



MOX-Report No. 25/2021

**PDE-aware deep learning for inverse problems in
cardiac electrophysiology**

Tenderini, R.; Pagani, S.; Quarteroni, A.; Deparis S.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it

<http://mox.polimi.it>

PDE-aware deep learning for inverse problems in cardiac electrophysiology

Riccardo Tenderini^{1,2,*}, Stefano Pagani², Alfio Quarteroni^{2,3}, Simone Deparis¹

1 *SCI-SB-SD, Institute of Mathematics, École Polytechnique Fédérale de Lausanne, EPFL, Station 8, CH-1015 Lausanne, Switzerland*

2 *MOX, Modeling and Scientific Computing, Department of Mathematics, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 - Milano, Italy*

3 *Institute of Mathematics, École Polytechnique Fédérale de Lausanne, EPFL, Av. Piccard, CH-1015 Lausanne, Switzerland (Professor Emeritus)*

* Corresponding author. E-mail: `riccardo.tenderini@epfl.ch`

Abstract

In this work, we present a PDE-aware deep learning (DL) model, named Space-Time Reduced Basis Deep Neural Network (ST-RB-DNN), for the numerical solution to the Inverse Problem of Electrocardiography. The main feature of the proposed neural network (NN) is that it both leverages data availability and exploits the knowledge of a physically-based mathematical model, expressed by means of partial differential equations (PDEs), to carry out the task at hand. The goal is to estimate the epicardial potential field from measurements of the electric potential at a discrete set of points on the body surface. Such a problem has become central in biomedical research, providing the theoretical basis for Electrocardiographic Imaging (ECGI), but it is extremely hard to solve because of its ill-posedness. The employment of deep learning techniques in this context is made difficult by the low amount of clinical data at disposal (*small data regime*), as measuring cardiac potentials requires invasive procedures. Suitably exploiting the underlying physically-based mathematical model allowed to circumvent the data availability issue and led to the development of fast-training and low-complexity PDE-aware DL models. In particular, physical-awareness has been pursued by means of two elements: the projection of the epicardial potential onto a Space-Time Reduced subspace, spanned by the numerical solutions of the governing PDEs, and the inclusion of a tensorial Reduced Basis (RB) solver of the Forward Problem in the network architecture. Numerical tests have been conducted only on synthetic data, obtained via a Full Order Model (FOM) approximation of the problem at hand, and two variants of the model have been addressed. Both proved to be accurate, up to an average ℓ^1 -norm relative error on epicardial activation maps of $\approx 3.5\%$, and both could be trained in ≈ 10 min. Nevertheless, some improvements, mostly concerning data generation, are necessary in order to bridge the gap with clinical applications.

Keywords: Cardiac Electrophysiology, Deep Learning, ECGI, Inverse Problem of Electrocardiography, Partial Differential Equations, PDE-aware Deep Learning

1 Introduction

Over the last 30 years, especially in the last decade, computational biomedical research has witnessed the development of ECG-Imaging (ECGI) [9–11, 26, 33, 35, 40], a novel imaging modality for non-invasive mapping of cardiac electrical activity, which makes use of body-surface ECG signals and of thoracic CT-scans. The success of ECGI can be recognized in its ability of providing relevant diagnostic information on the electrical activity of the heart, visualizing it directly at the heart level and easing in turn the readability and interpretability of the results. Additionally, ECGI has proved to achieve an accuracy that plain body surface signals (as ECG or BSPM) do not offer and that could be instead obtained only by means of invasive measurement techniques [26]. Despite being a recent field of research, the stunning advances made over the last decade have already allowed to validate its performances in a number of clinical practices [9, 11] and to effectively employ it in the clinical setting [34].

From a mathematical standpoint, the theoretical basis of ECGI is the so-called Inverse Problem of Electrocardiography; it consists in finding the epicardial potential field which gives rise to body surface signals as close as possible to some target ones, in the least-squares sense. Let us consider a geometry made by the torso domain Ω_T , with the body surface Γ_B representing an external boundary and the epicardial one Γ_H representing an internal boundary. Additionally, let $\Sigma \subset \Gamma_B$ be the portion of the body surface where electric potentials are recorded and call z such recordings. Now, let $v \in H^{\frac{1}{2}}(\Gamma_H)$ be the epicardial potential and define $y(v)$ as the unique solution in $H^1(\Omega_T)$ of the following stationary elliptic problem

$$\begin{cases} -\nabla \cdot (D_T \nabla y(v)) = 0 & \text{in } \Omega_T \\ y(v) = v & \text{on } \Gamma_H \\ \nabla y(v) \cdot \mathbf{n}_B = 0 & \text{on } \Gamma_B \end{cases} \quad (1)$$

being D_T the tensor expressing the electrical conductivity in the torso and \mathbf{n}_B the outward unit normal vector to Γ_B . Also, we define the operator $A : H^{\frac{1}{2}}(\Gamma_H) \longrightarrow H^{\frac{1}{2}}(\Sigma)$ such that

$$Av = y(v)|_{\Sigma} \quad (2)$$

and the cost functional

$$J(v) = \frac{1}{2} \int_{\Sigma} |y(v) - z|^2 d\sigma = \frac{1}{2} \|Av - z\|_{L^2(\Sigma)}^2 \quad (3)$$

for $z \in L^2(\Sigma)$. Then, the Inverse Problem of Electrocardiography can be written as:

$$\text{find } u_H \in H^{\frac{1}{2}}(\Gamma_H) : \quad J(u_H) = \min_{v \in H^{\frac{1}{2}}(\Gamma_H)} (J(v) + R(v)) \quad (4)$$

being $R(v)$ a suitable regularization term. Without stabilization (i.e. setting $R(v) = 0$) problem (4) is ill-posed in usual Sobolev spaces, which means that A admits an unbounded inverse operator in the spaces H^s , $\forall s \in \mathbb{R}$ (see [36, 46, 49]). This implies that small

perturbations (for instance due to measurement errors) in the observed surface potentials z may lead to much larger variations on the reconstructed epicardial potential u_H . To retrieve well-posedness, problem (4) must be therefore approximated with a family of stable problems. The regularization term $R(\cdot)$ can be computed as

$$R(v) = \frac{\alpha}{2} \|v\|_{H^{1/2}(\Gamma_H)}^2 \quad \text{or} \quad R(v) = \frac{\alpha}{2} \|y(v)\|_{H^1(\Omega_T)}^2 \quad (5)$$

being $y(v)$ the solution of system (1) with v as Dirichlet boundary datum on Γ_H and α a hyperparameter to be tuned. Two other well-established techniques are also *Tikhonov Regularization* [47], which consists in computing $R(\cdot)$ as the $L^2(\Gamma_H)$ -norm either of the epicardial potential or of its derivatives, and *Total Variation Regularization*, which instead defines $R(\cdot)$ as the $L^1(\Gamma_H)$ -norm of the normal derivative of the epicardial potential. Notice that all these strategies are subject to the choice of at least one hyperparameter; various techniques have been proposed to (sub-)optimize such a choice [29].

In the current work we have followed a different path, bypassing classical approaches thanks to two strategies. The first one is the exploitation of Machine Learning (ML) and, more specifically, of Deep Learning (DL), which has established itself as a pillar of a new generation of scientific development, thanks to the abundance of available data and to the progresses in terms of computational power and resources. In particular, we wish to overcome the ill-posedness of the Inverse Problem of Electrocardiography leveraging data availability, yet still somehow retaining knowledge of the physical laws underlying the problem at hand. Incidentally, if until few years ago ECGI could be regarded as almost free of ML-DL influence, recently few works pursuing an integration between the fields did start to appear. Other than neural network (NN) models yielding heartbeats classification from body surface signals [28,55], we ought to mention the works by *Giffard et al.* [20], *Wang et al.* [50], *Bacoyannis et al.* [2] and *Bujnarowski et al.* [5]. In particular, [20] employs kernel ridge regression and a 2-term exponential regression to perform a patient-specific estimation of the main parameters of the chosen cardiac model. [50] builds a NN that provides (sub-)optimal values for the regularization parameters involved in the Alternating Direction Method of Multipliers iterative optimization algorithm. [2] proposes a novel method based on Conditional β Variational Autoencoders (β -CVA) using Deep Convolutional Generative Neural Networks for data-driven integration of spatio-temporal correlations and imaging information in the ECGI problem. [5] proposes a Convolutional Deep Autoencoder which yields an imaging-free approximation of the mapping between recorded body surface potentials and epicardial ones.

The second key element is based on the observation that all classical methods to solve the Inverse Problem of Electrocardiography involve the computation of the transfer matrix **A**. Such matrix depends on the geometry of the torso and of the heart and on their electrical conductivities and it can thus be inferred only through thoracic CT-scans and subsequent imaging post-processing routines. Anyway, CT-scans imply a certain level of radiation and they are not even an available option for all the patients. Developing, then, a

CT-scan-free model, able to map body surface potentials to epicardial ones in a reliable and physically-consistent way, could be of great help and “could serve as a preliminary study of the patient’s condition, before a more thorough examination is performed” [5]. This work was developed in a *fixed geometrical setting*; with this expression, we mean that we employed the same geometry in all the numerical simulations and we kept it fixed over time, discarding variations in the shape, dimension, position and rotation of the heart or in the dimension and shape of the torso. Furthermore, we approximated the latter as a homogeneous and isotropic volume conductor, neglecting the presence of other organs than the heart, characterized by different electric conductivities. Despite being quite strong assumptions [12, 41, 54], they can be all justified considering this work as an initial proof-of-concept on the integration between DL and ECGI. Our ultimate goal, indeed, is to work in a dynamic and parametrized geometrical setting, designing a model which does not just reconstruct the epicardial potential field, but that also provides estimates of suitable parameters characterizing the geometry and the electric conductivity.

Despite [5] appears to have many points in common with this work, here a further key element is present: physical awareness. Indeed the model we propose has been designed so that, albeit in the framework of a classical Encoder-Decoder structure, some elements allow to drive the predicted solutions to be consistent with the problem physics, approximated by means of PDEs (from which the expression “PDE-aware Deep Learning”). Broadening the scope, the proposed model is an instance of a physically-informed deep learning model, i.e. a DL model which exploits the knowledge of some physical laws underlying the phenomenon of interest to ease its training and/or to ameliorate its predictive power. Recent years have witnessed a flourishing of works in this direction [7, 15, 24, 31, 39, 51]; among those a special mention is deserved by *Raissi et al.* [39], where the authors introduce Physically-Informed Neural Networks (PINNs), i.e. neural networks able to approximate solutions to a variety of PDEs by minimizing, in the loss function, the Mean Squared Error (MSE) on the residuals. Another work we ought to refer to is [15] by *Dal Santo et al.*; there the authors present DL models (called RB-DNNs) that, receiving in input samples of solutions to parametrized PDEs and embedding a Reduced Basis (RB) solver [38] as a deterministic layer inside the NN architecture, manage to reconstruct such solutions in the whole computational domain and further provide estimates of the characteristic parameters.

According to [39], physics-informed DL can perform well in the so-called *small-data* regime, i.e. in frameworks where the amount of data at disposal is either limited or partial or subject to a high degree of inaccuracy, because of the high cost and/or complexity of data acquisition procedures. Indeed, in such a context, classical DL methods feature severe problems in terms of robustness, generalization and convergence, due to the lack of data. However, if the phenomenon generating such data happens to be characterized by the presence of some underlying physical laws, expressible by means of PDEs/ODEs, then classical numerical methods can be integrated with DL-based ones. In particular the knowledge of the physics of the problem at hand can be made somehow available to

the DL model, acting as a physically-aware regularization agent, that (hopefully) eases model training and improves its performances. ECGI is clearly a context of such kind; indeed on the one side recordings of the epicardial potential are nowadays possible only via intrusive techniques and assembling a dataset made of a number of observations adequate for DL applications is unfeasible. On the other side, instead, a good understanding of the most important physical laws underlying cardiac electrophysiology is present [19, 45], so that exploiting it in the framework of physically-informed DL appears both doable and challenging.

Outline

Section *Materials and Methods* presents a physics-aware deep learning model, called Space-Time Reduced Basis Deep Neural Network (ST-RB-DNN), that provides reliable and physically-consistent solutions to the Inverse Problem of Electrocardiography and that furthermore is CT-scans-free. In particular: in Subsection *General model description* an overview on the ST-RB-DNN model is given, in order to highlight its structure and its main features. Subsection *Data generation and model order reduction* reports a brief description of the numerical modelling of cardiac electrophysiology and of the Forward Problem of Electrocardiography and an overview on the model order reduction techniques exploited in the ST-RB-DNN model; additional details on these topics can be found in S1 Appendix. In Subsection *Additional details and information* a more extensive and technical characterization of the proposed model is reported and two different versions of it are sketched. Section *Results* reports the numerical tests conducted with two different versions of the ST-RB-DNN model on the same test case; a brief description of the computational setting used to generate the dataset (geometries, meshes, training-test splitting, sources of randomness, etc.) is also given. Section *Discussion* evaluates the overall performances of the model and investigates the differences (in terms of accuracy and complexity) between the proposed versions of it; additionally, Subsection *Limitations* lists the main limitations and proposes some possible further developments. Section *Conclusion* provides a final summary.

2 Materials and methods

2.1 General model description

The proposed PDE-aware DL model for the Inverse Problem of Electrocardiography is named Space-Time Reduced Basis Deep Neural Network (ST-RB-DNN); its general structure is reported in the scheme of Fig 1.

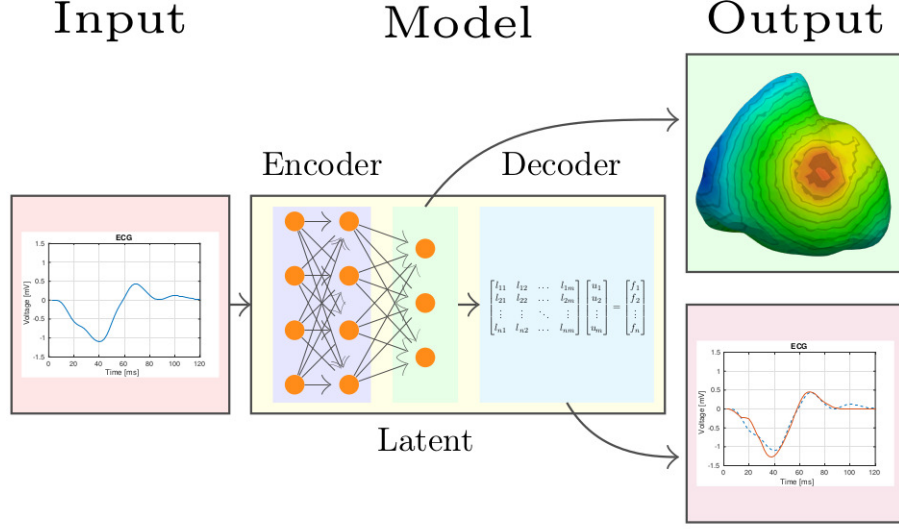


Fig 1. Basic structure of the ST-RB-DNN model

Three main components can be identified:

- **Input:** the input is constructed from body surface signals, recorded by electrodes placed on the patient's skin. The raw signals could possibly undergo a suitable pre-processing (lowpass/bandpass filtering, magnitude normalization, Discrete Fourier Transform (DFT), etc.). Also, all datapoints must contain the same number of signals, but it is possible to construct different datasets and to train different models, considering a different number of signals.
- **Model:** the model can be identified as an autoencoder, similar to the one presented in [5]. However, while the encoder consists of a neural network trained with the classical backpropagation algorithm [42], the decoder is deterministic, being an embedded tensorial Reduced Basis solver for Problem (1). In particular:
 - **Encoder:** the encoder is a deep neural network, which takes as input the body surface signals (or some quantities derived from them after some pre-processing routine) and learns a latent dimensionality-reduced representation of the epicardial potential field, responsible for the generation of the same signals given in input. To some extent, it can be stated that the encoder provides a data-driven approximation of $A^{-1} : L^2(\Sigma) \rightarrow L^2(\Gamma_H)$, i.e. the inverse of the transfer operator A defined in Eq (2). Two aspects are worth considering: the first one is that the architecture of the NN depends on the nature of the quantities that are provided in input and, furthermore, it can be optimized via a grid-search process, aimed at finding a (sub-)optimal balance between accuracy and complexity. The second is that the latent representation of the epicardial potential field is not uninterpretable and inexplicable, as the one proposed in [5]. Indeed, what the

Encoder learns are the coefficients arising from the projection of the epicardial potential field onto a spatio-temporal dimensionality-reduced subspace, generated from the same physical solutions included as output in the training dataset. The way such a subspace is generated will be sketched in Subsection *Data generation and model order reduction* and it is better detailed in S1 Appendix. A crucial remark is that body surface signals can be deterministically reconstructed from the estimated latent representation of the epicardial potential field. Moreover, such a representation features optimality properties that seamlessly depend on the method employed to generate the aforementioned dimensionality-reduced subspace. This entails that the epicardial potential field is encoded in the lowest possible number of coefficients, that guarantee a certain accuracy level in a suitable norm.

- **Decoder:** the decoder is responsible for the deterministic reconstruction of the body surface potentials, given the latent representation of the epicardial potential field estimated by the encoder. As such latent representation is actually a projection onto an optimal spatio-temporal dimensionality-reduced subspace, the decoder can be constructed as a tensorial (to be compliant with the backpropagation training algorithm) Reduced Order Model (ROM) solver of the Forward Problem of Electrocardiography (FPE) [3]. In particular, our model features a decoder that employs the Reduced Basis method to solve the FPE independently at all the discrete time instants. Additionally, the obtained solutions are suitably post-processed in the full order spatial domain, so that signals matching the ones provided in input can be computed. Such a choice presents a clear advantage with respect to the usage of a trainable decoder as in [5]; indeed no additional trainable parameters are added to the ones of the encoder and the overall model complexity is potentially halved. As a cons, anyway, designing the decoder as a ROM solver of the FPE makes it more difficult to achieve independence from imaging and CT-scans. In this paper, we have adopted a *fixed geometrical setting*, as described in Section *Introduction*; a natural future development would be to parametrize such setting and, consequently, to adapt and update the embedded RB solver.
- **Output:** the model output is made of two parts: the latent dimensionality-reduced representation of the epicardial potential field and the body surface signals reconstructed by the embedded RB solver of the FPE and matching the ones given in input. Consequently, the loss functional is constructed as a weighted sum between the estimation error on the dimensionality-reduced epicardial potential fields \mathcal{L}_{BC} and the reconstruction error on the body surface signals \mathcal{L}_{sig} , so that:

$$\mathcal{L}(\Theta) = \mathcal{L}_{BC}(\Theta) + \mathcal{L}_{sig}(\Theta) + \mathcal{L}_{reg}(\Theta) \quad (6)$$

Here Θ is the vector of the neural network trainable parameters and the additional term \mathcal{L}_{reg} represents a regularization term, which helps in preventing data overfitting.

The precise expressions of the three terms will be given in Subsection *Additional details and information*. We underline that, in principle, the ST-RB-DNN model can be trained as a pure autoencoder, by suppressing the loss term \mathcal{L}_{BC} related to the latent representation of the epicardial potential fields. Doing so, the non-invasive collection of body surface potentials would be sufficient for model training and all the invasively recorded cardiac potentials could be employed just for the sake of clinical validation. Anyway, preliminary numerical tests, conducted employing a numerically generated dataset, have shown that including a penalization of the epicardial potential fields in the loss dramatically improves model performance. As we employed only 400 datapoints (despite being them augmented on-the-fly via the superimposition of correlated white Gaussian noise - see Section *Results*), we expect to get better results by training the ST-RB-DNN model as an autoencoder on a bigger dataset and further investigations in this direction are planned. The implementation of additional regularization terms - as the ones reported in Eq (5) - to mitigate the ill-posedness of the underlying problem is also going to be analysed. For the moment, anyway, we decided to include a penalization of the epicardial potential fields in the loss and, consequently, to employ a numerically generated dataset both for training and testing. In view of this, the RB solver of the FPE should be interpreted, more than as a decoder, as a physically-aware regularization agent, which drives the predicted epicardial potential fields to belong to a lower dimensional and physically-consistent manifold.

Ultimately, two elements contribute to that physical awareness that makes the ST-RB-DNN model a PDE-aware DL model for ECGI:

1. The estimation of the epicardial potential field as projected onto a spatio-temporal dimensionality-reduced subspace, generated from the same solutions to the cardiac electrophysiology problem that represent the output part of the training dataset. In this way, indeed, the estimated epicardial potential field is forced to belong to a lower-dimensional manifold of physically-consistent solutions. To meet this requirement the training dataset should be made of realistic (if not real) data, suitable to capture the majority of the dynamics of the problem of interest.
2. The reconstruction and penalization of body surface signals equivalent to the ones provided in input, which are computed via an embedded RB solver of the FPE that acts as a deterministic decoder. Indeed, if the contribution of such reconstruction error in the loss is sufficiently high to be non-negligible, it drives the model to estimate epicardial potential fields which, at least in the considered *fixed geometrical setting*, determine the onset of body surface signals close to the observed ones. So, the space of admissible solutions for the Inverse Problem is further shrunked, as additional physical constraints have been added, and it ultimately reduces to a low-dimensional and physically-consistent manifold.

2.2 Data generation and model order reduction

Due to the paucity of clinical data relative to measurements of the electric potential at the epicardium, the datasets we employed for both training and testing the ST-RB-DNN model have been generated *in silico*. This involved the modelization and the numerical approximation of both cardiac electrophysiology and the FPE. All simulations have been carried out in a *fixed geometrical setting* (see Section *Results*) and the different datapoints have been generated by randomly changing the values of the main parameters characterizing the chosen cardiac model.

Cardiac electrophysiology has been modelled via the bidomain equations [48], coupled with the Aliev-Panfilov (AP) ionic model [1]. The latter is a phenomenological model that allows to approximate the current induced by the movement of electrically charged ions through the cell membrane. The former, instead, are characterized by a partitioning of the heart domain in an intracellular domain and an extracellular one, which however coexist on the same space upon a homogenization process [21]. The anisotropy of the cardiac tissue, dictated by the presence of conducting fibers, has been taken into account by setting different electric conductivities longitudinally and transversally to the fibers sheet [37].

In modelling the FPE, we have approximated the human torso as an isotropic and inhomogeneous volume conductor. The inhomogeneity derives from the presence of different organs (bones, lungs, blood vessels, etc.), characterized by different electric conductivities. To reduce the computational cost of the FOM simulations, we have decided to work under the *isolated heart assumption* (see [4]); in this way, solving the FPE amounts at computing the solution to a generalized Laplace equation, with homogeneous Neumann boundary conditions on the body surface and inhomogeneous Dirichlet ones at the epicardium.

In the ST-RB-DNN model, Reduced Order Model (ROM) techniques have been used both to estimate physically meaningful and dimensionality-reduced latent representations of the epicardial potential fields and to efficiently solve the FPE, in order to reconstruct signals matching the ones given in input. Considering that the FPE is a stationary problem, the latter task can be carried out by resorting to the classical RB method [38], i.e. by deriving a basis for a dimensionality-reduced subspace in space (via Proper Orthogonal Decomposition (POD) of the snapshots' tensor), performing a projection along the spatial dimension and solving the resulting reduced problem independently at all the discrete time instants. The former task, instead, can be much better accomplished by resorting to spatio-temporal ROM techniques, as they allow to encode the information coming from time-dependent fields into a very low number of coefficients, (almost) independent of the level of refinement of the grid along both the spatial and the temporal dimension. In particular, the space-time-reduced subspace where the epicardial potential fields have been projected has been generated using the *Tailored Temporal Subspaces via ST-HOSVD* approach proposed in [8]. This strategy is based on the computation of a tailored temporal basis for each element of the spatial one,

allowing for a dramatic reduction of the dimensionality of the resulting spatio-temporal subspace. We notice that estimating a FOM approximation of the epicardial potential field is not a viable option in the framework of the ST-RB-DNN model as that would significantly increase the overall complexity, severely hindering the training process. Reference [30] features a theoretical analysis on this aspect, that can be also adapted to our case.

2.3 Additional details and information

This subsection is devoted to a more detailed presentation of the ST-RB-DNN model architecture. In particular, the first part provides a description of the generic model structure (i.e. the general structure of the ST-RB-DNN model, independently of the nature of the signals given in input). The second and the third part, instead, focus on the two models that have been actually implemented and tested in this work, analysing their specific features with respect to the general case.

The generic ST-RB-DNN model

The architecture of the generic ST-RB-DNN model is reported in Fig 2. The model takes as input body surface signals or quantities directly derived from those (in red). The initial processing is done by a trainable NN (in blue) whose actual architecture depends on the nature of the quantities given in input and which acts as an encoder. Such NN is responsible for the estimation of two type of quantities:

1. **A latent representation of the epicardial potential field** (in orange): the main element which is estimated by the encoder is a latent representation of the epicardial potential field, which is obtained by means of a projection onto a space-time reduced subspace generated from physical solutions to the heart electrophysiology problem (see S1 Appendix). Notice that the latent representation estimated by the encoder, called $\bar{\mathbf{u}}_e^{a,\mu}$, is not the actual projection to be lately provided in input to the embedded RB solver of the FPE (in cyan), called $\hat{\mathbf{u}}_e^{a,\mu}$. Indeed, the former is processed by a deterministic *rescaler layer* (in yellow) to obtain the latter; the rescaling is affine and it has the following form:

$$\hat{\mathbf{u}}_e^{a,\mu} = \hat{\mathbf{u}}_e^{shift} + \hat{\mathbf{u}}_e^{scale} \bar{\mathbf{u}}_e^{a,\mu} \quad (7)$$

with $\hat{\mathbf{u}}_e^{shift}$ and $\hat{\mathbf{u}}_e^{scale}$ representing the additive and the multiplicative terms of the affine transformation. In simpler terms, the NN estimates how far is (in terms of the epicardial potential) the current datapoint from a target value, referred to as the "shift", with a scaling factor given by what is referred to as the "scale". Several choices can be made for the values of these two quantities; we decided to set $\hat{\mathbf{u}}_e^{shift}$ to $\mathbf{0}$ (so, corresponding to a totally inactivated heart) and to compute $\hat{\mathbf{u}}_e^{scale}$ as the standard deviation of the epicardial potentials in the training dataset, assuming them to have $\mathbf{0}$ mean. The rescaling is motivated by the fact that the values of the coefficients

arising from the projection onto the space-time-reduced subspace happen to span over a broad range of orders of magnitude and this severely hinders the training process. Exploiting Eq (7) significantly improves the performances.

2. **Torso conductivities** (in green): other than the epicardial potential field, the encoder may also learn the values of the electric conductivities of the different organs considered in the torso geometry, in case an inhomogeneous approximation of the latter is performed. To be precise, following the suggestions given in [15], such parameters are estimated as normalized in the interval $[0; 1]$ by employing a *sigmoid* as activation function. Additionally, working under the *isolated heart assumption*, what matters are the relative values of the conductivities and not their absolute ones; thus, one conductivity must be always normalized to 1 by default and so no parameters' estimation is needed if a homogeneous torso approximation is employed.

Finally, the quantities estimated by the encoder are fed to an embedded tensorial RB solver of the FPE (in cyan), which acts as an efficient and deterministic decoder. In particular:

- Given the estimated torso conductivities, it rescales them in the proper ranges via a *Min-Max rescaler* and it assembles the torso stiffness matrix exploiting affine parametrization (see S1 Appendix for additional details).
- It expands the estimated space-time-reduced epicardial potential along the temporal dimension only.
- It solves the FPE, reduced along the spatial dimension, at all time instants.
- It reconstructs the FOM approximation of the torso potential at the electrodes locations and it computes signals equivalent to the ones provided in input to the model. Those are then returned in output (in purple).

As reported in Eq (6), the loss functional is constructed as the sum of three contributions: the error on the dimensionality-reduced epicardial potential field, the error on the reconstruction of body surface signals and a *Ridge Regularization* term, which helps in preventing *overfitting*. More in detail, the loss functional is defined as follows:

$$\begin{aligned} \mathcal{L}(\Theta) &= \mathcal{L}_{BC}(\Theta) + \mathcal{L}_{sig}(\Theta) + \mathcal{L}_{reg}(\Theta) \\ &= w_{BC} \text{MAE}_{\sigma}(\hat{\mathbf{u}}_e^r, \hat{\mathbf{u}}_e^a(\Theta)) + w_{sig} \sum_{i=1}^{n_{sig}} \text{MSE}(\mathbf{S}_{\#i}^r, \mathbf{S}_{\#i}^a(\Theta)) + \lambda_r \|\Theta\|_2^2. \end{aligned} \quad (8)$$

The quantities with apex r denote target values, the quantities with apex a denote approximated values, $\mathbf{S}_{\#i}$ encodes the i^{th} input signal, from a total of n_{sig} signals, and w_{sig} , $w_{BC} \in \mathbb{R}$ represent the weights of the body surface signals and of the space-time-reduced epicardial potential field respectively. Incidentally, since what matters is the ratio between

w_{BC} and w_{sig} , the latter has been always set equal to 1. The choice of using the Mean Absolute Error (MAE) instead of the Mean Squared Error (MSE) for the loss term \mathcal{L}_{BC} aims at forcing to 0 the least relevant coefficients, as characteristic of ℓ^1 -norm penalization. Additionally, the MAE appearing in \mathcal{L}_{BC} is not only weighted by the non-negative scalar w_{BC} , but also by the singular values that, during the POD-construction of the two Reduced Bases, have been associated to the different coefficients. So,

$$MAE_{\sigma}(\hat{\mathbf{u}}_e^r, \hat{\mathbf{u}}_e^a(\Theta)) =: \frac{1}{b_s n_{st}} \sum_{i=1}^{b_s} \sum_{j=1}^{n_s^e} \sum_{k=1}^{n_t^j} \sqrt{\frac{\sigma_j^{e,s} \sigma_k^{e,t,j}}{\sigma_1^{e,s} \sigma_1^{e,t,1}}} |(\hat{\mathbf{u}}_e^r - \hat{\mathbf{u}}_e^a(\Theta))_{\mathcal{F}(j,k)}| \quad (9)$$

where $\{\sigma_j^{e,s}\}_{j=1}^{n_s^e}$ are the singular values arising from the (spatial) POD applied to the mode-1 unfolding of the epicardial potentials tensor, $\{\sigma_k^{e,t,j}\}_{k=1}^{n_t^j}$ are the singular values deriving from the (temporal) POD applied to the projection of the epicardial potentials tensor onto the 1-dimensional subspace spanned by the j -th spatial basis element, b_s is the training batch size and $\mathcal{F}(\cdot, \cdot)$ is the mapping from the Space and Time bases indexes to the Space-Time basis index. A more detailed and rigorous definition of all such quantities can be found in S1 Appendix. In this way, the weight of the error on a space-time-reduced epicardial potential coefficient in the loss gets lower and lower as its relevance decreases; this allows to obtain better estimates of the most relevant coefficients, easing the training process and improving model performances. Two remarks are worth to follow. Firstly, by resorting to ROM techniques, we managed to mitigate the effects of the ill-posedness of the inverse problem, as the dimensionality of the space where we seek for a solution is dramatically reduced. However, the risk of facing instability issues grows with the number of considered spatio-temporal basis functions (i.e. it increases as the POD tolerances decrease). Due to the small dimension of our dataset, we did not experience any problem in this sense; nevertheless, adding supplementary regularization terms - as the ones of Eq (5) - to the loss functional certainly allows to retain well-posedness. Secondly, the estimation error on the torso conductivities is not involved in the expression of the loss functional; as in [15], such values are obtained as a by-product.

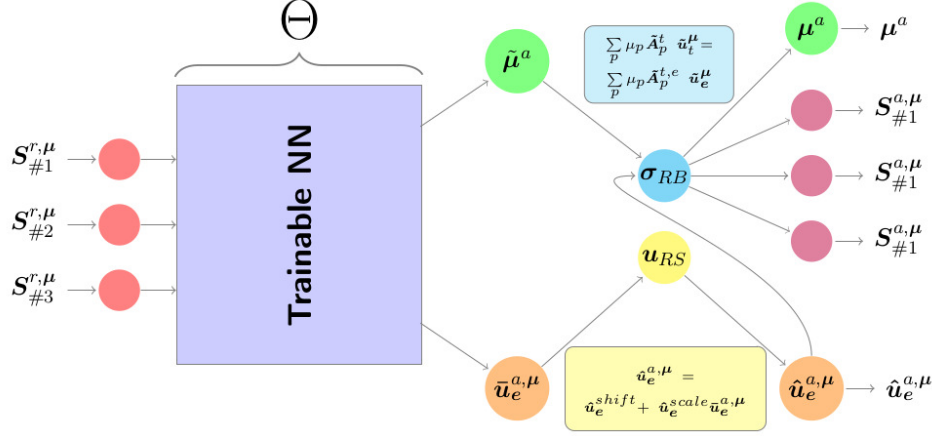


Fig 2. Scheme of the generic architecture of the ST-RB-DNN model

The time-series-based ST-RB-DNN model

The architecture of the time-series-based ST-RB-DNN model is shown in Fig 3. Its main feature, with respect to the generic model, is that it takes as input body surface signals organized in the form of time-series. The following aspects deserve to be discussed:

- As suggested in [28, 55], at pre-processing stage it is useful to remove noise components from body surface signals, since it eases the training process and it increases the representative power of the models. Many works regarding the denoising of real body surface potentials are available, as [43, 44]. Anyway in this project we considered only numerically generated signals with superimposed correlated white Gaussian noise at an average SNR of 26 dB; thus, we employed a standard low-pass Butterworth filter of order $n = 3$ [6]. Additionally, all signals have been normalized in the interval $[-1; 1]$ to ensure they are all within the same scale, as in [5].
- Since the input signals are provided as time-series, we chose to employ a Convolutional Neural Network (CNN) design for the encoder. In particular, the structure of the CNN can be divided in three parts:
 - *Pre-Convolutional Fully-Connected Layers*: as first step, the input signals are processed by some fully-connected layers; their aim is to extract features on top of which convolutional operations may exhibit better performances. It is relevant to notice that these layers only combine quantities related to the same time instant, using the same set of weights at all time instants. In other terms, the original input signals are non-linearly combined together, so that other signals, featuring a better information encoding, can be derived. The activation function is *ReLU* for all neurons.
 - *1D Convolutional + Max-Pooling Layers*: after the initial "pre-processing" stage, the resulting data are passed by to 1D convolutional layers, that convolve them

along the temporal dimension. The number of layers, the dimensionality of the convolutional kernel and the number of convolutional filters are three key hyperparameters of the model. Also, each 1D-convolutional layer is followed by a *Max-Pooling* layer, with pooling window of dimension 2×3 and with 2×2 stride. As a result, with proper padding choices, the dimensionality of the input is halved along both dimensions at each step, reducing in turn the number of trainable parameters of the model.

- *Post-Convolutional Fully-Connected Layers*: after 1D convolutions, the data are finally processed by a second set of fully-connected layers. Their aim is to extract relevant features to ultimately perform a good estimation of both the torso conductivities and the space-time-reduced epicardial potential. All neurons of all layers are *ReLU*-activated, except from the ones of the last layer, i.e. the one providing the final estimates. These ones indeed feature two different activation functions: *sigmoid* for the neurons storing the normalized values of the torso conductivities and *SELU* (i.e. an optimally scaled version of the Exponential Linear Unit (ELU)) for the ones storing the space-time-reduced epicardial potential. This last choice allows to circumvent the so-called "dying *ReLU*" problem [32], which has been experienced in some preliminary tests employing *ReLU* as activation function.

Additionally, the model is equipped with *dropout* layers with $p = 0.20$ to reduce the risk of *overfitting*.

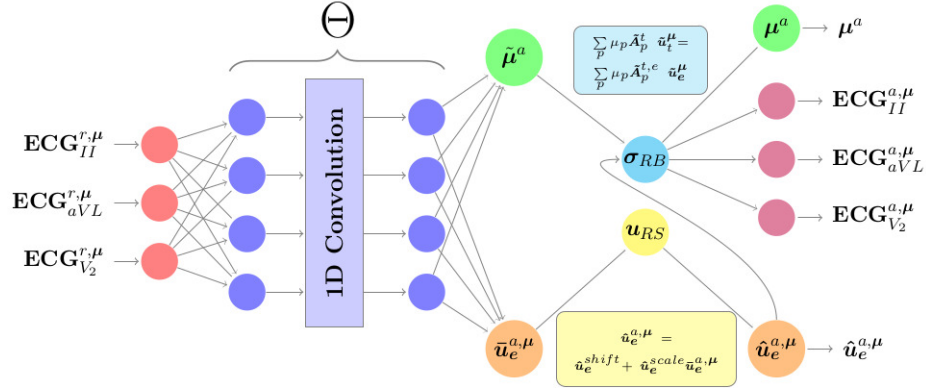


Fig 3. Scheme of the architecture of the Time-series-based ST-RB-DNN model. Notice that, as the input body surface signals are in the form of time series, the Trainable NN acting as decoder is a CNN, with 1D convolutional layers acting along the temporal dimension only.

The DFT-based ST-RB-DNN model

The architecture of the DFT-based ST-RB-DNN model can be visualized in Fig 4. The name stems from the usage of low-frequency Discrete Fourier Transform (DFT) coefficients

of the body surface signals to assemble the model input. The 1D DFT is the discrete counterpart of the uni-variate Fourier Transform and it turns a sequence of N complex numbers $\{x_n\}_{n=0}^{N-1}$, which is said to belong to the temporal domain, into another sequence of N complex numbers $\{X_k\}_{k=0}^{N-1}$, which belongs instead to the frequency domain. In particular:

$$X_k = \sum_{n=0}^{N-1} x_n e^{-\frac{2\pi i}{N}kn} = \sum_{n=0}^{N-1} x_n \left[\cos\left(\frac{2\pi}{N}kn\right) - i \sin\left(\frac{2\pi}{N}kn\right) \right] \quad (10)$$

As Eq (10) is hermitian on real inputs, the dimensionality of the DFT applied to body surface signals is $\bar{M} =: \lfloor N_t/2 \rfloor + 1$, being N_t the signals' length. The dimensionality of the model inputs is then $M \leq \bar{M}$, upon having excluded a certain portion of coefficients, related to the highest-frequency modes; the choice of M is a model hyperparameter. The DFT of body surface signals have been efficiently computed using the Fast Fourier Transform (FFT) algorithm [14]. The following aspects deserve to be underlined:

- A drawback of the time-series-based ST-RB-DNN model is that it can work only at a fixed acquisition frequency; in our test cases, for instance, we sampled body surface signals at 500 Hz . Thus, if the signals to be processed have been sampled at a lower frequency, they have to be interpolated in time before being provided to the model. Conversely, the lowest-frequency DFT coefficients feature a low sensitivity with respect to the acquisition frequency; thus, using them as input allows to avoid any interpolation at pre-processing stage. Notice, as a *caveat*, that the low sensitivity occurs only with respect to the lowest-frequency DFT coefficients; however highest-frequency ones are typically not much informative, so that model accuracy should not be too much affected by their removal.
- Since the application of a DFT at pre-processing stage already allows to capture temporal dynamics, a convolution inside the encoder is no longer needed. Thus, the latter simply consists of a Multiple Layer Perceptron (MLP), with several fully-connected layers flanked one after the other. Also, the encoder is equipped with *Dropout* layers with $p = 0.20$ to reduce the risk of *overfitting* and all its layers are *ReLU*-activated, except from the last one. Its neurons, indeed, feature either *sigmoid* or *SELU* as activation function, depending on the nature of the quantity they are estimating.

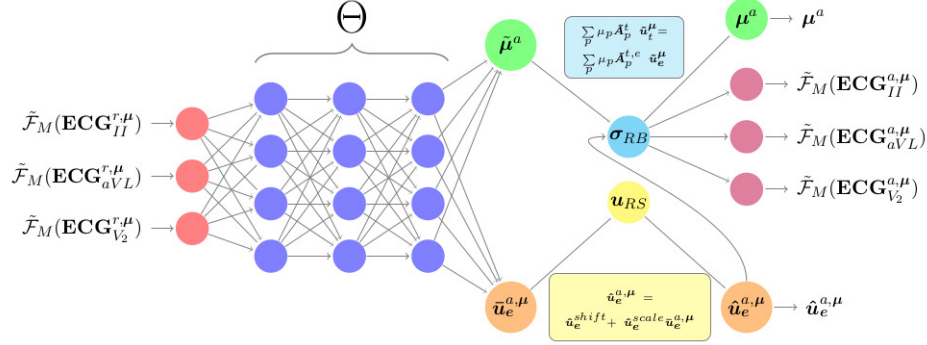


Fig 4. Scheme of the architecture of the DFT-based ST-RB-DNN model. Notice that, as the body surface signals are given in input by means of their lowest-frequency DFT coefficients, the Trainable NN acting as decoder is a MLP, since temporal dynamics have already been handled by the DFT itself at pre-processing stage.

3 Results

In this section, we present the numerical results got with the ST-RB-DNN model on a benchmark test case. The data used to train and test the developed models have been generated numerically. In particular, heart electrophysiology has been approximated by solving the bidomain equations coupled with the phenomenological AP ionic model, on a reference bi-ventricular geometry (taken from [25] - see Fig 5). The FPE required instead to solve a generalized Laplace equation in the human torso. Again, we considered a fixed geometry (taken from [18] - see Fig 6); moreover we modelled the torso as a homogeneous and isotropic volume conductor, discarding the presence of different organs (as bones, lungs or blood vessels) with different electric conductivities. Working under the *isolated heart assumption*, the two problems could be fully decoupled and solved sequentially, one after the other. All numerical simulations have been carried out employing the classical FE method, with elements of degree $p = 1$. Simulation duration has been set to 160 ms , so that only ventricular depolarization was captured. The computational mesh built on top of the bi-ventricular geometry was made of $221'088$ tetrahedral elements, resulting in $49'674$ vertices; the mesh constructed on top of the torso geometry, instead, was made of $498'992$ tetrahedral cells, which resulted in $94'976$ vertices. Three snapshots of a FOM solution to the problem are shown in Fig 7.




Fig 5. Geometry of the bi-ventricle employed in the FOM simulations, taken from [25]. The coloured spheres denote the sites of potential occurrence of epicardial breakthroughs in a healthy patient, according to [52]; the colours of the spheres are associated to the probability of EBT occurrence (refer to the colormap on the left side). (Left): anterior view. (Center-left): posterior view. (Center-right): left lateral view. (Right): inferior view.




Fig 6. Geometry of the human torso employed in the FOM simulations, taken from [18]. The black spheres denote the positions of the 3 peripheral electrodes used in the standard 12-lead ECG system. The coloured spheres denote the positions of the 6 precordial electrodes; their colours have been set accordingly to the American Heart Association (AHA) colour-coding system. The lilac squares, together with the coloured spheres, denote the positions of the electrodes on the "simulated" vest. (Left): anterior view. (Right): posterior view

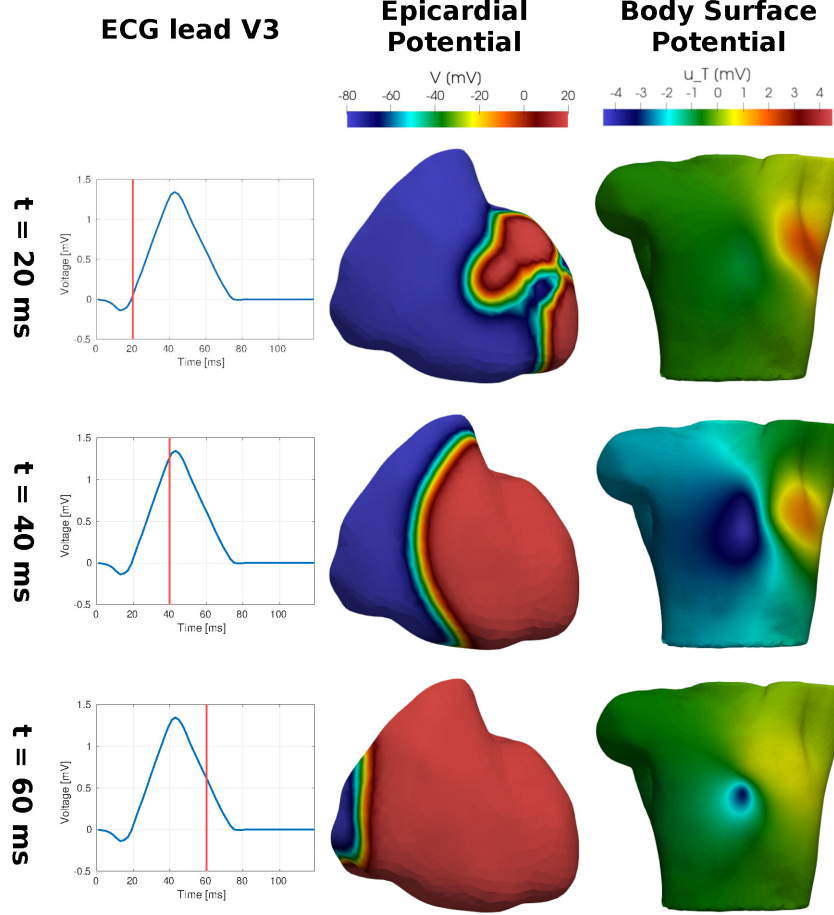


Fig 7. Snapshots of one FOM solution to the FPE at $t = 20, 40, 60$ ms. (Left): V_3 ECG precordial lead. (Center): epicardial extracellular potential field - anterior view. (Right): body surface potential field - anterior view.

A crucial step in order to generate a dataset suitable to train DL models is to add sufficient variability. Indeed, on the one side the training datapoints must differ one from the other, so that the model can learn from them the widest possible amount of dynamics and conditions; on the other side, also the testing datapoints should show differences both within each other and with respect to the training ones, so that a proper assessment of model performances can be made. We are aware of the fact that the simplifying modelling assumptions we adopted prevent us from employing our model with success on real data. However, we tried to enrich the data generation process with several sources of randomness, in order to construct a sufficiently variable dataset, able to challenge the capabilities of the ST-RB-DNN model.

Since the FPE consists only of a generalized Laplace equation to be solved in a ho-

mogeneous and isotropic volume conductor, no variability has been inserted explicitly at such stage. Regarding heart electrophysiology, instead, variability has been added in two ways. First, all cardiac conductivities (longitudinal and transversal to the fibers direction, intracellular and extracellular) have been sampled from a uniform distribution (see Tab 1); indeed they have been proved to severely influence the depolarization process and, in turn, the shape of body surface potentials [4].

$\sigma_{\mathbf{i}}[S\text{ cm}^{-1}]$	$\sigma_{\mathbf{t}_i}[S\text{ cm}^{-1}]$	$\sigma_{\mathbf{l}_e}[S\text{ cm}^{-1}]$	$\sigma_{\mathbf{t}_e}[S\text{ cm}^{-1}]$
$\mathcal{U}(1.77, 2.92) \cdot 10^{-4}$	$\mathcal{U}(0.66, 1.08) \cdot 10^{-4}$	$\mathcal{U}(1.65, 2.73) \cdot 10^{-3}$	$\mathcal{U}(0.96, 1.32) \cdot 10^{-3}$

Tab 1. Values of the cardiac electric conductivities employed to numerically approximate heart electrophysiology via the bidomain equations. $\mathcal{U}(a, b)$ denotes a uniform distribution over the interval $[a; b]$.

Second, different initial activation patterns have been considered; all such patterns are consistent with the findings of *Wyndham et al.* on EBTs localization in both healthy [52] and LBBB-affected [53] patients. In particular, a thin endocardial and sub-endocardial layer is activated, for a duration of 5 ms, in 10 possible different regions; this implies that EBTs could occur in 10 possible different locations, that mimic the effect of Purkinje fibers (see Fig 5). For healthy patients, all these regions could possibly activate, with the constraint that no more than 5 EBTs should appear. Instead, for patients affected by LBBB (RBBB), only the regions located on the right (left) ventricle can be involved in the stimulation protocol and no more than 3 EBTs should be observed. 50% of the patients in the dataset is healthy, 25% is LBBB-affected and 25% is RBBB-affected. Furthermore, the position of the point around which the initial stimulation is applied can vary within a sphere of 2 mm radius, obeying to a 3D uniform distribution, and the initial activation times are sampled from a uniform distribution, with ranges having been defined according to [52, 53].

Ultimately, the dataset we employed is made of $N_\mu = 400$ datapoints. Train and test datasets have been obtained from a 90% – 10% splitting of such dataset; the validation dataset has been assembled by picking the 10% of the training datapoints. Furthermore, training and validation datasets have been subject to data augmentation on-the-fly, by superimposing correlated white Gaussian noise so that an average Signal-to-Noise Ratio (SNR) of 26 dB is achieved. Additionally, starting from the same FOM solutions, two different datasets have been assembled: the first one contains 12-lead ECG signals as body surface signals, while the second one features 158 signals computed from measurements of the body surface potential in 155 different locations (see Fig 6). Both datasets have been suitably modified when the DFT-based model has been considered.

In order to reduce the computational burden of the offline stage of the method, the training and validation datasets have been used also to generate all the Reduced Bases. All POD have been performed with respect to the standard ℓ^2 -norm and the *Randomized SVD* algorithm [23] has been employed. In particular:

- The Reduced Basis in space for the torso potential has been computed by means of

a POD applied to the mode-1 unfolding of the corresponding tensor. A tolerance $\epsilon_{POD}^{t,s} = 10^{-3}$ has been chosen and 314 basis functions have been derived.

- The Reduced Basis in space for the epicardial potential has been computed by means of a POD applied to the mode-1 unfolding of the corresponding tensor. A tolerance $\epsilon_{POD}^{e,s} = 10^{-1}$ has been chosen and 101 basis functions have been derived.
- The dimensionality reduction along the temporal dimension for the epicardial potential has been realized according to the *Tailored Temporal Subspaces via ST-HOSVD* approach proposed in [8], thus computing for each element of the spatial Reduced Basis a tailored temporal one. The tolerances for the POD $\tilde{\epsilon}_{POD}^{e,t}$ have been set to drop from $5 \cdot 10^{-2}$ to $5 \cdot 10^{-1}$, depending on the singular values associated to the corresponding spatial Reduced Basis element. Ultimately, the computed Space-Time Reduced Basis was made of 766 elements.

A grid search has been conducted both on the time-series-based ST-RB-DNN model and on the DFT-based one, in order to identify optimal values for those hyperparameters having the largest impact on model performance. In particular, all hyperparameters related to the architecture of the trainable encoder have been involved in the grid search and the dataset featuring 12-lead ECG signals was employed. The ℓ^1 -norm relative error on the reconstructed epicardial activation maps has been chosen as evaluation metric. Such maps have been computed by assigning to each vertex of the computational mesh the time instant at which the time derivative of the epicardial extracellular potential was maximal. Results of the grid searches are shown in Tab 2 for the time-series-based model and in Tab 3 for the DFT-based one. Additionally, Tab 4 reports the complexities and the training/testing times for some of the trained models. Concerning other relevant hyperparameters: the learning rate has been set to $\nu = 10^{-3}$ (with a reduction of factor 4 if no improvement of the validation loss could be observed for 20 consecutive epochs); *Nadam* optimizer (i.e. an improved version of Adam optimizer incorporating Nesterov momentum) [16] was chosen; Ridge Regularization has been employed, with parameters $\lambda_r^{FC} = 10^{-7}$ in fully-connected layers and $\lambda_r^{Conv} = 10^{-4}$ in convolutional ones. Also, we set $w_{BC} = 10^3$ for the time-series-based model and $w_{BC} = 10^2$ for the DFT-based one; this allowed to get a loss split of approximately 95% – 5% between the error on epicardial potentials and the one on body surface signals. Fig 8 and Fig 9 show the results of the two best models on test datapoint #1 in terms of epicardial activation maps and 12-lead ECG signals, respectively. Fig 10 and Fig 11 show the results of the two best models on test datapoint #3 in terms of epicardial activation maps and 12-lead ECG signals, respectively; these results are relevant, as this datapoint identifies the worst case scenario for the majority of the trained models. Additionally, Fig 12 and Fig 13 report the results of the two best models on test datapoint #3 in case they are trained on the dataset featuring 158 body surface signals. Finally, Tab 5 provides a summary of the errors made by the two best models, comparing the performances in case either 12 or 158 body surface signals are provided in input.

All models training and testing as well as all post-processing routines have been carried out on a *Lenovo ThinkPad T490s* mounting *Ubuntu 20.04.1 LTS*, with 16 GB RAM and an *Intel i7-8565U* processor with 4 cores at 1.80 GHz. The numerical simulations needed to construct the dataset were instead performed on the *IHeart* cluster (*Lenovo SR950* 8x24-Core *Intel Xeon Platinum 8160*, 2.10 GHz and 1.7 TB RAM) at *MOX, Dipartimento di Matematica, Politecnico di Milano*.

Dense \ 1D Conv	(5) - 15	(5) - 25	(10) - 15	(10) - 25	(15) - 15	(15) - 25
(32) - (128,64,32,16,8,4)	6.43e-2	6.02e-2	6.52e-2	6.34e-2	6.71e-2	6.06e-2
(64) - (128,64,32,16,8,4)	5.78e-2	6.44e-2	6.95e-2	5.91e-2	8.28e-2	6.15e-2
(32) - (128,64,32,16,8)	5.15e-2	5.44e-2	4.94e-2	5.56e-2	5.40e-2	4.81e-2
(64) - (128,64,32,16,8)	4.51e-2	4.88e-2	5.10e-2	3.96e-2	4.65e-2	5.28e-2
(32) - (128,64,32,16)	5.20e-2	4.09e-2	4.76e-2	4.65e-2	4.32e-2	4.37e-2
(64) - (128,64,32,16)	4.57e-2	4.32e-2	4.23e-2	4.33e-2	4.35e-2	4.43e-2

Tab 2. Average relative errors in ℓ^1 -norm on the epicardial activation maps of the test dataset with the time-series-based ST-RB-DNN model. The green cell displays the best model; the red cell displays the worst model. Rows labels are in the form "Pre-Layers - Post-Layers", where the first entry defines the layers of the pre-convolutional fully-connected block and the second one the layers of the post-convolutional fully-connected block. Columns labels are of the form " N_F - K_{dim} ", being N_F the number of convolutional filters and K_{dim} the dimension of the 1D convolutional kernel.

Dense \ DFT modes	9	17	25	33
(64,64) - (256,256,128,64)	5.64e-2	5.57e-2	5.29e-2	5.46e-2
(128,128) - (256,256,128,64)	5.46e-2	5.58e-2	5.40e-2	4.06e-2
(256,256) - (256,256,128,64)	5.07e-2	5.08e-2	5.20e-2	4.61e-2
(64,64) - (256,256,128)	4.78e-2	4.54e-2	4.14e-2	5.08e-2
(128,128) - (256,256,128)	4.37e-2	4.28e-2	4.26e-2	4.05e-2
(256,256) - (256,256,128)	3.78e-2	3.81e-2	3.67e-2	3.79e-2
(64,64) - (256,256)	4.19e-2	3.91e-2	3.83e-2	4.85e-2
(128,128) - (256,256)	3.90e-2	3.83e-2	4.12e-2	3.90e-2
(256,256) - (256,256)	3.90e-2	3.70e-2	3.77e-2	3.86e-2
(64,64) - (256,512)	3.77e-2	3.79e-2	3.65e-2	5.07e-2
(128,128) - (256,512)	3.60e-2	3.49e-2	3.69e-2	3.75e-2
(256,256) - (256,512)	3.50e-2	3.51e-2	3.55e-2	3.70e-2

Tab 3. Average relative errors in ℓ^1 -norm on the epicardial activation maps of the test dataset with the DFT-based ST-RB-DNN model. The green cell displays the best model; the red cell displays the worst model; the yellow cells identify two DFT-based models that will be considered in Tab 4. Rows labels are of the form "Pre-Layers - Post-Layers", where the entries identify the number of neurons in the layers preceding and following a Flatten layer. Columns labels represent the number of DFT coefficients given as input to the model.

	AM error	# Params	T_{train}	T_{epoch}	$T_{predict}$	T_{test}
Best T-series	3.96e-2	213'264	603.52 s	11.3871 s	0.0288 s	4.3385 s
Best DFT	3.49e-2	622'590	586.93 s	10.2970 s	0.0329 s	4.6384 s
DFT Model 1	4.78e-2	239'998	446.73 s	9.1180 s	0.0253 s	4.5188 s
DFT Model 2	3.90e-2	358'654	451.92 s	9.2228 s	0.0306 s	4.5453 s

Tab 4. Average activation maps ℓ^1 -norm relative error, number of trainable parameters and training and testing times (in s) for four different ST-RB-DNN models. "Best T-series" and "Best DFT" identify the best time-series-based and DFT-based model respectively; "DFT Model 1" is a DFT-based model whose complexity is the closest to the one of "Best T-series" among the tested models; "DFT Model 2" is a DFT-based model whose average ℓ^1 -norm relative error on activation maps is the closest to the one of "Best T-series" model among the tested models. The architectures of the models can be derived from Tab 2 and Tab 3, following the colour-coded notation and the reported activation maps errors. T_{train} is the full training time; T_{epoch} is the average time per epoch; $T_{predict}$ is the time required by a single forward pass in the model; T_{test} is the full testing time on a single datapoint, including all post-processing routines.

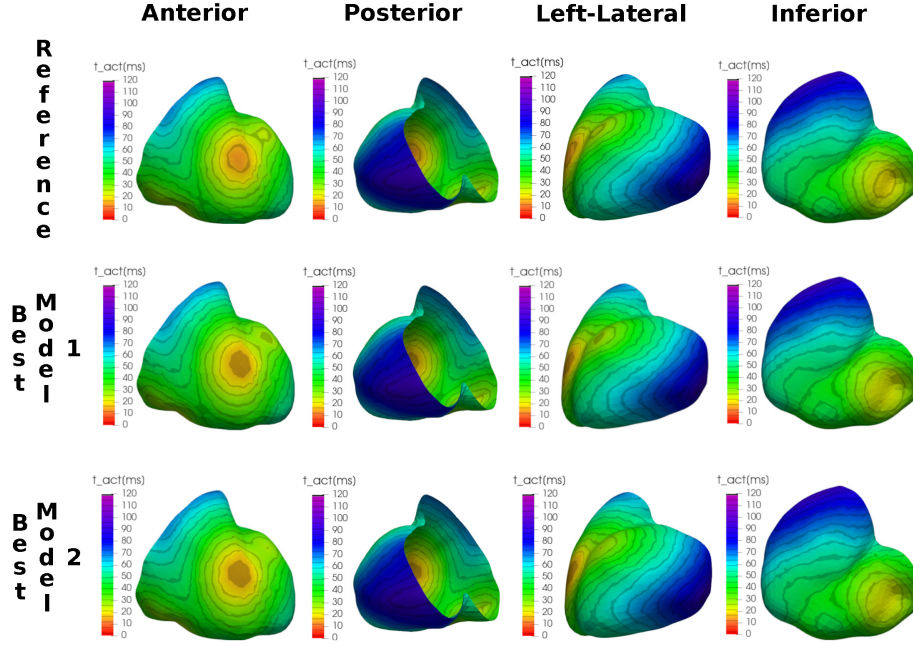


Fig 8. Epicardial activation maps for the test datapoint #1 (healthy), computed from three different potential fields: the one reconstructed from the exact ST-ROM approximation (top), the one estimated by the best time-series-based ST-RB-DNN model (center) and the one estimated by the best DFT-based ST-RB-DNN model (bottom). The errors (relative, ℓ^1 -norm) are 1.60% and 1.78% respectively. Both models get 12-lead ECG signals in input. The activation maps are shown in four different views.

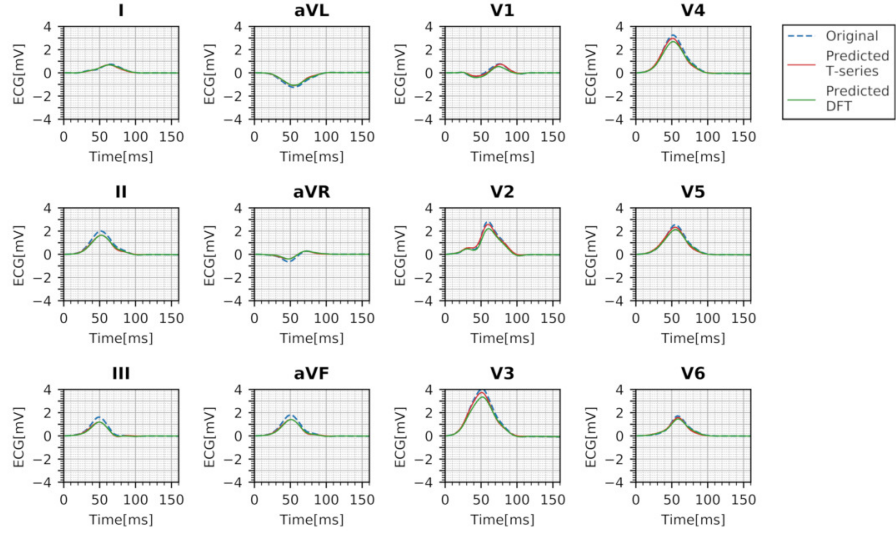


Fig 9. 12-lead ECG signals for test datapoint #1 (healthy) computed by solving the Forward Problem with the RB method. Three epicardial potential fields have been employed as Dirichlet boundary datum: the one reconstructed from the exact ST-ROM approximation, the one estimated by the best time-series-based ST-RB-DNN model and the one estimated by the best DFT-based ST-RB-DNN model. The errors (absolute, ℓ^1 -norm, in mV) are $5.63 \cdot 10^{-2}$ and $7.89 \cdot 10^{-2}$ respectively. Both models get 12-lead ECG in input.

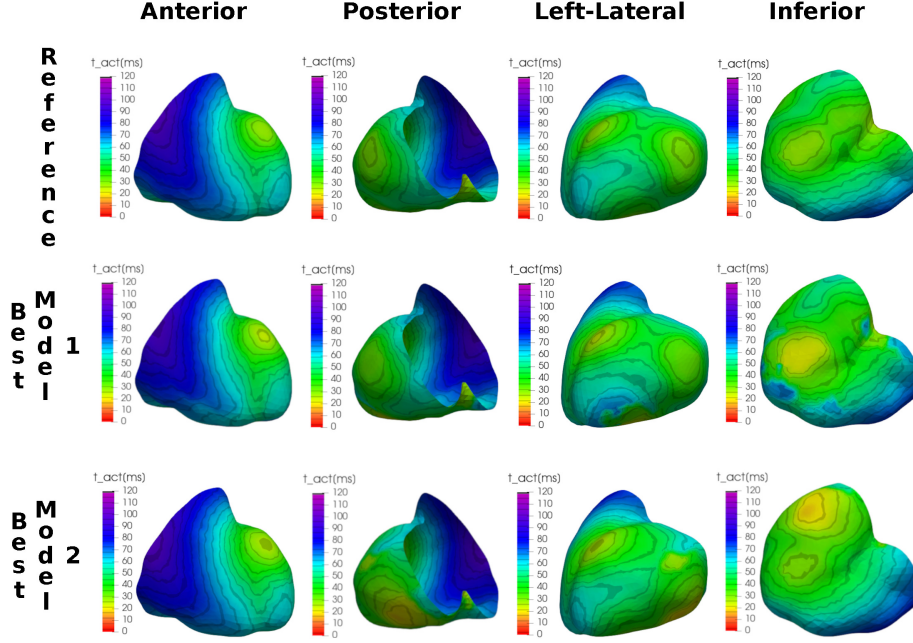


Fig 10. Epicardial activation maps for the test datapoint #3 (RBBB-affected), computed from three different potential fields: the one reconstructed from the exact ST-ROM approximation (top), the one estimated by the best time-series-based ST-RB-DNN model (center) and the one estimated by the best DFT-based ST-RB-DNN model (bottom). The errors (relative, ℓ^1 -norm) are 6.36% and 9.14% respectively. Both models get 12-lead ECG signals in input. The activation maps are shown in four different views.

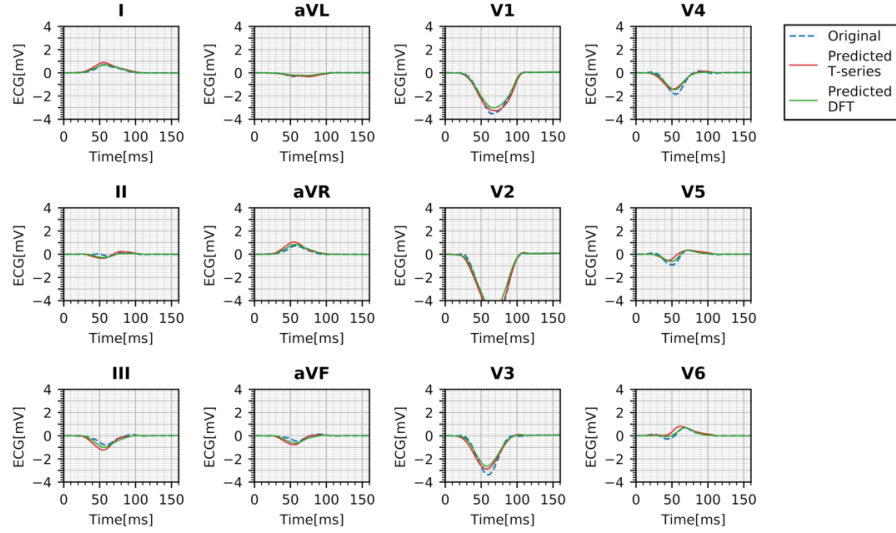


Fig 11. 12-lead ECG signals for test datapoint #3 (RBBB-affected) computed by solving the Forward Problem with the RB method. Three epicardial potential fields have been set as Dirichlet boundary datum: the one reconstructed from the exact ST-ROM approximation, the one estimated by the best time-series based ST-RB-DNN model and the one estimated by the best DFT-based ST-RB-DNN model. The errors (absolute, ℓ^1 -norm, in mV) are $11.51 \cdot 10^{-2}$ and $8.23 \cdot 10^{-2}$ respectively. Both models get 12-lead ECG in input.

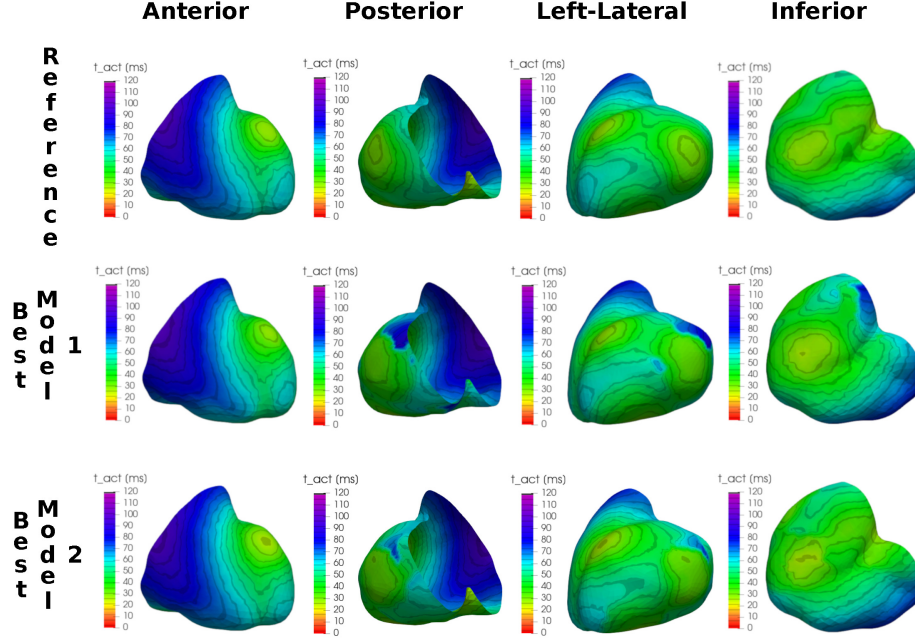


Fig 12. Epicardial activation maps for the test datapoint #3 (RBBB-affected), computed from three different potential fields: the one reconstructed from the exact ST-ROM approximation (top), the one estimated by the best time-series-based ST-RB-DNN model (center) and the one estimated by the best DFT-based ST-RB-DNN model (bottom). The errors (relative, ℓ^1 -norm) are 9.26% and 4.65% respectively. Both models get 158 body surface signals in input. The activation maps are shown in four different views.




Fig 13. 12-lead ECG signals for test datapoint #3 (RBBB-affected) computed by solving the Forward Problem with the RB method. Three epicardial potential fields have been set as Dirichlet boundary datum: the one reconstructed from the exact ST-ROM approximation, the one estimated by the best time-series based ST-RB-DNN model and the one estimated by the best DFT-based ST-RB-DNN model. The errors (absolute, ℓ^1 -norm, in mV) are $11.71 \cdot 10^{-2}$ and $8.19 \cdot 10^{-2}$ respectively. Both models get 158 signals in input.

Size \ Type	T-series	DFT	Size \ Type	T-series	DFT
12	3.96e-2	3.49e-2	12	6.52e-2 mV	6.82e-2 mV
158	3.98e-2	3.29e-2	158	7.55e-2 mV	6.82e-2 mV

Tab 5. Reconstruction errors on the epicardial activation maps (left) and on 12-lead ECG signals (right) made by the best model, depending on the nature and on the dimensionality of the input. Activation maps errors are computed as relative ℓ^1 -norm errors; ECG signals errors are computed as absolute (in mV) ℓ^1 -norm errors.

4 Discussion

Our study focused on the development of a DL model (called ST-RB-DNN) that allows to estimate physically-consistent solutions to the Inverse Problem of Electrophysiology not only leveraging data availability, but also exploiting the awareness of the physical laws underlying the phenomenon of interest, expressed by means of suitable PDEs. In light of the obtained numerical results, the main findings are the listed in the following.

1. The model provides reliable estimations of epicardial activation maps: the time-series-based ST-RB-DNN model trained with 12-lead ECG signals in input presents good performances in terms of accuracy. Tab 2 shows the results of the grid search process on architecture-related hyperparameters (in terms of the average ℓ^1 -norm relative error on the test set); the table is organized so that model complexity increases both row-wise (top to bottom) and column-wise (left to right). Remarkably, none of the considered models features an error higher than 10% and the best one, with an error of 3.96%, presents an intermediate level of complexity (213'264 trainable parameters, see Tab 4), proving that no overfitting did occur. Again referring to Tab 4, we can notice that the training time equals ≈ 604 s (11.4 s per epoch on average); testing on one fresh data sample, instead, takes 4.34 s, of which just 0.029 s are due to the forward pass in the model, while the others are related to the post-processing routines that allow to reconstruct the FOM solution and to compute the epicardial activation map. Activation maps are qualitatively well reconstructed by the "best model", with a good identification of both epicardial breakthrough (EBT) and latest epicardial activation (LEA) areas (see Fig 8, 10 - second row). The level of accuracy we attain allows then to capture a vast part of the variability of the phenomenon at hand, comprising both physiological and pathological (in our case, limited to LBBB and RBBB conduction defects) conditions. Additionally, body surface signals are well approximated in the autoencoding portion of the output (see Fig 9, 11 - red solid line), with an average ℓ^1 -norm absolute error on the test dataset of just $6.52 \cdot 10^{-2}$ mV; the presence of an embedded RB solver of the FPE as deterministic decoder happens to play a key role on this aspect.

2. Signals pre-processing via DFT improves model performances: if the ST-RB-DNN model is adapted so that it can take as input the lowest-frequency coefficients arising from the application of a DFT to 12-lead ECG signals, its performances improve. Tab 3

reports the results of the grid search process on architecture-related hyperparameters and on the number of DFT coefficients; the table is organized so that model complexity increases both row-wise and column-wise. Also in this case, all average ℓ^1 -norm relative errors are below the threshold of 10%; moreover, almost all models feature an error that is lower than 5% and that equals 3.49% for the best model (-0.47% with respect to the best time-series-based model). In this case, the time required by testing on a fresh data sample equals 4.64 s, of which only 0.033 s are related to the forward pass in the model (see Tab 4). About computational complexity, on the one side lower errors are achieved if a limited number of DFT coefficients is given in input. Indeed providing information on high-frequency modes happens to be not only useless, but even counterindicative, as ultimately a spatio-temporal reduced representation of the epicardial potential is estimated. On the other side, the best model has 622'590 trainable parameters and, fixing the input dimensionality, it is the second most complex among the considered ones. Anyway, by proceeding row-wise along each column of Tab 3, it can be noticed that errors seem to reach a plateau, so that no significant improvements should be expected by further increasing the number of hyperparameters; no additional tests have been made on this point. Noticeably, despite featuring a much higher complexity with respect to the best time-series-based model, the best DFT-based one could be trained in a slightly lower amount of time (see Tab 4); this is due to the fact that the trainable encoder of such model is no longer a 1D-CNN, but a simple MLP, hence the backpropagation algorithm can be performed much faster. Also, from Tab 4 it can be noticed that a DFT-based model as complex as the best time-series-based one features a much higher error (4.78% vs 3.96%), while a DFT-based model whose performances are comparable with the ones of the best time-series-based model exhibits a higher complexity (358'654 parameters vs 213'264 parameters), but it could be trained in a much lower amount of time (≈ 452 s vs ≈ 604 s). As expected, the epicardial activation maps estimated by the best DFT-based model are qualitatively close to the target ones (see Fig 8, 10 - third row), showing a good identification of both EBT and LEA locations. Actually, in the worst case scenario presented in Fig 10, the DFT-based model happens to estimate a non-physical EBT in the infero-lateral basal portion of the left ventricle; such misidentifications have been encountered only in a minority of the 40 testing datapoints (for both models) and they were always referred to low-occurrence EBTs. For this reason, we expect the problem to be avoided by either just increasing the dimensionality of the training dataset or upsampling the critical EBTs at the data generation stage. Finally, body surface signals have been well reconstructed at all leads in both reported test datapoints (see Fig 9, 11 - green solid line), with an average ℓ^1 -norm absolute error on the test dataset equalling $6.82 \cdot 10^{-2}$ mV ($+0.30 \cdot 10^{-2}$ mV with respect to the best time-series-based model).

3. Considering more body surface signals does not severely impact on model performances: the best model architectures, either trained with 158 body surface signals or with "standard" 12-lead ECG ones, exhibited comparable performances. In particular, the time-series-based model featured an average ℓ^1 -norm relative error of 3.98% ($+0.02\%$), while that drops to 3.29% for the DFT-based one (-0.20%). The reduced dimensionality

and heterogeneity of the training dataset is thought to exert a big impact on this aspect, as 12 body surface signals may be enough to encode the variability of the considered data. Thus, the effective usefulness of providing more body surface signals - which is common in ECGI clinical applications - should be assessed when dealing with bigger datasets. This consideration is enhanced by looking at Fig 12, which reports the epicardial activation maps estimated by the best models taking 158 body surface signals in input; indeed it can be noticed that the DFT-based model no longer estimates the non-physical EBT on the infero-lateral basal portion of the left ventricle, which was instead detected by the corresponding model taking 12-lead ECG signals in input (see Fig 10).

4. Exploiting physical-awareness allows to construct low-complexity and fast-training models in a *small data regime*: as reported in Tab 4, all the implemented ST-RB-DNN models, despite the small amount of synthetic data at our disposal, managed to provide good estimations of the epicardial potential field during ventricular depolarization and could be trained on a simple laptop in approximately 10 *min*. The two elements that allowed us to get physical-awareness play a key role in this, indeed:

- The projection of the epicardial potential onto a space-time reduced subspace, generated from the same physical FOM solutions that have been employed to construct the training dataset, dramatically reduces the number of coefficients to be estimated by the trainable encoder, other than allowing to achieve independence from the space-time mesh refinement and to constrain the solution to a lower-dimensional and physically-consistent manifold. In this way, the training process of the model, via backpropagation algorithm, becomes feasible, while it would have been extremely difficult, as well as severely expensive, if FOM approximations of the epicardial potential field had had to be predicted (see [30] for additional details). In addition, the reduced dimensionality of the dataset allowed us to train the models on a simple laptop. If much more data were at our disposal (as it would be necessary for effective clinical validation), we would need for more computational resources and for a more optimized training pipeline. However, the trainability of the models is not expected to change in any way.
- The insertion of an embedded RB solver for the FPE as deterministic decoder and its employment as a physically-aware regularization agent in the expression of the loss functional (see Eq (8)) allowed us to design a physics-aware NN, similar to the ones presented in [15]. As stated in [39], building ML models able to exploit knowledge of the most relevant physical laws underlying the phenomenon of interest proved to be effective in the *small data regime*, i.e. when the amount of data at disposal is limited due to the high cost/complexity of the acquisition procedures. In fact, just starting from a small dataset generated *in silico*, we managed to train low-complexity models, showing a remarkable representation power.

4.1 Limitations

The current work represents a first methodological attempt of tackling the Inverse Problem of Electrocardiography by combining knowledge of the main physical laws underlying the problem at hand, expressed via of suitable PDEs/ODEs, with deep learning techniques for the exploitation of data availability. Configuring as a preliminary analysis, we have decided to work under several simplifying assumptions that should be addressed in order to bridge the gap with clinical applications. In particular, the most relevant limitations are listed in the following:

1. Simplified Dataset: the ST-RB-DNN model training necessitates of many epicardial activation maps. Unfortunately, acquiring such maps from clinical measurements requires invasive procedures. To circumvent this shortcoming, we trained and tested our models with *in silico* data, obtained by numerically approximating the solutions to heart electrophysiology and to the FPE. Due to the high computational costs required by accurate FOM simulations, we made several simplifying assumptions as the usage of a cardiac geometry without atria, the employment of not too refined computational meshes and the neglect of conductivity inhomogeneities in the torso. Furthermore, we imposed a one-way coupling between the two aforementioned problems via the *isolated heart assumption*, discarding the continuity of the currents at the epicardial surface. Despite appearing relevant from the physical point of view, this last assumption is not expected to severely impact on our results. Indeed exposing the heart to insulating air has been proven, both numerically [4] and experimentally [22], to determine an increase in the epicardial potential magnitudes, but without affecting in most cases the activation pattern, that we wish to reconstruct. Even if we have made some efforts in reproducing body surface signals exhibiting the expected polarities and amplitudes, the results we got still show some differences with respect to data collected *in vivo*. As a consequence, the models we implemented cannot be directly employed with success on the latter.

2. Reference Geometry: all the training and testing datapoints have been generated on the same, reference, geometry. On the one side this allowed us to develop an imaging-free model, but on the other side it clearly represents a limitation. In general, it could be possible to derive a parametrization of the Forward Problem with respect to the heart and torso geometries; however, all geometric parameters would have to be estimated by the model, significantly increasing the complexity of the task to be performed. Even considering fixed geometries for both the heart and the torso, it is broadly known that the position and the rotation of the former inside the latter can vary considerably, both from patient to patient and, for the same patient, from heartbeat to heartbeat, exerting a non-negligible impact on body surface signals. These effects could be taken into account by using methods that allow to track the position of the heart from body surface potentials, as the one presented in [12]. In terms of the ST-RB-DNN model, just few modifications would be necessary as the adaptation of the architecture to a new input structure and the implementation of an embedded RB solver for the Forward Problem that suitably handles novel geometry-related

parameter values. Conversely, many problems could arise at data generation stage, as smart ways of both computing solutions to the FOM problem in continuously varying geometrical settings and deriving spatio-temporal Reduced Bases for the quantities of interest should be figured out.

3. White Noise: another key aspect that differentiates clinically measured body surface potentials from numerically approximated ones is the presence of noise. We tried to take this element into consideration by superimposing to artificial signals correlated white Gaussian noise at an average SNR of ≈ 26 dB. However, this is a very simplistic approach, as real body surface potentials feature the presence of different sources of noise, which are due to various aspects (muscular contraction/relaxation, breathing, measurement instruments and procedures etc.) and whose distribution is not necessarily Gaussian (see [43, 44]). Thus, an interesting development towards clinical application would be to superimpose to the numerically generated signals more realistic noise components, as for instance done in [17] by means of a tailored synthetic noise generator.

4. Limited Design Research: an additional limitation concerns the NN design, as only a limited amount of architectures has been tested for the encoder. In particular, we chose to employ a 1D-CNN in the time-series-based model and a MLP in the DFT-based one. Actually, during the first stages of development, we also considered other architectures as Residual Networks (*ResNet*) [28] or Convolutional Recurrent Neural Networks (CRNN) [55]; however, as no significant improvements were observed, no further efforts have been made in this direction. In general, the field offers plenty of possibilities in terms of architectural design and of optimization algorithms and additional efforts in this sense are very likely to bring to better models.

5 Conclusion

The present study concerned the development of a physics-aware deep learning model (called ST-RB-DNN), able to provide physically-consistent and data-driven solutions to the Inverse Problem of Electrocardiography. In the numerical test cases that have been conducted, all variants of the ST-RB-DNN model (distinguished based on the dimensionality and nature of the provided body surface signals) proved to be accurate, up to an average error of $\approx 3.5\%$ in relative ℓ^1 -norm, considering the epicardial activation maps of the test dataset. Additionally, despite working under several simplifying assumptions and employing just a small amount of numerically generated data for both training and testing, our models also proved to be highly efficient; their trainings could be all performed in ≈ 10 min on a simple laptop and a single forward pass only required ≈ 0.03 s. Finally, the obtained results indicate that our method has potential to be used in clinical practice, notwithstanding some necessary improvements, mostly concerning the data generation stage.

Supporting information

S1 Appendix. Physical modelling and model order reduction. This appendix contains further details on the physical modelling and numerical approximation of cardiac electrophysiology and of the Forward Problem of Electrocardiography; additionally the model order reduction techniques exploited in the ST-RB-DNN model are presented.

Acknowledgments

SP, AQ have received funding from the *European Research Council (ERC)* under the *European Union's Horizon 2020 research and innovation programme* (grant agreement No 740132, *iHEART - An Integrated Heart Model for the simulation of the cardiac function*, P.I. Prof. A. Quarteroni). RT, SD have received funding from the *Institute of Mathematics* of EPFL.

S1 Appendix.

Physical modelling and model order reduction

Physical modelling

Due to the paucity of clinical data relative to measurements of the electric potential at the epicardium, the datasets we employed for both training and testing the ST-RB-DNN model have been generated *in silico*. In particular, this involved the modelization and the numerical approximation of both cardiac electrophysiology and the propagation of the electric signal from the epicardium to the body surface through the torso (known as the Forward Problem of Electrocardiography [3]). As anticipated, all simulations have been carried out in a *fixed geometrical setting* (see Section *Results*) and the different datapoints have been generated by randomly changing the values of the main parameters characterizing the selected cardiac model.

Cardiac electrophysiology

Cardiac electrophysiology has been modelled via the bidomain equations [48], coupled with the Aliev-Panfilov (AP) ionic model [1]. The former are characterized by a partitioning of the heart domain Ω_H in an intracellular domain Ω_H^i and an extracellular one Ω_H^e , which however coexist on the same space upon a homogenization process [21]. This reflects in the definition of an intracellular potential u_i and of an extracellular potential u_e ; also the trans-membrane potential can be defined as $v := u_i - u_e$. Ultimately, the bidomain model is expressed via the following system of time-dependent PDEs:

$$\begin{cases} A_m \left(C_m \frac{\partial v}{\partial t} + I_{ion}(v, w; \boldsymbol{\mu}) \right) - \nabla \cdot (\mathbf{D}_i(\boldsymbol{\mu}) \nabla v) - \nabla \cdot (\mathbf{D}_e(\boldsymbol{\mu}) \nabla u_e) = A_m I_{app}(\boldsymbol{\mu}) & \text{in } \Omega_H \times [t_0, T] \\ -\nabla \cdot (\mathbf{D}_i(\boldsymbol{\mu}) \nabla v) - \nabla \cdot (\mathbf{D}_e(\boldsymbol{\mu}) \nabla u_e) = 0 & \text{in } \Omega_H \times [t_0, T] \end{cases} \quad (11)$$

Here A_m represents the area of cell membrane per unit volume, C_m the membrane capacitance per unit area, \mathbf{D}_i and \mathbf{D}_e the conductivity tensors in the intracellular and extracellular domains respectively (which may take into account the anisotropy of the cardiac tissue), I_{app} an externally applied current, typically responsible for the initial activation of the heart tissue, and I_{ion} the ionic current across the cell membrane. The simulations are conducted over the time interval $[t_0; T]$. Finally, $\boldsymbol{\mu}$ represents a vector which stores all the scalar parameters that characterize the problem; notice that the expressions of I_{app} , I_{ion} and $\mathbf{D}_{i,e}$ depend on $\boldsymbol{\mu}$, while the domain Ω_H is assumed to be fixed. System (11) must be coupled with proper initial (on v only) and boundary conditions. Concerning the latter, in particular, we have chosen to work under the *isolated heart assumption*, thus keeping the cardiac potentials independent of the one in the torso. This is achieved by imposing homogeneous Neumann boundary conditions at the epicardial surface Γ_H for

both u_e and v ; the impact of this simplifying assumption has been extensively analysed and discussed, for instance, in [4]. The anisotropy of the cardiac tissue, dictated by the presence of conducting fibers, has been taken into account; specifically, different electrical conductivities have been set longitudinally and transversally to the fibers orientation, both in the intracellular and in the extracellular domain [37]. For simplicity, we have supposed all cardiac conductivities to be constant in Ω_H (*homogeneous anisotropy assumption*) and the two transversal conductivities, coplanar and orthogonal to the fibers sheet, to be equal (*axially isotropic case*) [13]. Ultimately, the conductivity tensors could be written as:

$$\mathbf{D}_{i,e}(\mathbf{x}) = \sigma_{t_{i,e}} \mathbf{I} + (\sigma_{l_{i,e}} - \sigma_{t_{i,e}}) \mathbf{a}_l(\mathbf{x}) \mathbf{a}_l^T(\mathbf{x}) \quad \mathbf{x} \in \Omega_H \quad (12)$$

being $\mathbf{a}_l(\mathbf{x})$ a unit vector parallel to the fibers orientation at a point $\mathbf{x} \in \Omega_H$.

The AP model, instead, is a phenomenological ionic model that allows to approximate the current I_{ion} induced by the movement of electrically charged ions through the cell membrane. It configures as an extension of the FitzHugh-Nagumo (FHN) model [27] and it proved to exceed its performances on cardiomyocytes, leading to a much better approximation of the action potential shape and duration and of the cardiac tissue restitution curve. I_{ion} takes the form of a cubic non-linear function of v and a single (dimensionless) gating variable w plays the role of a recovery function, allowing to model the refractoriness of cells disregarding any sub-cellular process. From a mathematical standpoint, the AP model is expressed by the following equations:

$$\begin{cases} I_{ion}(v, w; \boldsymbol{\mu}) = Kv(v - a)(v - 1) + v & \text{in } \Omega_H \times [t_0, T] \\ \frac{\partial w}{\partial t} = g(v, w; \boldsymbol{\mu}) = C(w)(-w - Kv(v - b - 1)) & \text{in } \Omega_H \times [t_0, T] \end{cases} \quad (13)$$

where the parameters $K, a, b, \epsilon_0, \mu_1, \mu_2$ are all related to the cell. In particular, a acts as an oscillation threshold, above which an action potential is fired, and the weighting term

$$C(w) = \epsilon_0 + \frac{\mu_1 w}{v + \mu_2} \quad (14)$$

has been specifically added (compared to the FHN model) to account for a finer tuning of the myocardial restitution curve. Additionally, a suitable initial condition has to be imposed on w (typically $w(t = t_0) = 0$) and, in order to get the desired scaling properties, it is necessary to scale the time variable as $t[ms] = 12.9 t[t.u.]$.

Upon a Galerkin Finite Element (FE) approximation, the discrete-in-space continuous-in-time formulation of the bidomain equations coupled with any phenomenological ionic model reads as follows.

Problem 1. Given $\boldsymbol{\mu} \in \mathcal{P} \subset \mathbb{R}^p$, find $\mathbf{u}_{eh} = \mathbf{u}_{eh}(t; \boldsymbol{\mu})$, $\mathbf{v}_h = \mathbf{v}_h(t; \boldsymbol{\mu})$ and $\mathbf{w}_h = \mathbf{w}_h(t; \boldsymbol{\mu})$ such that:

$$\begin{cases} A_m C_m \mathbf{M} \frac{\partial \mathbf{v}_h}{\partial t} + \mathbf{A}_{in}(\boldsymbol{\mu}) \mathbf{v}_h + \mathbf{A}_{in}(\boldsymbol{\mu}) \mathbf{u}_{eh} + & t \in [t_0, T] \\ A_m \mathbf{I}_{ion}(\mathbf{v}_h, \mathbf{w}_h; \boldsymbol{\mu}) = A_m \mathbf{I}_{app}(t; \boldsymbol{\mu}) & \\ \mathbf{A}_{in}(\boldsymbol{\mu}) \mathbf{v}_h + (\mathbf{A}_{in}(\boldsymbol{\mu}) + \mathbf{A}_{ex}(\boldsymbol{\mu})) \mathbf{u}_{eh} = 0 & t \in [t_0, T] \\ \frac{\partial \mathbf{w}_h}{\partial t} = g(\mathbf{v}_h, \mathbf{w}_h; \boldsymbol{\mu}) & t \in [t_0, T] \\ \mathbf{v}_h(t_0; \boldsymbol{\mu}) = \mathbf{v}_0(\boldsymbol{\mu}); \quad \mathbf{w}_h(t_0; \boldsymbol{\mu}) = \mathbf{w}_0(\boldsymbol{\mu}) & \end{cases} \quad (15)$$

where \mathbf{M} is the mass matrix and \mathbf{A}_{in} , \mathbf{A}_{ex} are the intracellular and extracellular stiffness matrices, respectively.

The Forward Problem of Electrocardiography

In modelling the Forward Problem of Electrocardiography (FPE), we can approximate the human torso as an isotropic and inhomogeneous volume conductor; indeed, as reported in [45], the torso is not made of excitable cells, as the heart, and thus there is no need for a model able to simulate the behaviour of the transmembrane potential. The inhomogeneity derives from the presence of different organs (bones, lungs, blood vessels, etc.), which are characterized by different electric conductivities.

Under the *isolated heart assumption*, at any $t \in [t_0; T]$ the torso potential u_T can be computed by solving a generalized Laplace equation of the form

$$-\nabla \cdot (\mathbf{D}_T \nabla u_T) = 0 \quad \text{in } \Omega_T \quad (16)$$

coupled with a homogeneous Neumann boundary condition at the body surface Γ_B and with an inhomogeneous Dirichlet one at the epicardium Γ_H , where the extracellular potential u_e acts as boundary datum. In Eq (16) Ω_T is the reference torso domain and \mathbf{D}_T represents the torso conductivity tensor.

As discussed in [4], working under the *isolated heart assumption* provides the advantage of getting a *one-way coupling* between the heart and torso problems, easing and speeding up the computations. It also prevents from imposing the continuity of the electric potential fluxes at the heart-torso interface; in terms of body surface potentials, this results in signals that often show the correct shape, but that feature abnormal magnitudes.

Upon a Galerkin FE approximation, the problem writes as follows.

Problem 2. Given $\mu \in \mathcal{P} \subset \mathbb{R}^p$, find $\mathbf{u}_{t_h} = \mathbf{u}_{t_h}(t; \mu)$ such that:

$$\mathbf{A}^t(\mu) \mathbf{u}_{t_h}(t) = -\mathbf{A}_{\Gamma_H}^t(\mu) \mathbf{u}_{e_h}^{\Gamma_H}(t) \quad t \in [t_0, T] \quad (17)$$

where $\mathbf{A}_{\Gamma_H}^t(\mu) =: \mathbf{A}^t(\mu)[\cdot, \{j_{Dir}\}]$, being \mathbf{A}^t is the torso stiffness matrix and $\{j_{Dir}\}$ the set of Degrees of Freedom (DOFs) at which the Dirichlet boundary condition is imposed. Furthermore $\mathbf{u}_{e_h}^{\Gamma_H}(t)$ represents the FE discretization of the trace of the extracellular potential u_e on the epicardial surface Γ_H at the time instant t .

Model order reduction

In the ST-RB-DNN model, Reduced Order Model (ROM) techniques have been used both to estimate physically meaningful and dimensionality-reduced latent representations of the epicardial potential fields and to efficiently solve the generalized Laplace equation in the torso, in order to reconstruct signals matching the ones given in input. Considering that Problem 2 is stationary, the latter task can be carried out by resorting to the classical RB method [38], i.e. by deriving a basis for a dimensionality-reduced subspace in space (via Proper Orthogonal Decomposition (POD) of the snapshots' tensor), performing a projection along the spatial dimension and solving the resulting reduced problem independently at all the discrete time instants. The former task, instead, can be much better accomplished by resorting to spatio-temporal ROM techniques, that allow to encode the information coming from time-dependent fields into a very low number of coefficients, (almost) independent of the level of refinement of the grid along both the spatial and the temporal dimension. In particular, the space-time-reduced subspace where the epicardial potential fields have been projected has been generated using the *Tailored Temporal Subspaces via ST-HOSVD* approach proposed in [8]. This strategy is based on the computation of a tailored temporal basis for each element of the spatial one, allowing for a dramatic reduction of the dimensionality of the resulting spatio-temporal subspace. We notice that estimating a FOM approximation of the epicardial potential field is not a viable option in the framework of the ST-RB-DNN model as that would significantly increase the overall complexity, severely hindering the training process. Reference [30] features a detailed theoretical analysis on this aspect, that can be also adapted to the present case.

In the following, we employ the aforementioned ROM techniques in the ST-RB-DNN model, in order to efficiently solve the FPE inside an embedded tensorial solver. Consider the third order tensor $\mathcal{X}^t \in \mathbb{R}^{N_s^t \times N_t \times N_\mu}$ storing the solutions to (17) at all the Space-Time $N_s^t N_t$ DOFs and for N_μ different parameter values. The basis in space, encoded by the matrix $\mathbf{V}_s^t \in \mathbb{R}^{N_s^t \times n_s^t}$, is computed by applying a truncated POD to the mode-1 unfolding of \mathcal{X}^t . Once such basis is derived, by projecting all quantities onto the dimensionality-reduced subspace it spans, we get:

$$\tilde{\mathbf{A}}^t(\mu) \tilde{\mathbf{u}}_{t_h}^{(l)} = \tilde{\mathbf{u}}_{e_h}^{\Omega_T^{(l)}}(\mu) \quad l \in \{0, \dots, N_t\}, \quad (18)$$

where, being $\mathbf{A}^t(\boldsymbol{\mu})$ and $\mathbf{u}_{e_h}^{\Omega T}(\boldsymbol{\mu})$ defined as in Problem 2:

$$\tilde{\mathbf{A}}^t(\boldsymbol{\mu}) =: \mathbf{V}_s^{t^T} \mathbf{A}^t(\boldsymbol{\mu}) \mathbf{V}_s^t \in \mathbb{R}^{n_s^t \times n_s^t}, \quad (19a)$$

$$\tilde{\mathbf{u}}_{t_h}^{(l)} =: \mathbf{V}_s^{t^T} \mathbf{u}_{t_h}^{(l)} \in \mathbb{R}^{n_s^t} \quad l \in \{0, \dots, N_t\}, \quad (19b)$$

$$\tilde{\mathbf{u}}_{e_h}^{\Omega T(l)}(\boldsymbol{\mu}) =: -\mathbf{V}_s^{t^T} \mathbf{A}^t(\boldsymbol{\mu}) \mathbf{u}_{e_h}^{\Omega T(l)}(\boldsymbol{\mu}) \in \mathbb{R}^{n_s^t} \quad l \in \{0, \dots, N_t\}. \quad (19c)$$

Based on Eq (17), we can perform a dimensionality reduction step also on the epicardial extracellular potential. So, let $\mathcal{X}^e \in \mathbb{R}^{N_s^e \times N_t \times N_\mu}$ be the third-order tensor storing the values of the epicardial extracellular potential $\mathbf{u}_{e_h}^{\Gamma H}$ at all the $N_s^e N_t$ Space-Time DOFs of the epicardial surface and for N_μ different parameter values. Then a reduced basis in space, stored by the matrix $\mathbf{V}_s^e \in \mathbb{R}^{N_s^e \times n_s^e}$, can be computed by applying a truncated POD to the mode-1 unfolding of \mathcal{X}^e . Eq (17) can so be written in reduced form as:

$$\tilde{\mathbf{A}}^t(\boldsymbol{\mu}) \tilde{\mathbf{u}}_{t_h}^{(l)} = \tilde{\mathbf{A}}^{t,e}(\boldsymbol{\mu}) \tilde{\mathbf{u}}_{e_h}^{\Gamma H(l)}(\boldsymbol{\mu}) \quad l \in \{0, \dots, N_t\}, \quad (20)$$

where $\tilde{\mathbf{A}}^t(\boldsymbol{\mu})$ and $\tilde{\mathbf{u}}_{t_h}^{(l)}$ are defined as in (19a) and (19b) respectively, while

$$\tilde{\mathbf{A}}^{t,e}(\boldsymbol{\mu}) =: \mathbf{V}_s^{t^T} \mathbf{A}^t(\boldsymbol{\mu}) \mathbf{V}_s^e \in \mathbb{R}^{n_s^t \times n_s^e}, \quad (21a)$$

$$\tilde{\mathbf{u}}_{e_h}^{\Gamma H(l)}(\boldsymbol{\mu}) =: -\mathbf{V}_s^{e^T} \mathbf{u}_{e_h}^{\Gamma H(l)}(\boldsymbol{\mu}) \in \mathbb{R}^{n_s^e} \quad l \in \{0, \dots, N_t\}. \quad (21b)$$

As discussed before, in the framework of the ST-RB-DNN model the epicardial potential field is projected also along the temporal dimension, as it is extremely convenient to encode its dynamics in a low and grid-refinement-independent number of coefficients. Following the *Tailored Temporal Subspaces via ST-HOSVD* approach, n_s^e different truncated POD are applied to the projections of the snapshots' tensor \mathcal{X}^e onto the spaces spanned by the different spatial basis elements, i.e.

$$\mathcal{X}^e(\mathbf{V}_{s_i}^e) \in \mathbb{R}^{N_t \times N_\mu} \quad s.t. \quad (\mathcal{X}^e(\mathbf{V}_{s_i}^e))_{jk} = \sum_{l=1}^{N_s^e} \mathcal{X}_{ljk}^e \mathbf{V}_{s_{il}}^e, \quad i \in \{1, \dots, n_s^e\}. \quad (22)$$

These POD allow to compute the temporal bases $\mathbf{V}_{t_i}^e \in \mathbb{R}^{N_t \times n_t^i} \quad \forall i \in \{1, \dots, n_s^e\}$. The overall Space-Time Reduced Basis for the epicardial potential is then made of $n_{st} =: \sum_{i=1}^{n_s^e} n_t^i$ elements of dimension $N_s^e \times N_t$, each one being defined as the outer product between an element of the Reduced Basis in space \mathbf{V}_s^e and an element of the associated Reduced Basis in time $\{\mathbf{V}_{t_i}^e\}_{i=1}^{n_s^e}$.

The embedded reduced solver acting as decoder in the ST-RB-DNN model computes the ROM-in-space FOM-in-time solution to the generalized Laplace equation in the torso (see Eq (20)) receiving as input the space-time projection of the epicardial potential field.

Leveraging orthonormality due to the POD, such projection can be expressed in terms of the FOM approximation as

$$\left(\hat{\mathbf{u}}_{e_h}^{\Gamma_H}(\boldsymbol{\mu})\right)_k = \sum_{i=0}^{N_s^e} \sum_{j=0}^{N_t} \boldsymbol{\pi}_{ij}^k \left(\mathbf{u}_{e_h}^{\Gamma_H}(\boldsymbol{\mu})\right)_{ij} \quad k \in \mathbb{N}(n_{st}), \quad (23)$$

where $\boldsymbol{\pi}_{ij}^k$ denotes the element in position (i, j) of the k -th Space-Time basis function. Therefore, the ROM-in-space FOM-in-time epicardial potential $\tilde{\mathbf{u}}_{e_h}^{\Gamma_H}(\boldsymbol{\mu})$ can be derived by expanding in time $\hat{\mathbf{u}}_{e_h}^{\Gamma_H}(\boldsymbol{\mu})$ as:

$$\left(\tilde{\mathbf{u}}^{\Gamma_H(l)}(\boldsymbol{\mu})\right)_i = \sum_{j=0}^{n_t^i} (\mathbf{V}_{t_i}^e)_{jl} \hat{\mathbf{u}}_{\mathcal{F}(i,j)}^{\Gamma_H}(\boldsymbol{\mu}) \quad i \in \{1, \dots, n_s^e\} \quad l \in \{0, \dots, N_t\}, \quad (24)$$

where

$$\mathcal{F} : (\mathbb{N}(n_s^e), \mathbb{N}(n_t^i)) \ni (i, j) \rightarrow \sum_{k=1}^{i-1} n_t^k + j \in \mathbb{N}(n_{st}) \quad (25)$$

is the mapping from the Space and Time bases indexes to the Space-Time basis index.

Remark: if the torso is approximated as an inhomogeneous volume conductor, featuring the presence of organs with different electric conductivities, then it is possible to take advantage of an affine parametrization to efficiently assemble the stiffness matrix during the online stage. For instance, if we suppose the torso to be partitioned into N_p^t parts with conductivities $\boldsymbol{\mu}^t =: \{\sigma_p^t\}_{p=1}^{N_p^t}$, then the stiffness matrix can be written as

$$\mathbf{A}^t(\boldsymbol{\mu}^t) = \sum_{p=1}^{N_p^t} \sigma_p^t \mathbf{A}_p^t \quad (26)$$

where \mathbf{A}_p^t is the stiffness matrix associated to the DOFs belonging to the part p of the torso and computed assuming a default conductivity of 1. Readily, an equivalent affine decomposition applies also to the reduced torso stiffness matrix $\tilde{\mathbf{A}}^t(\boldsymbol{\mu}^t)$.

References

1. Rubin R Aliev and Alexander V Panfilov. A simple two-variable model of cardiac excitation. *Chaos, Solitons & Fractals*, 7(3):293–301, 1996.
2. Tania Bacoyannis, Julian Krebs, Nicolas Cedilnik, Hubert Cochet, and Maxime Sermesant. Deep learning formulation of ECGI for data-driven integration of spatiotemporal correlations and imaging information. In *International Conference on Functional Imaging and Modeling of the Heart*, pages 20–28. Springer, 2019.
3. Laura R Bear, Leo K Cheng, Ian J LeGrice, Gregory B Sands, Nigel A Lever, David J Paterson, and Bruce H Smaill. Forward problem of electrocardiography: is it solved? *Circulation: Arrhythmia and Electrophysiology*, 8(3):677–684, 2015.
4. Muriel Boulakia, Serge Cazeau, Miguel A Fernández, Jean-Frédéric Gerbeau, and Nejib Zemzemi. Mathematical modeling of electrocardiograms: a numerical study. *Annals of biomedical engineering*, 38(3):1071–1097, 2010.
5. Kamil Bujnarowski, Pietro Bonizzi, Matthijs Cluitmans, Ralf Peeters, and Joel Karel. CT-scan free neural network-based reconstruction of heart surface potentials from ECG recordings. 2020.
6. Stephen Butterworth et al. On the theory of filter amplifiers. *Wireless Engineer*, 7(6):536–541, 1930.
7. Shengze Cai, Zhicheng Wang, Lu Lu, Tamer A Zaki, and George Em Karniadakis. Deepm&mnet: Inferring the electroconvection multiphysics fields based on operator approximation by neural networks. *arXiv preprint arXiv:2009.12935*, 2020.
8. Youngsoo Choi and Kevin Carlberg. Space-Time Least-Squares Petrov-Galerkin projection for nonlinear model reduction. *SIAM Journal on Scientific Computing*, 41(1):A26–A58, 2019.
9. Matthijs Cluitmans, Dana H Brooks, Rob MacLeod, Olaf Dössel, María S Guillem, Peter M van Dam, Jana Svehlikova, Bin He, John Sapp, Linwei Wang, et al. Validation and opportunities of electrocardiographic imaging: from technical achievements to clinical applications. *Frontiers in physiology*, 9:1305, 2018.
10. Matthijs J M Cluitmans, RLM Peeters, RL Westra, and PGA Volders. Noninvasive reconstruction of cardiac electrical activity: update on current methods, applications and challenges. *Netherlands Heart Journal*, 23(6):301–311, 2015.
11. Matthijs JM Cluitmans, Pietro Bonizzi, Joël MH Karel, Marco Das, Bas LJH Kietselaer, Monique MJ de Jong, Frits W Prinzen, Ralf LM Peeters, Ronald L Westra, and Paul GA Volders. In vivo validation of electrocardiographic imaging. *JACC: Clinical Electrophysiology*, 3(3):232–242, 2017.
12. Jaume Coll-Font and Dana H Brooks. Tracking the position of the heart from body surface potential maps and electrograms. *Frontiers in physiology*, 9:1727, 2018.
13. Piero Colli Franzone and Luca F Pavarino. A parallel solver for reaction–diffusion systems in computational electrocardiology. *Mathematical models and methods in applied sciences*, 14(06):883–911, 2004.
14. James W Cooley and John W Tukey. An algorithm for the machine calculation of complex Fourier series. *Mathematics of computation*, 19(90):297–301, 1965.

15. Niccolò Dal Santo, Simone Deparis, and Luca Pegolotti. Data driven approximation of parametrized PDEs by Reduced Basis and Neural Networks. *Journal of Computational Physics*, page 109550, 2020.
16. Timothy Dozat. Incorporating Nesterov momentum into Adam. *Dostupné z: http://cs229.stanford.edu/proj2015/054_report.pdf*, 2016.
17. Oussama El B'charri, Rachid Latif, Khalifa Elmansouri, Abdenbi Abenaou, and Wissam Jenkal. Ecg signal performance de-noising assessment based on threshold tuning of dual-tree wavelet transform. *Biomedical engineering online*, 16(1):26, 2017.
18. Ana Ferrer, Rafael Sebastián, Damian Sanchez-Quintana, Jose F Rodriguez, Eduardo J Godoy, Laura Martinez, and Javier Saiz. Detailed anatomical and electrophysiological models of human atria and torso for the simulation of atrial activation. *PloS one*, 10(11):e0141573, 2015.
19. Piero Colli Franzone, Luca Franco Pavarino, and Simone Scacchi. *Mathematical cardiac electrophysiology*, volume 13. Springer, 2014.
20. Sophie Giffard-Roisin, Thomas Jackson, Lauren Fovargue, Jack Lee, Hervé Delingette, Reza Razavi, Nicholas Ayache, and Maxime Sermesant. Noninvasive personalization of a cardiac electrophysiology model from body surface potential mapping. *IEEE Transactions on Biomedical Engineering*, 64(9):2206–2218, 2016.
21. Erik Grandelius and Kenneth H Karlsen. The cardiac bidomain model and homogenization. *arXiv preprint arXiv:1811.07524*, 2018.
22. LS Green, Bruno Taccardi, PR Ershler, and RL Lux. Epicardial potential mapping. effects of conducting media on isopotential and isochrone distributions. *Circulation*, 84(6):2513–2521, 1991.
23. Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
24. Jan S Hesthaven and Stefano Ubbiali. Non-intrusive reduced order modeling of nonlinear problems using neural networks. *Journal of Computational Physics*, 363:55–78, 2018.
25. Corné Hoogendoorn, Nicolas Duchateau, Damian Sanchez-Quintana, Tristan Whitmarsh, Federico M Sukno, Mathieu De Craene, Karim Lekadir, and Alejandro F Frangi. A high-resolution atlas and statistical model of the human heart from multislice CT. *IEEE transactions on medical imaging*, 32(1):28–44, 2012.
26. Anselma Intini, Robert N Goldstein, Ping Jia, Charulatha Ramanathan, Kyungmoo Ryu, Bartolomeo Giannattasio, Robert Gilkeson, Bruce S Stambler, Pedro Brugada, William G Stevenson, et al. Electrocardiographic imaging (ECGI), a novel diagnostic modality used for mapping of focal left ventricular tachycardia in a young athlete. *Heart Rhythm*, 2(11):1250–1252, 2005.
27. Eugene M Izhikevich and Richard FitzHugh. FitzHugh-Nagumo model. *Scholarpedia*, 1(9):1349, 2006.
28. Mohammad Kachuee, Shayan Fazeli, and Majid Sarrafzadeh. ECG heartbeat classification: a deep transferable representation. In *2018 IEEE International Conference on Healthcare Informatics (ICHI)*, pages 443–444. IEEE, 2018.

29. Amel Karoui, Laura Bear, Pauline Migerditichan, and Nejib Zemzemi. Evaluation of fifteen algorithms for the resolution of the electrocardiography imaging inverse problem using ex-vivo and in-silico data. *Frontiers in Physiology*, 9:1708, 2018.
30. Gitta Kutyniok, Philipp Petersen, Mones Raslan, and Reinhold Schneider. A theoretical analysis of deep neural networks and parametric PDEs. *arXiv preprint arXiv:1904.00377*, 2019.
31. Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019.
32. Lu Lu, Yeonjong Shin, Yanhui Su, and George Em Karniadakis. Dying relu and initialization: Theory and numerical examples. *arXiv preprint arXiv:1903.06733*, 2019.
33. BJ Messinger-Rapport and Yoram Rudy. Noninvasive recovery of epicardial potentials in a realistic heart-torso geometry. Normal sinus rhythm. *Circulation research*, 66(4):1023–1039, 1990.
34. Justus Obergassel, Ilaria Cazzoli, Silvia Guarguagli, Venkat D. Nagarajan, and Sabine Ernst. Noninvasive 3d mapping and ablation of epicardial premature ventricular contractions from the endocardial aspect of the left atrial appendage. *JACC: Case Reports*, 2(11):1776–1780, 2020.
35. Howard S Oster, Bruno Taccardi, Robert L Lux, Philip R Ershler, and Yoram Rudy. Non-invasive electrocardiographic imaging: reconstruction of epicardial potentials, electrograms, and isochrones and localization of single and multiple electrocardiac events. *Circulation*, 96(3):1012–1024, 1997.
36. Lawrence Edward Payne. *Improperly posed problems in Partial Differential Equations*. SIAM, 1975.
37. Roberto Piersanti, Pasquale C Africa, Marco Fedele, Christian Vergara, Luca Dedè, Antonio F Corno, and Alfio Quarteroni. Modeling cardiac muscle fibers in ventricular and atrial electrophysiology simulations. *Computer Methods in Applied Mechanics and Engineering*, 373:113468, 2021.
38. Alfio Quarteroni, Andrea Manzoni, and Federico Negri. *Reduced Basis methods for Partial Differential Equations: an introduction*, volume 92. Springer, 2015.
39. Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
40. Charulatha Ramanathan, Raja N Ghanem, Ping Jia, Kyungmoo Ryu, and Yoram Rudy. Noninvasive electrocardiographic imaging for cardiac electrophysiology and arrhythmia. *Nature medicine*, 10(4):422–428, 2004.
41. Charulatha Ramanathan and Yoram Rudy. Electrocardiographic imaging: I. effect of torso inhomogeneities on body surface electrocardiographic potentials. *Journal of cardiovascular electrophysiology*, 12(2):229–240, 2001.
42. David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.

43. Reza Sameni, Mohammad B Shamsollahi, Christian Jutten, and Gari D Clifford. A non-linear bayesian filtering framework for ECG denoising. *IEEE Transactions on Biomedical Engineering*, 54(12):2172–2185, 2007.
44. Omid Sayadi and Mohammad Bagher Shamsollahi. ECG denoising and compression using a modified extended Kalman filter structure. *IEEE Transactions on Biomedical Engineering*, 55(9):2240–2248, 2008.
45. Joakim Sundnes, Glenn Terje Lines, Xing Cai, Bjørn Frederik Nielsen, Kent-Andre Mardal, and Aslak Tveito. *Computing the electrical activity in the heart*, volume 1. Springer Science & Business Media, 2007.
46. Nikolai Nikolaevich Tarkhanov. *The Cauchy problem for solutions of elliptic equations*, volume 7. Vch Pub, 1995.
47. Andrey N Tikhonov and Vasilii Y Arsenin. Solutions of ill-posed problems. *New York*, pages 1–30, 1977.
48. Leslie Tung. *A bi-domain model for describing ischemic myocardial DC potentials*. PhD thesis, Massachusetts Institute of Technology, 1978.
49. Dafang Wang, Robert M Kirby, and Chris R Johnson. Resolution strategies for the finite-element-based solution of the ECG inverse problem. *IEEE Transactions on biomedical engineering*, 57(2):220–237, 2009.
50. Liansheng Wang, Weiqi Wu, Yiping Chen, and Changhua Liu. An ADMM-net solution to inverse problem of electrocardiology. In *2018 5th International Conference on Systems and Informatics (ICSAI)*, pages 565–569. IEEE, 2018.
51. Qian Wang, Jan S Hesthaven, and Deep Ray. Non-intrusive reduced order modeling of unsteady flows using artificial neural networks with application to a combustion problem. *Journal of computational physics*, 384:289–307, 2019.
52. CHRISTOPHER R Wyndham, MOOIDEEN K Meeran, THOMAS Smith, AMARKANTH Saxena, RICHARD M Engelman, SIDNEY Levitsky, and KM Rosen. Epicardial activation of the intact human heart without conduction defect. *Circulation*, 59(1):161–168, 1979.
53. CR Wyndham, Thomas Smith, Mooideen K Meeran, ROBERT Mammana, SIDNEY Levitsky, and KM Rosen. Epicardial activation in patients with left bundle branch block. *Circulation*, 61(4):696–703, 1980.
54. Nejib Zemzemi, Cecile Dobrzynski, Laura Bear, Mark Potse, Corentin Dallet, Yves Coudière, Remi Dubois, and Josselin Duchateau. Effect of the torso conductivity heterogeneities on the ECGI inverse problem solution. In *2015 Computing in Cardiology Conference (CinC)*, pages 233–236. IEEE, 2015.
55. Martin Zihlmann, Dmytro Perekrstenko, and Michael Tschannen. Convolutional recurrent neural networks for electrocardiogram classification. In *2017 Computing in Cardiology (CinC)*, pages 1–4. IEEE, 2017.

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 24/2021** Regazzoni, F.; Chapelle, D.; Moireau, P.
Combining Data Assimilation and Machine Learning to build data-driven models for unknown long time dynamics - Applications in cardiovascular modeling
- 22/2021** Domanin, M.; Bennati, L.; Vergara, C.; Bissacco, D.; Malloggi, C.; Silani, V.; Parati, G.; Trimani, A.
Fluid structure interaction analysis to stratify the behavior of different atheromatous carotid plaques
- 23/2021** Scimone, R.; Taormina, T.; Colosimo, B. M.; Grasso, M.; Menafoglio, A.; Secchi, P.
Statistical modeling and monitoring of geometrical deviations in complex shapes with application to Additive Manufacturing
- 20/2021** Pasquale, A.; Ammar, A.; Falcó, A.; Perotto, S.; Cueto, E.; Duval, J.-L.; Chinesta, F.
A separated representation involving multiple time scales within the Proper Generalized Decomposition framework
- 21/2021** Torti, A.; Galvani, M.; Menafoglio, A.; Secchi, P.; Vantini, S.
A General Bi-clustering Algorithm for Hilbert Data: Analysis of the Lombardy Railway Service
- 16/2021** Salvador, M.; Dede', L.; Manzoni, A.
Non intrusive reduced order modeling of parametrized PDEs by kernel POD and neural networks
- 19/2021** Gillard, M.; Benacchio, T.
FT-GCR: a fault-tolerant generalized conjugate residual elliptic solver
- 18/2021** Gigante, G.; Vergara, C.
On the choice of interface parameters in Robin-Robin loosely coupled schemes for fluid-structure interaction
- 17/2021** Chew, R.; Benacchio, T.; Hastermann, G.; Klein, R.
Balanced data assimilation with a blended numerical model
- 13/2021** Ferro, N.; Perotto, S.; Cangiani, A.
An anisotropic recovery-based error estimator for adaptive discontinuous Galerkin methods