that the causal graph has been effective in identifying the main parameters which directly inform on the drug safety classification.

**Table 6.** Mean values of the AUC, accuracy, sensitivity and specificity of classification based on the ion channels selected with DirectLiNGAM compared to other state-of-the-art methods which use ion channels for proarrhythmic risk classification.

Ion Combination	AUC	Accuracy	Specificity	Sensitivity	Ref.
$I_{Kr}$ , $I_{CaL}$ , $I_{Na}$	0.91	0.91	0.88	0.87	[36]
$I_{Kr}$ , $I_{CaL}$ , $I_{Na}$	NA	0.87	0.73	0.89	[37]
$I_{Kr}$ , $I_{CaL}$ , $I_{NaL}$	0.94	0.94	0.92	0.95	Ours

Table 6 finally summarizes the mean performances of our causality-based classifications with respect to some state-of-the-art results [36,37]. Our ion-combination selection shows the best mean AUC and specificity scores, outperforming the other methods.

These results emphasize the importance of taking into consideration ion currents  $I_{CaL}$  and  $I_{NaL}$  in addition to  $I_{Kr}$  to improve drug induced TdP-risk classification, and highlights the relevance of a causal discovery approach to infer ion channels selection for the drug-induced TdP risk. Our results are consistent with previous works (e.g., [45]) where the authors showed the crucial role of  $I_{Kr}$ ,  $I_{CaL}$  and  $I_{NaL}$  in the computation of *in-silico* arrhythmogenic biomarkers proposed for TdP risk assessment.

## 4.2 Multimodal causal VAE with *in-silico* electrophysiological biomarkers

### 4.2.1 Objectives

The increasing complexity and efficacy of electrophysiological models has motivated the analysis of multiple indices for TdP-risk assessment, which has shown promising. Nevertheless, the question of effectively combining the different sources of information (*channels*) to achieve a more informative TdP-risk assessment remains open, since the process of analyzing multi-channel data is challenging by nature, due to data heterogeneity and the potential presence of redundant shared information.

We propose a novel model based on Variational Autoencoders (VAEs) [46], called Multi-Channel Causal Variational Autoencoder (MC<sup>2</sup>VAE), to identify the hidden (latent) causal relationships between ionic currents blockades, torsadogenic indices, and electrophysiological biomarkers, considered here as three distinct channels which can contribute to inform on drug-induced TdP risk. Our main hypothesis is that discovering causal relationships between the different channels could help our understanding of the underlying mechanisms leading to TdP, and improve its characterization. As a final validation, we analyze the causal latent projections of the considered

channels in terms of their ability to discriminate between safe and unsafe drugs (considering the same binarized label as done in Section 4.1).

#### 4.2.2 Methodology

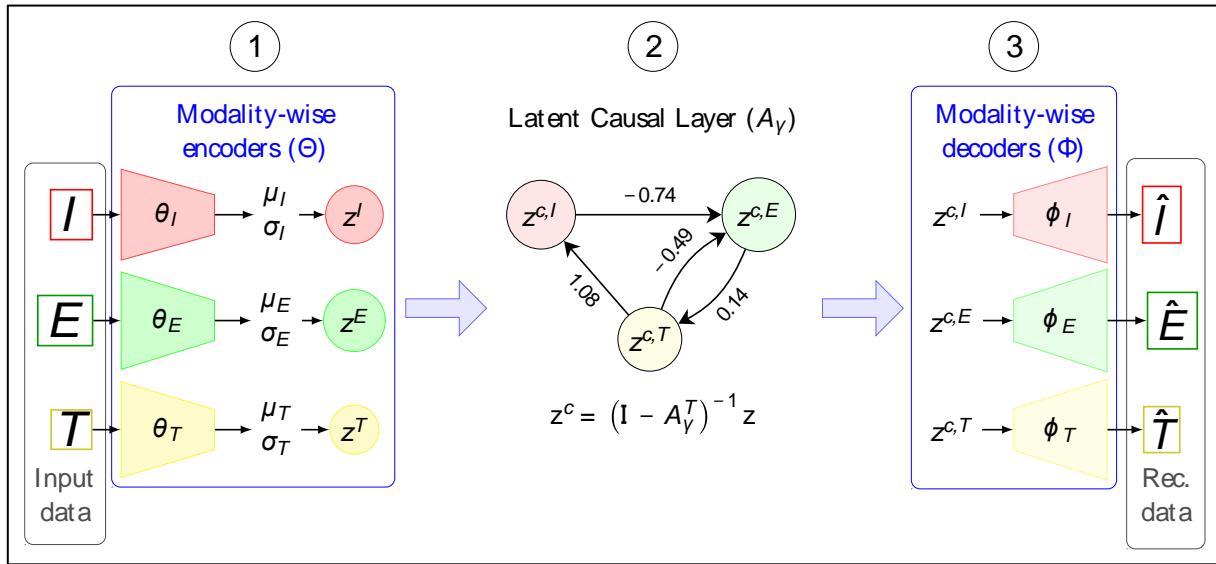
We consider the dataset of 109 drugs from [CredibleMeds](#) with known torsadogenic risk, as in Section 4.1. We perform electrophysiological cellular simulations with each drug using a modified version of the human endocardial ventricular O'Hara et al. model [47]. For each drug we dispose overall of data coming from 3 channels, denoted by  $I$ ,  $E$ ,  $T$ , and described below.

- $I$ : The drug-induced blockade of the seven most important ionic currents according to the CiPA initiative, computed using the simple pore block model, at the effective free therapeutic plasma concentration (EFTPC) [42].
- $E$ : In-silico electrophysiological biomarkers that consist of direct features from the action potential (APD90, APD50, Tri9050, Tri9030, qNet) and from the calcium transient signals (Casyst, Cadiast, CaTD90, CaTD50), including a surrogate of the electromechanical window ( $Emw = CaTD90 - APD90$ ) [41].
- $T$ : In-silico derived features that have been proposed as torsadogenic indices:  $T_x$ , the ratio between the concentration of a drug that provokes a 10% prolongation of the APD90 in control conditions and the EFTPC;  $T_{qNet}$ , the ratio between the net charge carried by the net current when exposed to 10 times the EFTPC with respect to the net charge in control conditions; and  $T_{triang}$ , the ratio between Tri9030 for a drug concentration of 10 times EFTPC and triangulation in control conditions [45].

A power transformation has been applied to  $I$  and a logarithmic transformation to  $T$  and  $E$ , followed by standardization.

We denote by  $X = \{I, E, T\}$  the dataset consisting of the three channels mentioned above;  $X_i := (I_i, E_i, T_i)$  represents the dataset of the  $i$ -th drug.

To analyze such data, we propose a method based on Variational Autoencoders (VAEs), Bayesian generative models which are composed of two main blocks, an encoder which projects data into a meaningful lower-dimensional latent space and learns its distribution, and a decoder which performs the inverse transformation and generate back the data in their original space from their latent representation. VAEs are highly flexible models, capable of adapting to a wide variety of data types and structures, and have already shown promise in the context of multi-channel data analysis [48].



**Figure 21.** The structure of MC<sup>2</sup>VAE, and its application to ionic currents blockades ( $I$ ), electrophysiological biomarkers ( $E$ ), and torsadogenic indices ( $T$ ). Encoders and decoders are neural networks (in this case, simple linear layers) parameterized by  $\Theta$  and  $\Phi$  respectively. They're outputs are a mean and a standard deviation ( $\mu, \sigma$ ) of an estimated gaussian distribution.  $A_\gamma$  is the weighted matrix describing the linear causal dependencies between the latent modality-specific projections.

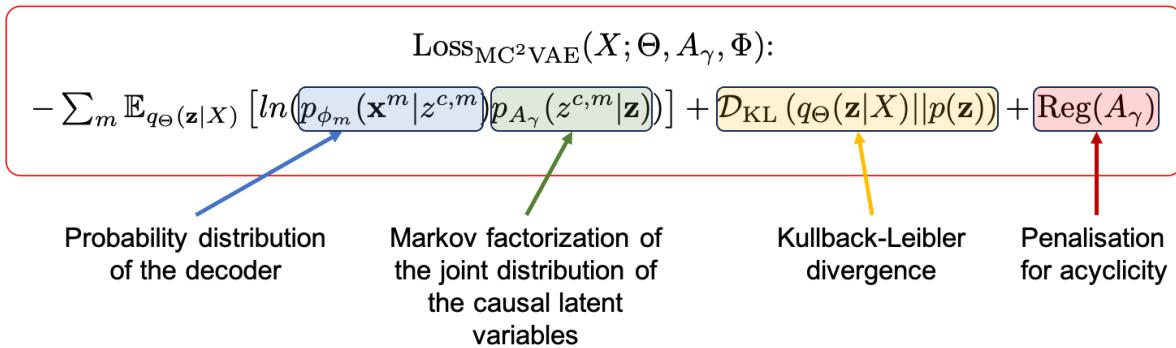
MC<sup>2</sup>VAE is an unsupervised model, whose structure, conversely to standard VAEs, is composed of three main components: 1) encoding, 2) causal layer, and 3) decoding (Figure 21). Each channel is projected into a one-dimensional latent space ( $z := (z^m)_{m=I,E,T}$ ) through its channel-specific encoder, a Neural Network (NN) parameterized by  $\Theta := (\theta_m)_{m=I,E,T}$ . The obtained latent variables are fed to the causal layer. We assume independence between the drugs and a normal prior for their respective noise terms,  $p(z)$ . Moreover, we hypothesize linear causal relationships across the channel-specific latent variables. These assumptions define the following latent structural causal model, whose learnable parameters are contained in the  $m$ -dimensional matrix  $A_\gamma$ :

$$z^c = A_\gamma^T z^c + z = (I - A_\gamma^T)^{-1} z, \quad z \sim \mathcal{N}(0, I)$$

$(A_\gamma)_{ij}$  provides the strength of the causal linear relationships of the  $i$ -th latent parent variable on the  $j$ -th latent children variable, with  $i, j$  in  $\{I, E, T\}$ .

Finally, the causal latent variables  $z^c := (z^{c,m})_{m=I,E,T}$  are fed to channel-specific decoders, NNs parameterized by  $\Phi := (\phi_m)_{m=I,E,T}$ , for the reconstruction.

To optimize the parameters of MC<sup>2</sup>VAE, we should maximize the marginal log-likelihood of  $X$ ,  $\mathcal{L}(X; \Phi, \Theta, A_\gamma)$ , and derive the true posterior  $p(z|X)$  over the latent space. Due to analytical intractability, we apply variational Bayes and introduce a tractable posterior  $q_\Theta(z|X)$  which approximates  $p(z|X)$  [49]. We finally get to the following loss:

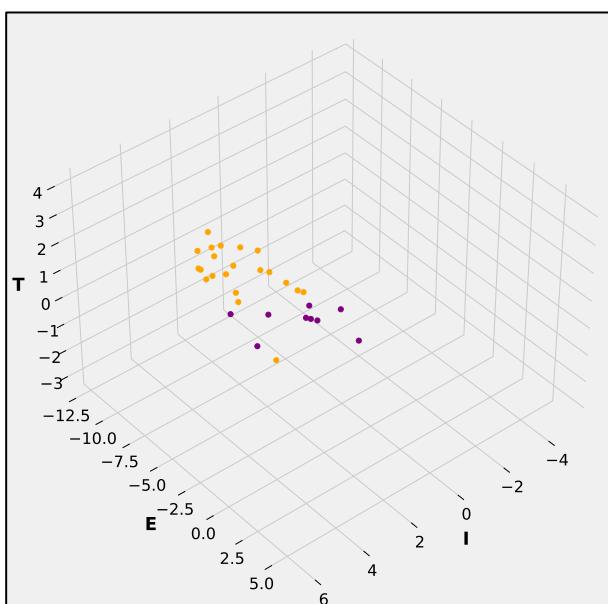


Of note, since we seek for directional cause-and-effect relationships among causal latent variables, we enforce the acyclicity of the causal graph through a penalization over term  $A_\gamma$  inspired by [50].

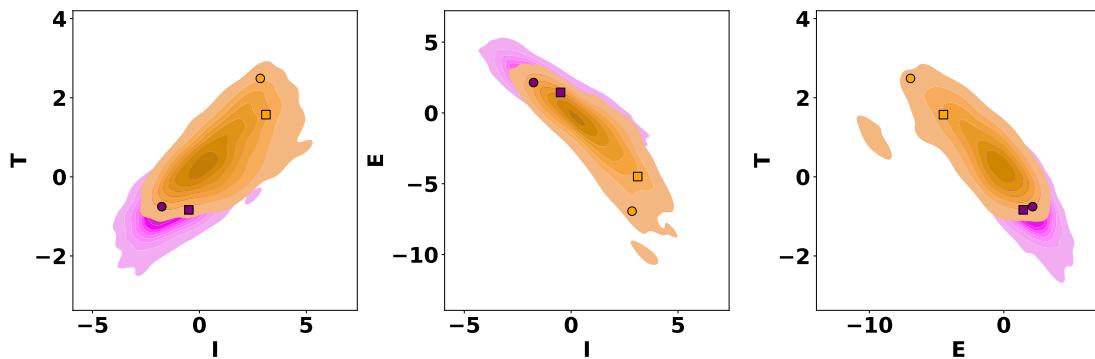
All encoders and decoders consist in this work of a single linear layer. MC<sup>2</sup>VAE training is efficiently carried out through stochastic gradient descent using the Adam optimizer. We perform a 5-fold cross-validation strategy and run for each experiment 1000 epochs with an initial learning rate of 1e-2 which allows us to reach convergence.

#### 4.2.3 Results and conclusions

In Figure 21, middle panel, we show the causal graph obtained by applying MC<sup>2</sup>VAE on our 3-channels dataset. Latent causal relationships are captured between the three channels, as reasonably expected. In particular, a causal relationship from the ionic currents towards electrophysiological biomarkers has been highlighted, which is in accordance with the fact that the blockade of ionic currents is used as input of the biophysical model from which these outputs are computed. We also noticed a bidirectional causal relationship between the electrophysiological biomarkers and the torsadogenic indices, suggesting a reciprocal influence between both set of biomarkers. We reported the mean values obtained for the  $A_\gamma$  weights over the folds, which quantify the strengths of each causal link. Of note, looking at the bidirectional relationship between the  $T$  and  $E$  channels one can note that  $T \rightarrow E$  has a stronger absolute weight: by the way, if one would decide to privilege a posteriori one direction based on the weight, removing  $E \rightarrow T$  would ensure the acyclicity of the final DAG.



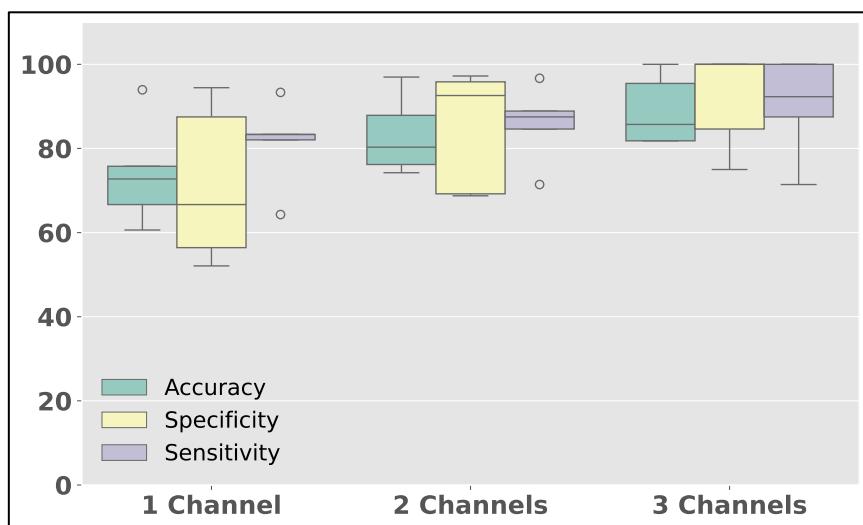
**Figure 22.** Scatter plot of all drugs after projection in the latent space. Unsafe drugs are in orange; safe drugs in purple.



**Figure 23.** Kernel density distribution of drugs for each pairwise combination of latent channels.

To quantify the impact of the discovered causal graph in defining a meaningful latent representation of the channels, we evaluate MC<sup>2</sup>VAE's ability to reconstruct the input data, measured through mean squared error (MSE) and assess that the causal learning step enables efficient reconstruction of the input data, providing very good reconstruction performances. Further, in Figure 22 we show the 3-dimensional latent representation of each drug obtained by sampling 50 times from their learned distribution, each axis being the latent coordinate for a specific channel. Despite MC<sup>2</sup>VAE is a fully unsupervised method, we can clearly see the separation of the drugs in the latent space with respect to their known TdP risk (considered here as binary: unsafe, in orange, for known or possible TdP risk drugs, and safe, in purple, for the remaining categories). For the sake of clarity, we further project these representations in each 2D plane (Figure 23), and show two unsafe and two safe drugs – ibutilide (orange circle), disopyramide (orange square), loratadine (purple circle), and diltiazem (purple square) – as an example.

Finally, we challenge our latent representation for the downstream binary classification task, and study the impact of an increasing number of channels to characterize drug-induced TdP risk. Figure 24 shows that the inclusion of each of the considered channels brings additional valuable information for TdP risk assessment, generating evidence and a strong rational for including them in the analysis.



**Figure 24.** Binary TdP-risk classification performance using a different number of latent channel-specific features.

In conclusion, thanks to MC<sup>2</sup>VAE we were able to reveals and quantifies the hidden causal relationships among the different sets of considered biomarkers, and helps to justify the integration of these three channels, as their joint causal analysis shows that they produce a much more effective and actionable characterization of TdP risk than if we were to consider one channel at a time.

## 5. Conclusions

Deliverable D5.5 presents a collection of works realized by the SimCardioTest consortium under Task 5.6 (M24-M54) on advanced data science for in-silico trials. The main results described here concern 1) the development of novel (unsupervised) statistical and machine learning-based approaches to leverage in-silico generated evidence (sometimes combined with other sources of information, in a multimodal framework) to gain better risk stratification, and 2) the development of novel pipelines to investigate the behavior of complex and highly non-linear models, needed to perform in-silico trials and account for high-dimensional observations.

**Model reduction and sensitivity analysis.** The increasing complexity of cardiac models needed for the integration and simulation of high-dimensional observations requires the development of *ad hoc* methods to statistically analyze and control their behavior from a large parameter space. Indeed, classical sensitivity analysis approaches may fail to provide a reliable solution, or require a prohibitive amount of model simulations to capture the complex interactions between model parameters and outputs: the development of novel sensitivity analysis approaches tailored to complex high-dimensional mechanistic models is still an active field of research. One direction pursued within the SimCardioTest consortium proposes to tackle this problem through model reduction (for PDE-based models), from 3D to 0D. Despite this strategy has shown effective by providing reliable information for 3D model calibration, several questions remain to be addressed, notably how to evaluate the credibility of this analysis (which is not part of V&V40)? How can we establish generic methodologies to guarantee the right features of the reduction processes and statistical analyses? How to properly characterize the relationship between the original and reduced model? An alternative direction explored within the context of SimCardioTest consists in leveraging the probabilistic theory of causality, and establishing causal relationships between parameters and outputs, for instance through the performance of regressions and conditional independence tests. The promising results obtained so far for single-point model results encourage us to move forward in this direction, where additional challenges need to be addressed, for example, how can we account for time-varying outcomes? Which is the most appropriate choice of causal discovery method?

**Digital twins with in-silico data.** The identification of risk factors of thrombus formation, and patient stratification into phenogroups associated with cardioembolic thrombus is essential for improving stroke treatment and prevention, thus supporting clinical decision-making. Nowadays, we dispose of mechanistic models based on computational fluid dynamics (CFD) which enable a reliable simulation of blood flows in patient-specific atrial geometries, from which in-silico hemodynamic indices can be generated. Within SimCardioTest we demonstrate the importance of integrating clinical, morphological, and hemodynamic data for a better understanding of the multifactorial drivers of thrombus formation in patients with advanced atrial fibrillation. To do that, we leverage

unsupervised machine learning (multiple kernel learning - MKL) to perform dimensionality reduction of the multidomain dataset, hence clustering in the lower-dimensional space: clusters were clinically interpreted, with a particular focus on the proportion of patients who had experienced thrombus formation. Future work should focus on the validation of the obtained phenogroups clusters against larger and external patient cohorts, and on the refinement of the proposed integrative framework, for instance by incorporating categorical variables into the analysis or by relying on more sophisticated mechanistic models of left atrial hemodynamics, to develop novel risk scores and predictive tools for thrombus formation.

**Causal learning in healthcare.** Causal learning is a very active and rapidly evolving area of research, which is now attracting growing interest in the healthcare community: it aims at exploiting directed causal relationships between observations, going beyond traditional statistical association, and ultimately improving interpretability, explainability and actionability. Causal discovery is a research direction in causal learning whose goal is to retrieve the causal structuring underlying the data generating process, which ultimately leads to the observations. Classical approaches for causal discovery are typically well suited to deal with relatively low-dimensional data, while more recent attempts which couple causal discovery and machine learning techniques to cope with higher dimensional and complex datasets (Causal Disentangled Representation Learning - CDRL) have shown promising results. Within SimCardioTest we propose to take a step further and investigate the possibility of extending CDRL to the multimodal scenario, thus taking advantage of multiple data sources to better characterize a given observed phenomena (drug-induced Torsade-de-Pointes). We show that our CDRL strategy, fully unsupervised, were effective in building an informative latent representation of the original multimodal data. Several challenges arise from this first study, including the possibility of modeling the causal relationships at multiple levels of granularity, and use the estimated latent causal graph to perform inference and prediction under varying conditions.

The results described in D5.5 have already been published and presented at international conferences:

- **Section 2.1:** Pannetier, Valentin, et al. "Towards validation of two computational models of artificial pacemakers" *Q. Z. Radomír Chabiniok, Functional Imaging and Modeling of the Heart*. Springer Nature Switzerland. Available July, 27th 2025. 2025. (ref. [9])
- **Section 2.2:** Al-Ali, Safaa, et al. "Cardiac Electromechanical Model Sensitivity Analysis using Causal Discovery", *Functional Imaging and Modeling of the Heart*. Springer Nature Switzerland. Available July, 27th 2025. 2025. (ref. [12])
- **Section 3:** Saiz-Vivó, Marta, et al. "Digital twin integrating clinical, morphological and hemodynamic data to identify stroke risk factors." *npj Digital Medicine* 8.1 (2025): 1-14. (ref. [26])
- **Section 4.1:** Al-Ali, Safaa, et al. "A causal discovery approach to streamline ionic currents selection to improve drug-induced tdp risk assessment." *2023 Computing in Cardiology (CinC)*. Vol. 50. IEEE, 2023. (ref. [42])
- **Section 4.2:** Al-Ali, Safaa, et al. "Assessing Ionic Current Blockades and Electromechanical Biomarkers' Interrelations Through a Novel Multi-Channel Causal Variational Autoencoder.". *2024 Computing in Cardiology (CinC)*. IEEE, 2024. (ref. [43])

## 6. Bibliography

- [1] Salvador, Matteo, et al. "Electromechanical modeling of human ventricles with ischemic cardiomyopathy: numerical simulations in sinus rhythm and under arrhythmia." *Computers in Biology and Medicine* 136 (2021): 104674.
- [2] Oomen, Pim JA, et al. "A rapid electromechanical model to predict reverse remodeling following cardiac resynchronization therapy." *Biomechanics and modeling in mechanobiology* (2022): 1-17.
- [3] Petras, Argyrios, et al. "Mechanolectric effects in healthy cardiac function and under left bundle branch block pathology." *Computers in Biology and Medicine* 156 (2023): 106696.
- [4] Rodríguez-Padilla, Jairo, et al. "Impact of intraventricular septal fiber orientation on cardiac electromechanical function." *American Journal of Physiology-Heart and Circulatory Physiology* 322.6 (2022): H936-H952.
- [5] Saltelli, Andrea, et al. *Sensitivity analysis in practice: a guide to assessing scientific models*. Vol. 1. New York: Wiley, 2004.
- [6] Saltelli, Andrea. "Sensitivity analysis for importance assessment." *Risk analysis* 22.3 (2002): 579-590.
- [7] Martins, Joaquim RRA, and Andrew Ning. *Engineering design optimization*. Cambridge University Press, 2021.
- [8] Pannetier, Valentin, et al. "Sensitivity Analysis in a Rescaled 0D Pacemaker-Bimembrane Model." *CINC 2024-Computing In Cardiology*. 2024.
- [9] Pannetier, Valentin, et al. "Towards validation of two computational models of artificial pacemakers" Q. Z. Radomír Chabiniok, *Functional Imaging and Modeling of the Heart*. Springer Nature Switzerland. Available July, 27th 2025. 2025.
- [10] Pannetier, Valentin, et al. "Modeling cardiac stimulation by a pacemaker, with accurate tissue-electrode interface." *International Conference on Functional Imaging and Modeling of the Heart (FIMH)*. Cham: Springer Nature Switzerland, 2023.
- [11] Desrues, Gaëtan, et al. "Personal-by-design: a 3D Electromechanical Model of the Heart Tailored for Personalisation." *International Conference on Functional Imaging and Modeling of the Heart*. Cham: Springer International Publishing, 2021.
- [12] Al-Ali, Safaa, et al. "Cardiac Electromechanical Model Sensitivity Analysis using Causal Discovery", *Functional Imaging and Modeling of the Heart*. Springer Nature Switzerland. Available July, 27th 2025. 2025.
- [13] Pianosi, Francesca, and Thorsten Wagener. "A simple and efficient method for global sensitivity analysis based on cumulative distribution functions." *Environmental Modelling & Software* 67 (2015): 1-11.
- [14] Sobol, Ilya M. "Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates." *Mathematics and computers in simulation* 55.1-3 (2001): 271-280.
- [15] Vowels, Matthew J., Necati Cihan Camgoz, and Richard Bowden. "D'ya like dags? a survey on structure learning and causal discovery." *ACM Computing Surveys* 55.4 (2022): 1-36.
- [16] Nogueira, Ana Rita, et al. "Methods and tools for causal discovery and causal inference." *Wiley interdisciplinary reviews: data mining and knowledge discovery* 12.2 (2022): e1449.
- [17] Saltelli, Andrea, et al. "Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index." *Computer physics communications* 181.2 (2010): 259-270.
- [18] Shimizu, Shohei, et al. "A linear non-Gaussian acyclic model for causal discovery." *Journal of Machine Learning Research* 7.10 (2006).
- [19] Shimizu, Shohei, et al. "DirectLiNGAM: A direct method for learning a linear non-Gaussian structural equation model." *Journal of Machine Learning Research-JMLR* 12.Apr (2011): 1225-1248.
- [20] Faure, François, et al. "Sofa: A multi-model framework for interactive physical simulation." *Soft tissue biomechanical modeling for computer assisted surgery* (2012): 283-321.
- [21] Aguado, Ainhoa M., et al. "In silico optimization of left atrial appendage occluder implantation using interactive and modeling tools." *Frontiers in physiology* 10 (2019): 237.

[22] Guerrero-Hurtado, Manuel, et al. "Hemodynamics Affects Factor XI/XII Anticoagulation Efficacy in Patient-Specific Left Atrial Models." *bioRxiv* (2024): 2024-08.

[23] Albors, Carlos, et al. "Impact of occluder device configurations in in-silico left atrial hemodynamics for the analysis of device-related thrombus." *PLOS Computational Biology* 20.9 (2024): e1011546.

[24] Quarteroni, Alfio, and Luca Formaggia. "Mathematical modelling and numerical simulation of the cardiovascular system." *Handbook of numerical analysis* 12 (2004): 3-127.

[25] Masci, Alessandro, et al. "The impact of left atrium appendage morphology on stroke risk assessment in atrial fibrillation: a computational fluid dynamics study." *Frontiers in physiology* 9 (2019): 1938.

[26] Saiz-Vivó, Marta, et al. "Digital twin integrating clinical, morphological and hemodynamic data to identify stroke risk factors." *npj Digital Medicine* 8.1 (2025): 1-14.

[27] Sanchez-Martinez, Sergio, et al. "Characterization of myocardial motion patterns by unsupervised multiple kernel learning." *Medical image analysis* 35 (2017): 70-82.

[28] Mill, Jordi, et al. "The role of the pulmonary veins on left atrial flow patterns and thrombus formation." *Scientific Reports* 14.1 (2024): 5860.

[29] Stekhoven, Daniel J., and Peter Bühlmann. "MissForest—non-parametric missing value imputation for mixed-type data." *Bioinformatics* 28.1 (2012): 112-118.

[30] Nerbonne, Jeanne M., and Robert S. Kass. "Molecular physiology of cardiac repolarization." *Physiological reviews* 85.4 (2005): 1205-1253.

[31] Witchel, Harry J. "Drug-induced hERG block and long QT syndrome." *Cardiovascular therapeutics* 29.4 (2011): 251-259.

[32] Guideline, ICH Harmonised Tripartite. "The non-clinical evaluation of the potential for delayed ventricular repolarization (Qt Interval Prolongation) by human pharmaceuticals." S7B (<http://www.ich.org/products/guidelines/safety/article/safety-guidelines.html>) (2005).

[33] Li, Zhihua, et al. "Assessment of an in silico mechanistic model for proarrhythmia risk prediction under the CiPA initiative." *Clinical Pharmacology & Therapeutics* 105.2 (2019): 466-475.

[34] Colatsky, Thomas, et al. "The comprehensive in vitro proarrhythmia assay (CiPA) initiative—update on progress." *Journal of pharmacological and toxicological methods* 81 (2016): 15-20.

[35] Mistry, Hitesh B. "Complex versus simple models: ion-channel cardiac toxicity prediction." *PeerJ* 6 (2018): e4352.

[36] Kramer, James, et al. "MICE models: superior to the HERG model in predicting Torsade de Pointes." *Scientific reports* 3.1 (2013): 2100.

[37] Zhou, Xin, et al. "Blinded in silico drug trial reveals the minimum set of ion channels for torsades de pointes risk assessment." *Frontiers in Pharmacology* 10 (2020): 1643.

[38] Parikh, Jaimit, et al. "Global sensitivity analysis of ventricular myocyte model-derived metrics for proarrhythmic risk assessment." *Frontiers in pharmacology* 10 (2019): 1054.

[39] Lancaster, M. Cummins, and E. A. Sobie. "Improved prediction of drug-induced Torsades de Pointes through simulations of dynamics and machine learning algorithms." *Clinical Pharmacology & Therapeutics* 100.4 (2016): 371-379.

[40] Passini, Elisa, et al. "Drug-induced shortening of the electromechanical window is an effective biomarker for in silico prediction of clinical risk of arrhythmias." *British journal of pharmacology* 176.19 (2019): 3819-3833.

[41] Llopis-Lorente, Jordi, Beatriz Trenor, and Javier Saiz. "Considering population variability of electrophysiological models improves the in silico assessment of drug-induced torsadogenic risk." *Computer Methods and Programs in Biomedicine* 221 (2022): 106934.

[42] Al-Ali, Safaa, et al. "A causal discovery approach to streamline ionic currents selection to improve drug-induced tdp risk assessment." *2023 Computing in Cardiology (CinC)*. Vol. 50. IEEE, 2023.

[43] Al-Ali, Safaa, et al. "Assessing Ionic Current Blockades and Electromechanical Biomarkers' Interrelations Through a Novel Multi-Channel Causal Variational Autoencoder.". *2024 Computing in Cardiology (CinC)*. IEEE, 2024.



- [44] Ruta, Dymitr, and Bogdan Gabrys. "Classifier selection for majority voting." *Information fusion* 6.1 (2005): 63-81.
- [45] Llopis-Lorente, Jordi, et al. "In silico classifiers for the assessment of drug proarrhythmicity." *Journal of Chemical Information and Modeling* 60.10 (2020): 5172-5187.
- [46] Kingma, Diederik P., and Max Welling. "Auto-encoding variational bayes." 20 Dec. 2013.
- [47] O'Hara, Thomas, et al. "Simulation of the undiseased human cardiac ventricular action potential: model formulation and experimental validation." *PLoS computational biology* 7.5 (2011): e1002061.
- [48] Antelmi, Luigi, et al. "Sparse multi-channel variational autoencoder for the joint analysis of heterogeneous data." *International Conference on Machine Learning*. PMLR, 2019.
- [49] Blei, David M., Alp Kucukelbir, and Jon D. McAuliffe. "Variational inference: A review for statisticians." *Journal of the American statistical Association* 112.518 (2017): 859-877.
- [50] Zheng, Xun, et al. "Dags with no tears: Continuous optimization for structure learning." *Advances in neural information processing systems* 31 (2018).



*This project received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101016496*