



MOX-Report No. 40/2020

Deep learning-based reduced order models in cardiac electrophysiology

Fresca, S.; Manzoni, A.; Dedè, L.; Quarteroni, A.

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

mox-dmat@polimi.it

<http://mox.polimi.it>

Deep learning-based reduced order models in cardiac electrophysiology

Stefania Fresca¹, Andrea Manzoni^{1*}, Luca Dedé¹, Alfio Quarteroni^{1,2},

1 MOX - Dipartimento di Matematica, Politecnico di Milano, Milano, Italy

2 Mathematics Institute, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

* andrea1.manzoni@polimi.it

Abstract

Predicting the electrical behavior of the heart, from the cellular scale to the tissue level, relies on the formulation and numerical approximation of coupled nonlinear dynamical systems. These systems describe the cardiac action potential, that is the polarization/depolarization cycle occurring at every heart beat that models the time evolution of the electrical potential across the cell membrane, as well as a set of ionic variables. Multiple solutions of these systems, corresponding to different model inputs, are required to evaluate outputs of clinical interest, such as activation maps and action potential duration. More importantly, these models feature coherent structures that propagate over time, such as wavefronts. These systems can hardly be reduced to lower dimensional problems by conventional reduced order models (ROMs) such as, e.g., the reduced basis (RB) method. This is primarily due to the low regularity of the solution manifold (with respect to the problem parameters) as well as to the nonlinear nature of the input-output maps that we intend to reconstruct numerically. To overcome this difficulty, in this paper we propose a new, nonlinear approach which exploits deep learning (DL) algorithms to obtain accurate and efficient ROMs, whose dimensionality matches the number of system parameters. Our DL approach combines deep feedforward neural networks (NNs) and convolutional autoencoders (AEs). We show that the proposed DL-ROM framework can efficiently provide solutions to parametrized electrophysiology problems, thus enabling multi-scenario analysis in pathological cases. We investigate three challenging test cases in cardiac electrophysiology and prove that DL-ROM outperforms classical projection-based ROMs.

Introduction

The electrical activation of the heart is the main responsible of its contraction, is the result of two processes: at the microscopic scale, the generation of ionic currents through the cellular membrane producing a local action potential; and at the macroscopic scale, the propagation of the action potential from cell to cell in the form of a transmembrane potential [1–3]. This latter process can be described by means of partial differential equations (PDEs), suitably coupled with systems of ordinary differential equations (ODEs) modeling the ionic currents in the cells.

Solving this system using a high-fidelity, full order model (FOM) such as, e.g., the finite element (FE) method, is computationally demanding. Indeed, the propagation of the electrical signal is characterized by the fast dynamics of very steep fronts, thus requiring very fine space and time discretizations. [3–5]. Using a FOM may quickly become unaffordable if such a coupled system must be solved for several values of

parameters representing either functional or geometric data such as, e.g., material properties, initial and boundary conditions, or the shape of the domain. Multi-query analysis is relevant in a variety of situations: when analysing multiple scenarios, when dealing with sensitivity analysis and uncertainty quantification (UQ) problems in order to account for inter-subject variability [6–8], for parameter estimation or data assimilation, in which some unknown (or unaccessible) quantities characterizing the mathematical model must be inferred from a set of measurements [9–13].

Conventional projection-based reduced order models (ROMs) built, e.g., through the reduced basis (RB) method [14], yields inefficient ROMs when dealing with nonlinear time-dependent parametrized PDE-ODE system as the one arising from cardiac electrophysiology. The three major computational bottlenecks shown by such kind of ROMs for cardiac electrophysiology are due the fact that:

- the linear superimposition of modes, on which they are based, would cause the dimension of the ROM to be excessively large to guarantee an acceptable accuracy;
- evaluating the ROM requires the solution of a dynamical system, which might be unstable unless the size of time step Δt is very small;
- the ROM must also account for the dynamics of the gating variables, even when aiming at computing just the electrical potential. This fact entails an extremely intrusive and costly hyper-reduction stage to reduce the solution of the ODE system to a few, selected mesh nodes [15].

To overcome the limitations of projection-based ROMs, we propose a new, non-intrusive ROM technique based on deep learning (DL) algorithms, which we refer to as DL-ROM. Combining in a suitable way a convolutional autoencoder (AE) and a deep feedforward neural network (DFNN), the DL-ROM technique enables the construction of an efficient ROM, whose dimension is as close as possible to the number of parameters upon which the solution of the differential problem depends. A preliminary numerical assessment of our DL-ROM technique has already been presented in [16], albeit on simpler – yet challenging – test cases.

The proposed DL-ROM technique is a combination of a data-driven with a physics based model approach. Indeed, it exploits snapshots taken from a set of FOM solutions (for selected parameter values and time instances) and deep neural network architectures to learn, in a non-intrusive way, both *(i)* the nonlinear trial manifold where the ROM solution is sought, and *(ii)* the nonlinear reduced dynamics. In a linear ROM built, e.g., through proper orthogonal decomposition (POD), the former quantity is nothing but a set of basis functions, while the latter task corresponds to the projection stage in the subspace spanned by these basis functions. Here, our goal is to show that DL-ROM can be effectively used to handle parametrized problems in cardiac electrophysiology, accounting for both physiological and pathological conditions, in order to provide fast and accurate solutions. The proposed DL-ROM is computationally efficient during the testing stage, that is for any new scenario unseen during the training stage. This is particularly useful in view of the evaluation of patient-specific features to enable the integration of computational methods in current clinical platforms.

DL techniques for parametrized PDEs have previously been proposed in other contexts. In [17–20] feedforward neural networks have been employed to model the reduced dynamics in a less intrusive way, that is, avoiding the costs entailed by projection-based ROMs, but still relying on a linear trial manifold built, e.g., through POD. In [21–23] the construction of ROMs for nonlinear, time-dependent problems has been replaced by the evaluation of ANN-based regression models. In [24, 25] the reduced trial manifold where the approximation is sought has been modeled through ANNs thus avoiding the linear superimposition of POD modes, on a minimum residual formulation to derive the

ROM [25], or without considering an explicit parameter dependence in the differential problem that is considered [24]. In all these works, coupled problems have never been considered. Moreover, very often DL techniques have been exploited to address problems which require only a moderate dimension of projection-based ROMs. We demonstrate that our DL-ROM provides accurate results by constructing ROMs with extremely low-dimension in prototypical test cases. These tests exhibit all the relevant physical features which make the numerical approximation of parametrized problems in cardiac electrophysiology a challenging task.

Materials and methods

Cardiac electrophysiology

Muscle contraction and relaxation drive the pump function of the heart. In particular, tissue contraction is triggered by electrical signals self-generated in the heart and propagated through the myocardium thanks to the excitability of the cardiac cells, the cardiomyocytes [3, 26]. When suitably stimulated, cardiomyocytes produce a variation of the potential across the cellular membrane, called *transmembrane potential*. Its evolution in time is usually referred to as *action potential*, involving a polarization and a depolarization in the early stage of every heart beat. The action potential is generated by several ion channels (e.g., calcium, sodium, potassium) that open and close, and by the resulting ionic currents crossing the membrane. For instance, coupling the so-called monodomain model for the transmembrane potential $u = u(\mathbf{x}, t)$ with a phenomenological model for the ionic currents – involving a single gating variable $w = w(\mathbf{x}, t)$ – in a domain Ω representing, e.g., a portion of the myocardium, results in the following nonlinear time-dependent system

$$\begin{cases} \frac{\partial u}{\partial t} - \operatorname{div}(\mathbf{D}\nabla u) + I_{ion}(u, w) = I_{app}(\mathbf{x}, t) & (\mathbf{x}, t) \in \Omega \times (0, T), \\ \frac{\partial w}{\partial t} + g(u, w) = 0 & (\mathbf{x}, t) \in \Omega \times (0, T), \\ \nabla u \cdot \mathbf{n} = 0 & (\mathbf{x}, t) \in \partial\Omega \times (0, T), \\ u(\mathbf{x}, 0) = 0, w(\mathbf{x}, 0) = 0 & \mathbf{x} \in \Omega. \end{cases} \quad (1)$$

Here t denotes a rescaled time¹, \mathbf{n} denotes the outward directed unit vector normal to the boundary $\partial\Omega$ of Ω , whereas I_{app} is an applied current representing, e.g., the initial activation of the tissue. The nonlinear diffusion-reaction equation for u is two-ways coupled with the ODE system, which must be in principle solved at any point $\mathbf{x} \in \Omega$; indeed, the reaction term I_{ion} and the function g depend on both u and w . The most common choices for the two functions I_{ion} and g in order to efficiently reproduce the action-potential are, e.g., the FitzHugh-Nagumo [28, 29], the Aliev-Panfilov [27, 30] or the Mitchell and Schaeffer models [31]. The diffusivity tensor \mathbf{D} usually depends on the fibers-sheet structure of the tissue, affecting directional conduction velocities and directions. In particular, by assuming an axisymmetric distribution of the fibers, the conductivity tensor takes the form

$$\mathbf{D}(\mathbf{x}) = \sigma_t \mathbf{I} + (\sigma_l - \sigma_t) \mathbf{f}_0 \otimes \mathbf{f}_0, \quad (2)$$

where σ_l and σ_t are the conductivities in the fibers and the transversal directions.

When a simple phenomenological ionic model is considered, such as the FitzHugh-Nagumo or the Aliev-Panfilov (A-P) model, the ionic current takes the form of a cubic

¹Dimensional times and potential [27] are given by $\tilde{t}[ms] = 12.9t$ and $\tilde{u}[mV] = 100u - 80$. The transmembrane potential ranges from the resting state of -80 mV to the excited state of $+20$ mV.

nonlinear function of u and a single (dimensionless) gating variable plays the role of a recovery function, allowing to model refractoriness of cells. In this paper, we focus on the Aliev-Panfilov model, which consists in taking

$$\begin{aligned} I_{ion}(u, w) &= Ku(u - a)(u - 1) + uw, \\ g(u, w) &= \left(\epsilon_0 + \frac{c_1 w}{c_2 + u} \right) (-w - Ku(u - b - 1)). \end{aligned} \quad (3)$$

The parameters $K, a, b, \epsilon_0, c_1, c_2$ are related to the cell. Here a represents an oscillation threshold, whereas the weighting factor $\epsilon_0 + \frac{c_1 w}{c_2 + u}$ was introduced in [27] to tune the restitution curve to experimental observations by adjusting the parameters c_1 and c_2 ; see, e.g., [1–3, 32] for a detailed review. In the remaining part of the paper, we denote by $\boldsymbol{\mu} \in \mathcal{P} \subset \mathbb{R}^{n_\mu}$ a parameter vector listing all the n_μ input parameters characterizing physical (and, possibly, geometrical) properties we might be interested to vary; \mathcal{P} is a subset of \mathbb{R}^{n_μ} , denoting the parameter space. Relevant physical situations are those in which input parameters affect the diffusivity matrix \mathbf{D} (through the conduction velocities) and the applied current I_{app} ; previous analyses focused instead on the gating variable dynamics (through g) and the ionic current I_{ion} , see [15].

Projection-based ROMs

From an algebraic standpoint, the spatial discretization of system (1) through the Galerkin-finite element (FE) approximation [33] yields the following nonlinear dynamical system for $\mathbf{u} = \mathbf{u}(t; \boldsymbol{\mu})$, $\mathbf{w} = \mathbf{w}(t; \boldsymbol{\mu})$, representing our full order model (FOM):

$$\begin{cases} \mathbf{M}(\boldsymbol{\mu}) \frac{\partial \mathbf{u}}{\partial t} = \mathbf{A}(\boldsymbol{\mu}) \mathbf{u} + \mathbf{I}_{ion}(t, \mathbf{u}, \mathbf{w}; \boldsymbol{\mu}) + \mathbf{I}_{app}(t; \boldsymbol{\mu}), & t \in (0, T), \\ \frac{\partial \mathbf{w}}{\partial t}(t; \boldsymbol{\mu}) = \mathbf{g}(t, \mathbf{u}, \mathbf{w}; \boldsymbol{\mu}), & t \in (0, T), \\ \mathbf{u}(0) = \mathbf{u}_0, \quad \mathbf{w}(0) = \mathbf{w}_0. \end{cases} \quad (4)$$

Here $\mathbf{A}(\boldsymbol{\mu}) \in \mathbb{R}^{N \times N}$ is a matrix arising from the diffusion operator (thus including the conductivity tensor $\mathbf{D}(\boldsymbol{\mu}) = \mathbf{D}(\mathbf{x}; \boldsymbol{\mu})$, which can vary within the myocardium due to fiber orientation and conditions, such as the possible presence of ischemic regions); $\mathbf{M}(\boldsymbol{\mu}) \in \mathbb{R}^{N \times N}$ is the mass matrix; $\mathbf{I}_{ion}, \mathbf{g} \in \mathbb{R}^N$ are vectors arising from the nonlinear terms; $\mathbf{I}_{app} \in \mathbb{R}^N$ is a vector collecting the applied currents; finally, $\mathbf{u}_0, \mathbf{w}_0 \in \mathbb{R}^N$ are the initial data, possibly depending on $\boldsymbol{\mu}$. The dimension N is related to the dimension of the FE space and, ultimately, depends on the size $h > 0$ of the computational mesh used to discretize the domain Ω . Note that the system of ODEs arises from the collocation of the ODE (1)₂ at the nodes used for the numerical integration.

The intrinsic dimension of the solution manifold

$$\mathcal{S} = \{ \mathbf{u}(t; \boldsymbol{\mu}) \mid t \in [0, T] \text{ and } \boldsymbol{\mu} \in \mathcal{P} \subset \mathbb{R}^{n_\mu} \} \subset \mathbb{R}^N, \quad (5)$$

obtained by solving (4) when $(t; \boldsymbol{\mu})$ varies in $[0, T] \times \mathcal{P}$, is usually much smaller than N and, under suitable conditions, is at most $n_\mu + 1 \ll N$, where n_μ is the number of parameters – in this respect, the time independent variable plays the role of a parameter. For this reason, ROMs attempt at approximating \mathcal{S} by introducing a suitable trial manifold of lower dimension. The most popular approach is *proper orthogonal decomposition* (POD), which exploits a *linear trial manifold* built through the singular value decomposition of a matrix $\mathbf{S} \in \mathbb{R}^{N \times N_s}$ collecting a set of FOM snapshots

$$\mathbf{S} = [\mathbf{u}(t^1; \boldsymbol{\mu}_1) \mid \dots \mid \mathbf{u}(t^{N_t}; \boldsymbol{\mu}_1) \mid \dots \mid \mathbf{u}(t^1; \boldsymbol{\mu}_{N_{train}}) \mid \dots \mid \mathbf{u}(t^{N_t}; \boldsymbol{\mu}_{N_{train}})] ;$$

this is a set of solutions obtained for N_{train} selected input parameters at (a subset, possibly, of) the time instants $\{t^k\}_{k=1}^{N_t}$ in which $(0, T)$ is partitioned for the sake of time discretization. The most common choice is to set $t^k = k\Delta t$ where $\Delta t = T/(N_t - 1)$.

When using a projection-based ROM, the approximation of $\mathbf{u}(t; \boldsymbol{\mu})$ is sought as a linear superimposition of modes, under the form

$$\mathbf{u}(t; \boldsymbol{\mu}) \approx \mathbf{V}\mathbf{u}_n(t; \boldsymbol{\mu}), \quad (6)$$

thus yielding a linear ROM, in which the columns of the matrix $\mathbf{V} = [\boldsymbol{\zeta}_1, \dots, \boldsymbol{\zeta}_n] \in \mathbb{R}^{N \times n}$ form an orthonormal basis of a space V_n , an n -dimensional subspace of \mathbb{R}^N . In the case of POD, V_n provides the best n -rank approximation of \mathbf{S} in the Frobenius norm, that is, $\boldsymbol{\zeta}_1, \dots, \boldsymbol{\zeta}_n$ are the first n (left) singular vectors of \mathbf{S} corresponding to the n largest singular values $\sigma_1, \dots, \sigma_n$ of \mathbf{S} , such that the projection error is smaller than a desired tolerance ε_{POD} . To meet this requirement, it is sufficient to choose n as the smallest integer such that

$$\frac{\sum_{i=1}^N \sigma_i^2}{\sum_{i=1}^{N_s} \sigma_i^2} > 1 - \varepsilon_{POD}^2,$$

i.e., the energy retained by the last $N_s - n$ POD modes is equal or smaller than ε_{POD}^2 . The approximation of \mathbf{w} is given instead by its restriction

$$\mathbf{w}(t; \boldsymbol{\mu}) \approx \mathbf{P}\mathbf{w}_m(t; \boldsymbol{\mu}),$$

to a (possibly, small) subset of m degrees of freedom, where $m \ll n$, at which the nonlinear term \mathbf{I}_{ion} is interpolated exploiting a problem-dependent basis, spanned by the columns of a matrix $\boldsymbol{\Phi} \in \mathbb{R}^{N \times m}$, which is built according to a suitable *hyper-reduction* strategy; see, e.g., [15] for further details. Here $\mathbf{P} = [\mathbf{e}_1, \dots, \mathbf{e}_m] \in \mathbb{R}^{N \times m}$ denotes a matrix formed by the columns of the $N \times N$ identity matrix corresponding to the m selected degrees of freedom.

A Galerkin-POD ROM for system (1) is then obtained by (i) first, substituting equation (6) into equation 4 and projecting it onto V_n ; then, (ii) solving the system of ODEs at m selected degrees of freedom, thus yielding the following nonlinear dynamical system for $\mathbf{u}_n = \mathbf{u}_n(t; \boldsymbol{\mu})$ and the selected components $\mathbf{P}^T \mathbf{w} = \mathbf{P}^T \mathbf{w}(t; \boldsymbol{\mu})$ of \mathbf{w} :

$$\begin{cases} \mathbf{V}^T \mathbf{M}(\boldsymbol{\mu}) \mathbf{V} \frac{\partial \mathbf{u}_n}{\partial t} + \mathbf{V}^T \mathbf{A}(\boldsymbol{\mu}) \mathbf{V}^T \mathbf{u}_n \\ \quad + \mathbf{V}^T \boldsymbol{\Phi} (\mathbf{P}^T \boldsymbol{\Phi})^{-1} \mathbf{I}_{ion}(t, \mathbf{P}^T \mathbf{V} \mathbf{u}_n, \mathbf{P}^T \mathbf{w}; \boldsymbol{\mu}) - \mathbf{V}^T \mathbf{I}_{app}(t; \boldsymbol{\mu}) = \mathbf{0}, & t \in (0, T), \\ \mathbf{P}^T \frac{\partial \mathbf{w}}{\partial t} + \mathbf{g}(t, \mathbf{P}^T \mathbf{V} \mathbf{u}_n, \mathbf{P}^T \mathbf{w}; \boldsymbol{\mu}) = \mathbf{0}, & t \in (0, T), \\ \mathbf{u}_n(0) = \mathbf{V}^T \mathbf{u}_0, \quad \mathbf{P}^T \mathbf{w}(0) = \mathbf{P}^T \mathbf{w}_0. \end{cases} \quad (7)$$

This strategy is the essence of the reduced basis (RB) method for nonlinear time-dependent parametrized PDEs. However, using (7) as an approximation to (4) is known to suffer from several problems. First of all, an extensive hyper-reduction stage (exploiting, e.g., the discrete empirical interpolation method (DEIM)) must be performed in order to be able to evaluate any $\boldsymbol{\mu}$ - or \mathbf{u} -dependent quantities appearing in (7), that is, without relying on N -dimensional arrays. Moreover, whenever the solution of the differential problem features coherent structures that propagate over time, such as steep wavefronts, the dimension n of the projection-based ROM (7) might easily become very large, due to the basic linearity assumption, by which the solution is given by a linear superimposition of POD modes, thus severely degrading the computational efficiency of the ROM. A possible way to overcome this bottleneck is to rely on local reduced bases, built through POD after the set of snapshots has been split into $N_c > 1$ clusters, according to suitable clustering (or unsupervised learning) algorithms [15].

Deep learning-based reduced order modeling (DL-ROM)

To overcome the limitations of linear ROMs, we consider a new, nonlinear ROM technique based on deep learning models. First introduced in [16] and assessed on one-dimensional benchmark problems, the DL-ROM technique aims at learning both the nonlinear trial manifold (corresponding to the matrix \mathbf{V} in the case of a linear ROM) in which we seek the solution to the parametrized system (1) and the nonlinear reduced dynamics (corresponding to the projection stage in a linear ROM). This method is not intrusive; it relies on DL algorithms trained on a set of FOM solutions obtained for different parameter values.

We denote by N_{train} and N_{test} the number of training and testing parameter instances, respectively; the ROM dimension is again denoted by $n \ll N$. In order to describe the system dynamics on a suitable reduced nonlinear trial manifold (a task which we refer to as *reduced dynamics learning*), the intrinsic coordinates of the ROM approximation are defined as

$$\mathbf{u}_n(t; \boldsymbol{\mu}, \boldsymbol{\theta}_{DF}) = \boldsymbol{\phi}_n^{DF}(t; \boldsymbol{\mu}, \boldsymbol{\theta}_{DF}), \quad (8)$$

where $\boldsymbol{\phi}_n^{DF}(\cdot; \cdot, \boldsymbol{\theta}_{DF}) : \mathbb{R}^{(n_\mu+1)} \rightarrow \mathbb{R}^n$ is a deep feedforward neural network (DFNN), consisting in the subsequent composition of a nonlinear activation function, applied to a linear transformation of the input, multiple times [34]. Here $\boldsymbol{\theta}_{DF}$ denotes the vector of parameters of the DFNN, collecting all the corresponding weights and biases of each layer of the DFNN.

Regarding instead the description of the reduced nonlinear trial manifold, approximating the solution one, $\tilde{\mathcal{S}} \approx \mathcal{S}$ (a task which we refer to as *reduced trial manifold learning*) we employ the *decoder function of a convolutional autoencoder*² (AE) [35, 36]. More precisely, $\tilde{\mathcal{S}}$ takes the form

$$\tilde{\mathcal{S}} = \{\mathbf{f}^D(\mathbf{u}_n(t; \boldsymbol{\mu}, \boldsymbol{\theta}_{DF}); \boldsymbol{\theta}_D) \mid \mathbf{u}_n(t; \boldsymbol{\mu}, \boldsymbol{\theta}_{DF}) \in \mathbb{R}^n, t \in [0, T] \text{ and } \boldsymbol{\mu} \in \mathcal{P} \subset \mathbb{R}^{n_\mu}\} \quad (10)$$

where $\mathbf{f}^D(\cdot; \boldsymbol{\theta}_D) : \mathbb{R}^n \rightarrow \mathbb{R}^N$ consists in the decoder function of a convolutional AE. This latter results from the composition of several layers (some of which are convolutional), depending upon a vector $\boldsymbol{\theta}_D$ collecting all the corresponding weights and biases.

As a matter of fact, the approximation $\tilde{\mathbf{u}}(t; \boldsymbol{\mu}) \approx \mathbf{u}(t; \boldsymbol{\mu})$ provided by the DL-ROM technique is defined as

$$\tilde{\mathbf{u}}(t; \boldsymbol{\mu}, \boldsymbol{\theta}_{DF}, \boldsymbol{\theta}_D) = \mathbf{f}^D(\boldsymbol{\phi}_n^{DF}(t; \boldsymbol{\mu}, \boldsymbol{\theta}_{DF}); \boldsymbol{\theta}_D). \quad (11)$$

The encoder function of the convolutional AE can then be exploited to map the FOM solution associated to $(t, \boldsymbol{\mu})$ onto a low-dimensional representation

$$\tilde{\mathbf{u}}_n(t; \boldsymbol{\mu}, \boldsymbol{\theta}_E) = \mathbf{f}_n^E(\mathbf{u}(t; \boldsymbol{\mu}); \boldsymbol{\theta}_E); \quad (12)$$

$\mathbf{f}_n^E(\cdot; \boldsymbol{\theta}_E) : \mathbb{R}^N \rightarrow \mathbb{R}^n$ denotes the encoder function, depending upon a vector $\boldsymbol{\theta}_E$ of parameters.

Computing the DL-ROM approximation of $\mathbf{u}(t; \boldsymbol{\mu}_{test})$, for any possible $t \in (0, T)$ and $\boldsymbol{\mu}_{test} \in \mathcal{P}$, corresponds to the testing stage of a DFNN and of the decoder function of

²The AE is a particular type of neural network aiming at learning the identity function

$$\mathbf{f}^{AE}(\cdot; \boldsymbol{\theta}_E, \boldsymbol{\theta}_D) : \mathbf{x} \mapsto \tilde{\mathbf{x}} \quad \text{with} \quad \tilde{\mathbf{x}} \simeq \mathbf{x}. \quad (9)$$

. It is composed by two main parts:

- the *encoder* function $\mathbf{f}_n^E(\cdot; \boldsymbol{\theta}_E) : \mathbf{x} \mapsto \tilde{\mathbf{x}}_n = \mathbf{f}_n^E(\mathbf{x}; \boldsymbol{\theta}_E)$, where $\mathbf{f}_n^E(\cdot; \boldsymbol{\theta}_E) : \mathbb{R}^N \rightarrow \mathbb{R}^n$ and $n \ll N$, mapping the high-dimensional input \mathbf{x} onto a low-dimensional code $\tilde{\mathbf{x}}_n$;
- the *decoder* function $\mathbf{f}^D(\cdot; \boldsymbol{\theta}_D) : \tilde{\mathbf{x}}_n \mapsto \tilde{\mathbf{x}} = \mathbf{f}^D(\tilde{\mathbf{x}}_n; \boldsymbol{\theta}_D)$, where $\mathbf{f}^D(\cdot; \boldsymbol{\theta}_D) : \mathbb{R}^n \rightarrow \mathbb{R}^N$, mapping the low-dimensional code $\tilde{\mathbf{x}}_n$ to an approximation of the original high-dimensional input $\tilde{\mathbf{x}}$.

a convolutional AE; this does not require the evaluation of the encoder function. We remark that our DL-ROM strategy overcomes the three major computational bottlenecks implied by the use of projection-based ROMs, since:

- the dimension of the DL-ROM can be kept extremely small;
- the time resolution required by the DL-ROM can be chosen to be larger than the one required by the numerical solution of dynamical systems in cardiac electrophysiology;
- the DL-ROM can be queried at any desired time instant, without requiring the solution of a dynamical system until that time;
- the DL-ROM does not require to account for the dynamics of the gating variables, thus avoiding any hyper-reduction stage. This advantage, already visible when employing a single gating variable as in the test cases addressed later in this paper, might become even more effective when dealing with more realistic ionic models (the so-called *I* and *II generation models*), when dozens of additional variables in the system of ODEs must be accounted for [3].

The training stage consists in solving the following optimization problem, in the variable $\boldsymbol{\theta} = (\boldsymbol{\theta}_E, \boldsymbol{\theta}_{DF}, \boldsymbol{\theta}_D)$, after the snapshot matrix \mathbf{S} has been formed:

$$\min_{\boldsymbol{\theta}} \mathcal{J}(\boldsymbol{\theta}) = \min_{\boldsymbol{\theta}} \frac{1}{N_s} \sum_{i=1}^{N_{train}} \sum_{k=1}^{N_t} \mathcal{L}(t^k, \boldsymbol{\mu}_i; \boldsymbol{\theta}), \quad (13)$$

where $N_s = N_{train}N_t$ and

$$\begin{aligned} \mathcal{L}(t^k, \boldsymbol{\mu}_i; \boldsymbol{\theta}) &= \frac{\omega_h}{2} \|\mathbf{u}(t^k; \boldsymbol{\mu}_i) - \tilde{\mathbf{u}}(t^k; \boldsymbol{\mu}_i, \boldsymbol{\theta}_{DF}, \boldsymbol{\theta}_D)\|^2 \\ &+ \frac{1 - \omega_h}{2} \|\tilde{\mathbf{u}}_n(t^k; \boldsymbol{\mu}_i, \boldsymbol{\theta}_E) - \mathbf{u}_n(t^k; \boldsymbol{\mu}_i, \boldsymbol{\theta}_{DF})\|^2, \end{aligned} \quad (14)$$

with $\omega_h \in [0, 1]$. The *per-example* loss function (14) combines the reconstruction error (that is, the error between the FOM solution and the DL-ROM approximation) and the error between the intrinsic coordinates and the output of the encoder.

The architecture of DL-ROM is the one shown in Fig 1. The encoder function is used only during the training and validation steps; it is instead discarded during the testing phase. See [16] for further algorithmic details about the training and the testing algorithms required to build and evaluate a DL-ROM.

We highlight that the DL-ROM technique does not require to solve a (reduced) nonlinear dynamical system for the reduced degrees of freedom as in (7); rather, it evaluates a nonlinear map for any given couple $(t, \boldsymbol{\mu}_{test})$, for each $t \in (0, T)$. Numerical results are extremely accurate, the mean relative error is indeed below 1% (see, e.g., Test 2), even if the causality intrinsic to the parabolic nature of the diffusion-reaction equation providing the monodomain model is broken when computing the DL-ROM approximation. Moreover, the map features an extremely low dimension, in the most favorable scenario equal to $n_{\boldsymbol{\mu}} + 1$. From a computational perspective, remarkable gains and simplifications can be obtained against a linear ROM, since (i) no hyper-reduction is required to enhance the evaluation of any $\boldsymbol{\mu}$ - or \mathbf{u} -dependent quantity, and (ii) even more interestingly, there is no need to evaluate the dynamics of the recovery variable w if one is only interested in the electrical potential.




Fig 1. DL-ROM architecture. DL-ROM architecture used during the training phase. In the red box, the DL-ROM to be queried for any new selected couple $(t, \boldsymbol{\mu})$ during the testing phase. The FOM solution $\mathbf{u}(t; \boldsymbol{\mu})$ is provided as input to block (A) which outputs $\tilde{\mathbf{u}}_n(t; \boldsymbol{\mu})$. The same parameter instance associated to the FOM, i.e. $(t, \boldsymbol{\mu})$, enters block (B) which provides as output $\mathbf{u}_n(t; \boldsymbol{\mu})$ and the error between the low-dimensional vectors (dashed green box) is accumulated. The intrinsic coordinates $\mathbf{u}_n(t; \boldsymbol{\mu})$ are given as input to block (C) returning the ROM approximation $\tilde{\mathbf{u}}(t; \boldsymbol{\mu})$. Then the reconstruction error (dashed black box) is computed.

Results and discussion

We now assess the computational performances of the proposed DL-ROM strategy on three relevant test cases in cardiac electrophysiology. Our choice of the numerical tests is aimed at highlighting the performance of our DL-ROM method in challenging electrophysiology problems, namely pathological cases in portion of cardiac tissues or physiological scenarios on realistic left ventricle geometries.

The architecture used to perform all the numerical tests is the one reported in the SI Appendix. To solve the optimization problem (13)-(14) we use the ADAM algorithm [37] with a starting learning rate equal to $\eta = 10^{-4}$. Moreover, we perform cross-validation by splitting the data in training and validation and following a proportion 8:2 and we implement an early-stopping regularization technique to reduce overfitting [34].

To evaluate the performance of the DL-ROM, we use the loss function (14) and on an error indicator defined as

$$\epsilon_{rel} = \frac{1}{N_{test}} \sum_{i=1}^{N_{test}} \left(\frac{\sqrt{\sum_{k=1}^{N_t} \|\mathbf{u}^k(\boldsymbol{\mu}_{test,i}) - \tilde{\mathbf{u}}^k(\boldsymbol{\mu}_{test,i})\|^2}}{\sqrt{\sum_{k=1}^{N_t} \|\mathbf{u}^k(\boldsymbol{\mu}_{test,i})\|^2}} \right). \quad (15)$$

Neural networks required by our DL-ROM technique have been implemented by means of the **Tensorflow** deep learning framework [38]; numerical simulations have been carried out on a workstation equipped with an Nvidia GeForce GTX 1070 8 GB GPU.

Test 1: Two-dimensional slab with ischemic region

We consider the computation of the transmembrane potential in a square slab $\Omega = (0, 10 \text{ cm})^2$ of cardiac tissue in presence of an ischemic (non-conductive) region. The ischemic region may act as anatomical driver of cardiac arrhythmias like tachycardias and fibrillations. The system we want to solve is a slight modification of equations (1), accounting for the presence of a non-conductive region which affects both the conductivity tensor and the ionic current term. The ischemic portion of the domain is modeled by replacing the conductivity tensor $\mathbf{D}(\mathbf{x})$, defined in (2), with $\bar{\mathbf{D}}(\mathbf{x}; \boldsymbol{\mu}) = \sigma(\mathbf{x}, \boldsymbol{\mu})\mathbf{D}(\mathbf{x})$, where the function $\sigma(\mathbf{x}, \boldsymbol{\mu})$ is given by

$$\begin{aligned}\sigma(\mathbf{x}; \boldsymbol{\mu}) &= \rho(\mathbf{x}; \boldsymbol{\mu}) + \sigma_0(1 - \rho(\mathbf{x}; \boldsymbol{\mu})), \\ \rho(\mathbf{x}; \boldsymbol{\mu}) &= 1 - \exp\left(-\frac{(x_1 - \mu_1)^4 + (x_2 - \mu_2)^4}{2\alpha^2}\right).\end{aligned}\quad (16)$$

In this case, $n_\mu = 2$ parameters are considered, representing the coordinates of the center of the scar, belong to the parameter space $\mathcal{P} = [3.5, 6.5 \text{ cm}]^2$. Moreover, $\alpha = 7 \text{ cm}^2$, $\sigma_0 = 10^{-4}$, the transversal and longitudinal conductivities are $\sigma_t = 12.9 \cdot 0.1 \text{ cm}^2/\text{ms}$ and $\sigma_l = 12.9 \cdot 0.2 \text{ cm}^2/\text{ms}$, respectively, and $\mathbf{f}_0 = (1, 0)^T$, meaning that the tissue fibers are parallel to the x -axis. Similarly, the ionic current $I_{ion}(u, w)$ in (1) is replaced by $\bar{I}_{ion}(u, w; \boldsymbol{\mu}) = \rho(\mathbf{x}; \boldsymbol{\mu})I_{ion}(u, w)$. The applied current takes the form

$$I_{app}(\mathbf{x}, t) = C \exp\left(-\frac{\|\mathbf{x}\|^2}{\beta}\right) \mathbf{1}_{[0, \bar{t}]}(t),$$

where $C = 100 \text{ mA}$, $\beta = 0.02 \text{ cm}^2$ and $\bar{t} = 2 \text{ ms}$. The parameters appearing in (3) are set to $K = 8$, $a = 0.01$, $b = 0.15$, $\varepsilon_0 = 0.002$, $c_1 = 0.2$, and $c_2 = 0.3$, see [39]. The equations have been discretized in space through linear finite elements by considering $N = 64 \times 64 = 4096$ grid points. For the time discretization and the treatment of nonlinear terms, we use a one-step, semi-implicit, first order scheme (see [15] for further details) by considering a time step $\Delta t = 0.1/12.9$ over $(0, T)$ with $T = 400 \text{ ms}$.

For the training phase, we uniformly sample $N_t = 1000$ time instances over $(0, T)$ and consider $N_{train} = 49$ training-parameter instances, with $\boldsymbol{\mu}_{train} = (3.5 + i0.5, 3.5 + j0.5)$, $i, j = 0, \dots, 6$. The maximum number of epochs is set equal to $N_{epochs} = 10000$, the batch size is $N_b = 40$ and, regarding the early-stopping criterion, we stop the training if the loss function does not decrease in 500 epochs. For the testing phase, $N_{test} = 36$ testing-parameter instances $\boldsymbol{\mu}_{test} = (3.75 + i0.5, 3.75 + j0.5)$, $i, j = 0, \dots, 5$, have been considered.

In Figs 2 and 3 we show the FOM and the DL-ROM solutions, the latter obtained with $n = 3$ for the testing-parameter instance $\boldsymbol{\mu}_{test} = (6.25, 6.25) \text{ cm}$ at $\bar{t} = 100$ and 356 ms, respectively, together with the relative error $\boldsymbol{\epsilon}_k \in \mathbb{R}^N$, for $k = 1, \dots, N_t$, defined as

$$\boldsymbol{\epsilon}_k = \frac{|\mathbf{u}^k(\boldsymbol{\mu}_{test}) - \tilde{\mathbf{u}}^k(\boldsymbol{\mu}_{test})|}{\sqrt{\frac{1}{N_t} \sum_{k=1}^{N_t} \|\mathbf{u}^k(\boldsymbol{\mu}_{test})\|^2}}.\quad (17)$$

While (15) is a synthetic indicator, the quantity defined in (17) is instead a function of the space independent variable. In Fig 2 the tissue is depolarized except for the region occupied by scar and surrounding it, which is clearly characterized by a slower conduction. In Fig 3 the tissue is starting to repolarize and even if the shape of the ischemic region is not sharply reproduced, the DL-ROM solution is able to capture the diseased (non-conductive) nature of this portion of tissue.

In Fig 5 we show the action potentials (APs) computed at the six points P_1, \dots, P_6 reported in Fig 4. The DL-ROM is able to provide an accurate reconstruction of the




Fig 2. Test 1: comparison between FOM and DL-ROM solutions for a testing-parameter instance. FOM solution (left), DL-ROM solution with $n = 3$ (center) and relative error ϵ_k (right) for the testing-parameter instance $\boldsymbol{\mu}_{test} = (6.25, 6.25)$ cm at $\tilde{t} = 100$ ms. The maximum of the relative error ϵ_k is 10^{-3} and it is associated to the diseased tissue.




Fig 3. Test 1: comparison between FOM and DL-ROM solutions for a testing-parameter instance. FOM solution (left), DL-ROM solution with $n = 3$ (center) and relative error ϵ_k (right) for the testing-parameter instance $\boldsymbol{\mu}_{test} = (6.25, 6.25)$ cm at $\tilde{t} = 356$ ms. The maximum of the relative error ϵ_k is 10^{-3} and it is associated to the diseased tissue.

AP at almost all points; the maximum error is associated to the point P_3 , the closest one to the center of the scar, for $\tilde{t} \geq 200$ ms. However, even in this case, the DL-ROM technique is able to capture the difference, in terms of AP values, between the diseased and the healthy tissue.




Fig 4. Test 1: location of points P_i . FOM solution evaluated for $\boldsymbol{\mu}_{test} = (6.25, 6.25)$ cm at $\tilde{t} = 400$ ms together with the points P_1, \dots, P_6 .

The AP variability across the parameter space characterizing both the FOM and the DL-ROM solutions is shown in Fig 6. Still with a DL-ROM dimension $n = 3$, we report the APs for $\boldsymbol{\mu}_{test} = (\mu_{test}, \mu_{test})$ cm, with $\mu_{test} = 3.75, 4.25, 4.75, 5.25, 5.75, 6.25$, evaluated at $P = (7.46, 6.51)$ cm. The DL-ROM is able to capture such variability over \mathcal{P} ; moreover, the larger $\boldsymbol{\mu}_{test}$, the smaller the distance between the point P and the scar,




Fig 5. Test 1: comparison between the FOM and DL-ROM APs at P_i . APs evaluated for $\mu_{test} = (6.25, 6.25)$ cm at points P_1, \dots, P_6 . The DL-ROM, with $n = 3$, is able to sharply reconstruct the AP in almost all the points and the main features are captured also for the points close to the scar.

with their proximity impacting on the shape and the values of the AP. In particular, for $\mu_{test} = 6.25$, the point P falls into the *grey zone*.




Fig 6. Test 1: variability of the FOM and DL-ROM solutions over the parameter space. FOM (right) and DL-ROM (left) AP variability over \mathcal{P} at $P = (7.46, 6.51)$ cm. The DL-ROM sharply reconstructs the FOM variability over \mathcal{P} .

By using the DL-ROM technique and setting the dimension of the nonlinear trial manifold equal to the dimension of the solution manifold, i.e. $n = 3$, we obtain an error indicator (15) of $\epsilon_{rel} = 2.01 \cdot 10^{-2}$. In order to assess the computational efficiency of DL-ROM, we compare it with the POD-Galerkin ROM relying on N_c local reduced bases; we report in Table 1 the maximum and minimum number of basis functions, among all the clusters, required by the POD-Galerkin ROM [14, 15] to achieve the same accuracy.

In Fig 7 we compare the CPU time required to solve the FOM (through linear finite elements) over the time interval $(0, T)$, with the one needed by DL-ROM with $n = 3$, and the POD-Galerkin ROM with $N_c = 6$ local reduced bases, at testing time, by varying the FOM dimension N . Here, with testing time we refer, both for the DL-ROM and the POD-Galerkin ROM, to the time needed to query the ROM over the whole interval $(0, T)$,

Table 1. Test 1: dimensions of the POD-Galerkin ROM linear trial manifolds by varying the number of clusters.

$N_c = 1$	$N_c = 2$	$N_c = 4$	$N_c = 6$
250	219	200	193
	107	35	26

Maximum and minimum dimensions of the local reduced bases (that is, linear trial manifolds) built by the POD-Galerkin ROM for different numbers N_c of clusters.

by using for each technique the proper time resolution, for a given testing-parameter instance. Since the DL-ROM solution can be queried at a given time without requiring any solution of a dynamical system to recover the former time instances, the DL-ROM can employ larger time windows compared to the time steps required by the solution of the FOM and POD-Galerkin ROM dynamical systems for the cases at hand. This fact also has a positive impact on the data used during the training phase³. The speed-up obtained, for each value of N considered, is reported in Table 1. Both the DL-ROM and the POD-Galerkin ROM allow us to decrease the computational costs associated to the computation of the FOM solution for a testing-parameter instance. However, for a desired level of accuracy, CPU times required by the POD-Galerkin ROM during the testing phase are remarkably higher than the ones required by a DL-ROM with $n = 3$.




Fig 7. Test 1: FOM, DL-ROM and POD-Galerkin ROM CPU computational times. CPU time required to solve the FOM, by DL-ROM at testing time with $n = 3$ and by the POD-Galerkin ROM at testing time with $N_c = 6$ vs. N . The DL-ROM provides the smallest testing computational time for each N considered.

Both the DL-ROM and the POD-Galerkin ROM depend on the FOM dimension N . In the case of DL-ROM, the dependency on N at testing time, for a fixed value of Δt , is due to the presence of the decoder function; indeed, the process of learning the reduced dynamics (and so the dimension of the nonlinear trial manifold) does not depend on the FOM dimension. On the other hand, the dependence of the POD-Galerkin ROM on the FOM dimension also impacts on the dimension of the local linear trial manifolds: in general, by increasing N the dimension of each local linear subspace also increases. Referring to Fig 7 and Table 2, the CPU time required by the DL-ROM at

³Indeed, in order to build the snapshot matrix, we uniformly sample N_t time instances of the FOM solution over $T/\Delta t = 4000$ time steps; for each training parameter instance, only 25% of 4000 snapshots are retained from the FOM solution in the DL-ROM case, against 4000 snapshots in the POD-Galerkin ROM case.

testing time scales linearly with N , instead the one required by the POD-Galerkin ROM scales linearly with \sqrt{N} . In particular, even for the larger FOM dimension considered ($N = 16384$ for this test case), our DL-ROM is 19 times faster than the POD-Galerkin ROM. We are not able to run simulations for $N > 16384$, because of the limitation of the computing resources we have at our disposal. Despite the trend in Fig 7 is apparently not favorable for the DL-ROM technique, practice indicates that the CPU time for DL-ROM is smaller than the one for the POD-Galerkin ROM for *small* values of N , in other words only with very large values of N the POD-Galerkin ROM outperforms the DL-ROM strategy. Indeed, a linear fitting of the DL-ROM and the POD-Galerkin ROM CPU times⁴ in Fig 7 highlights that for $N = 65536$ and $N = 262144$, DL-ROM could be almost 10 and 5 times, respectively, faster than the POD-Galerkin ROM for the same values of N . Note that the results of this section have been obtained by employing the DL-ROM on a single CPU, an architecture which is not favorable to neural networks⁵. Further improvements are expected when employing our DL-ROM on a GPU for a given testing-parameter instance.

Table 2. Test 1: DL-ROM and POD-Galerkin ROM vs. FOM speed-up.

	$N = 256$	$N = 1024$	$N = 4096$	$N = 16384$
FOM vs. DL-ROM	472	536	539	412
FOM vs. POD-Galerkin ROM	3	6	12	22

DL-ROM and POD-Galerkin ROM vs. FOM speed-up by varying N . The DL-ROM speed-up is remarkably higher than the one obtained by using the POD-Galerkin ROM.

In Figs 8 and 9 we show the feature maps of the first convolutional layer of the encoder function $\sigma_1(W_1^k * \mathbf{u}^1(\boldsymbol{\mu}_{test}) + b_1^k)$, for $k = 1, \dots, 8$, in the DL-ROM neural network when the FOM solution for the testing-parameter instances $\boldsymbol{\mu}_{test} = (3.75, 3.75)$ cm and $\boldsymbol{\mu}_{test} = (6.25, 6.25)$ cm at $t = 0.2$ ms, are provided as inputs. At this stage, the feature maps retain most of the information present in the FOM solution. Moreover, by considering the two testing-parameter instances, we observe the translation equi-variance property [34] that convolutional layers hold when applied to the part of cardiac tissue corresponding to the scar. Moving to deeper layers, feature maps become increasingly abstract, and less visually interpretable; however, the extracted high-level features are still related both to the ischemic region and the electrical activation pattern.

We highlight that since the DL-ROM solution can be evaluated at any desired time instance without *solving* any dynamical system, the resulting computational time entailed by the DL-ROM at testing time are drastically reduced compared to the ones required by the FOM or the POD-Galerkin ROM to compute solutions at a particular time instance. In Fig 10 we show the DL-ROM, FOM and POD-Galerkin ROM CPU time needed to compute the approximated solution at \bar{t} , for $\bar{t} = 1, 10, 100$ and 400 ms averaged over the testing set and with $N = 4096$. We perform the training phase of the POD-Galerkin ROM over the original time interval $(0, T)$ ms and we report the results for $N_c = 6$, the number of clusters for which the smallest computational time is obtained. The DL-ROM CPU time to compute $\hat{\mathbf{u}}(\bar{t}; \boldsymbol{\mu}_{test})$ does not vary over \bar{t} and, by choosing $\bar{t} = T$, the DL-ROM speed-ups are equal to 7.3×10^4 and 6.5×10^3 with respect to the FOM and the POD-Galerkin ROM, with $N_c = 6$, computational times.

⁴ $N = 65536$ and $N = 262144$ for this test case represent FOM dimensions corresponding to mesh sizes h needed to solve, by means of linear finite elements, the problem on a 3D slab geometry both for physiological and pathological electrophysiology in the case a ten Tusscher-Panfilov ionic model [40] is used. This latter would indeed require smaller values of h compared to the Aliev-Panfilov model, due to the shape of the AP. See, e.g., [41, 42] for further details.

⁵Indeed, all tests are performed on a node (20 Intel[®] Xeon[®] E5-2640 v4 2.4GHz cores), using 5 cores, of our in-house HPC cluster.




Fig 8. Test 1: activations of the first convolutional layer of the encoder function for a testing-parameter instance. Feature maps of the first convolutional layer of the encoder function in the DL-ROM neural network for the testing-parameter instance $\boldsymbol{\mu}_{test} = (3.75, 3.75)$ cm at $\tilde{t} = 0.2$ ms.




Fig 9. Test 1: activations of the first convolutional layer of the encoder function for a testing-parameter instance. Activations of the first convolutional layer of the encoder function in the DL-ROM neural network for the testing-parameter instance $\boldsymbol{\mu}_{test} = (6.25, 6.25)$ cm at $\tilde{t} = 0.2$ ms.

Test 2: Two-dimensional slab with figure of eight re-entry

The most recognized cellular mechanisms sustaining atrial tachycardia is re-entry [43]. The particular kind of re-entry we deal with in this test case is called *figure of eight* re-entry, and can be obtained by solving equations (1). To induce the re-entry, we apply a classical S1-S2 protocol [3,44]. In particular, we consider a square slab of cardiac tissue $\Omega = (0, 2 \text{ cm})^2$ and apply an initial stimulus at the bottom edge of the domain, i.e.

$$I_{app}^1(\mathbf{x}, t) = \mathbf{1}_{\Omega_1}(\mathbf{x}) \mathbf{1}_{[t_1^i, t_1^f]}(\tilde{t}), \quad (18)$$

where $\Omega_1 = \{\mathbf{x} \in \Omega : y \leq 0.1\}$, $t_1^i = 0$ ms and $t_1^f = 5$ ms. A second stimulus under the form

$$I_{app}^2(\mathbf{x}, t; \mu) = \mathbf{1}_{\Omega_2(\mu)}(\mathbf{x}) \mathbf{1}_{[t_2^i, t_2^f]}(\tilde{t}), \quad (19)$$

with $\Omega_2(\mu) = \{\mathbf{x} \in \Omega : (x - 1)^2 + (y - \mu)^2 \leq (0.2)^2\}$, $t_2^i = 70$ ms and $t_2^f = 75$ ms, is then applied. The parameter μ , consisting in the y -coordinate of the center of the second




Fig 10. Test 1: FOM, POD-Galerkin ROM and DL-ROM CPU computational times. FOM, POD-Galerkin ROM and DL-ROM CPU computational times to compute $\tilde{\mathbf{u}}(\bar{t}; \boldsymbol{\mu}_{test})$ vs. \bar{t} averaged over the testing set. Thanks to the fact that the DL-ROM can be queried at any time instance it is extremely efficient in computing $\tilde{\mathbf{u}}(\bar{t}; \boldsymbol{\mu}_{test})$ with respect to both the FOM and the POD-Galerkin ROM.

circular stimulus, ranges in the parameter space $\mathcal{P} = [0.8, 1.1]$ cm. This choice has been made to obtain a re-entry elicited and sustained until $T = 175$ ms.

We restrict our study to the time interval [95, 175] ms, i.e. we do not consider the first time instances in which the re-entry has not arisen yet, being them equal over \mathcal{P} . The time step is $\Delta t = 0.2/12.9$. We consider $N = 256 \times 256 = 65536$ grid points, implying a mesh size $h = 0.0784$ mm; this mesh size is recognized to correctly solve the tiny transition front developing during depolarization of the tissue, as highlighted in [41, 42]. The fibers are parallel to the x -axis and the conductivities in the longitudinal and transversal directions to the fibers are $\sigma_l = 2 \times 10^{-3}$ cm²/ms and $\sigma_t = 3.1 \times 10^{-4}$ cm²/ms, respectively. The parameters appearing in (3) are set to $K = 8$, $a = 0.1$, $b = 0.1$, $\varepsilon_0 = 0.01$, $c_1 = 0.14$, and $c_2 = 0.3$, see [45].

The snapshot matrix is built by solving problem (1), completed with the applied currents 18 and 19, by means of linear finite elements and a semi-implicit scheme, over $N_t = 400$ time instances. Moreover, we consider $N_{train} = 13$ training-parameter instances uniformly distributed in the parameter space and $N_{test} = 12$ testing-parameter instances, each of them corresponding to the midpoint of two consecutive training-parameter instances. The maximum number of epochs is set equal to $N_{epochs} = 6000$, the batch size is $N_b = 3$, due to the high GPU memory occupation of each sample. Regarding the early-stopping criterion, we stop the training if the loss does not decrease in 1000 epochs.

In Fig 11 we show the FOM solution and the DL-ROM one obtained by setting the reduced dimension to $n = 5$, for the testing-parameter instance $\mu_{test} = 0.9625$ cm, at $\tilde{t} = 141.2$ ms and $\tilde{t} = 157.2$ ms, together with the relative error $\boldsymbol{\epsilon}_k^s \in \mathbb{R}^N$, for $k = 1, \dots, N_t$, defined as

$$\boldsymbol{\epsilon}_k^s = \frac{|\mathbf{u}^k(\mu_{test}) - \tilde{\mathbf{u}}^k(\mu_{test})|}{\|\mathbf{u}^k(\mu_{test})\|_1}. \quad (20)$$

The trend of the relative error (20) over time, for the selected testing-parameter instance $\mu_{test} = 0.9625$ cm, is depicted in Fig 12; we highlight that the error is always smaller than 1%. In particular, in Fig 12 we show the mean (over the domain), the median, and the first and third quartile of the relative error, as well as its minimum. The interquartile range (IQR) shows that the distribution of the error is almost uniform




Fig 11. Test 2: comparison between FOM and DL-ROM solutions for a testing-parameter instance. FOM solution (left), DL-ROM one (center) with $n = 5$, and relative error ϵ_k^s (right) at $\tilde{t} = 141.2$ ms (top) and $\tilde{t} = 157.2$ ms (bottom), for the testing-parameter instance $\mu_{test} = 0.9625$ cm. The relative error ϵ_k^s is below the 2% both for $\tilde{t} = 141.2$ ms and $\tilde{t} = 157.2$ ms, the maximum value of the error being associated to very few points of the domain.

over time, and that the maximum error is associated to the first time instant – this latter being the time instant at which the solution is most different over \mathcal{P} .




Fig 12. Test 2: trend of the relative error over time. Relative error ϵ_k^s vs. \tilde{t} with $n = 5$ for the testing-parameter instance $\mu_{test} = 0.9625$ cm (the red band indicates the IQR). The distribution of the error maintains uniform over time.

In Fig 13 we show the FOM and the DL-ROM solutions, the latter obtained by setting $n = 5$, for the last time instance, i.e. at $\tilde{t} = 175$ ms, for $\mu_{test} = 0.8125$ cm and $\mu_{test} = 1.0625$ cm, in order to point out the variability of the solution over the parameter space and the ability of DL-ROM to capture it.

We now compare the computational times required by the FOM, the POD-Galerkin ROM (for different values of N_c) and the DL-ROM, keeping for all the same degree of accuracy achieved by DL-ROM, i.e. $\epsilon_{rel} = 7.87 \times 10^{-3}$, and running the code on




Fig 13. Test 2: comparison between FOM and DL-ROM solutions for different testing-parameter instances. FOM solution (left), DL-ROM one (center) with $n = 5$, and relative error ϵ_k^s (right) at $\tilde{t} = 175$ ms, for the testing-parameter instance $\mu = 0.8125$ cm (top) and $\mu = 1.0625$ cm (bottom). The relative error ϵ_k^s is below the 2.8% both for $\mu = 0.8125$ cm and $\mu = 1.0625$ cm, the maximum value of the error being associated to very few points of the domain.

the hardware each implementation is optimized for – a CPU for the FOM and the POD-Galerkin ROM, a GPU⁶ for the DL-ROM. In Table 3 we report the CPU time needed to compute the FOM solution, and the POD-Galerkin ROM (at the testing phase), both on a full 64 GB node (20 Intel[®] Xeon[®] E5-2640 v4 2.4GHz cores), and the GPU time required by the DL-ROM to compute 875 time instances (the same number of time instants considered in the solution of the dynamical systems associated to the FOM and the POD-Galerkin ROM) at testing time, by fixing its dimension to $n = 5$, on an Nvidia GeForce GTX 1070 8 GB GPU. For the sake of completeness, we also report the computational time required by the DL-ROM when employing a single CPU node. It is evident that a POD-Galerkin ROM, built employing a global reduced basis ($N_c = 1$), is not amenable to a complex and challenging pathological cardiac electrophysiology problem like the figure of eight re-entry. Using a nonlinear approach, for which the solution manifold is approximated through a piecewise linear trial manifold (as in the case of $N_c = 2$ or $N_c = 4$ local reduced bases) reduces the online computational time. However, the DL-ROM still confirms to provide a more efficient ROM, almost 5 (or 2) times faster on the CPU, and 39 (or 19) faster on the GPU, than the POD-Galerkin ROM with $N_c = 2$ (or $N_c = 4$) local reduced bases.

In Fig 14 we show the trend of the error indicator (15) over the testing set versus the CPU computational time both for the DL-ROM and the POD-Galerkin ROM at testing phase. Slight improvements of the performance of DL-ROM, in terms of accuracy, are obtained for a small increase of the DL-ROM dimension n , coherently with our previous findings reported in [16]. Indeed, the DL-ROM is able, also in this case, to accurately represent the solution manifold by a reduced nonlinear trial manifold of dimension

⁶Indeed, at each layer of a neural network thousands of identical computations must be performed. The most suitable hardware architectures to carry out this kind of operations are GPUs because (i) they have more computational units (cores) and (ii) they have a higher bandwidth to retrieve from memory. Moreover, in applications requiring image processing, as CNNs, the graphics specific capabilities can be further exploited to speed up calculations.

Table 3. Test 2: FOM, POD-Galerkin ROM and DL-ROM computational times.

	time [s]	FOM/ROM dimensions
FOM (CPU)	382	$N = 65536$
DL-ROM (CPU/GPU)	15/1.2	$n = 5$
POD-Galerkin ROM $N_c = 1$ (CPU)	103	$n = 1538$
POD-Galerkin ROM $N_c = 2$ (CPU)	70	$n = 1158, 751$
POD-Galerkin ROM $N_c = 4$ (CPU)	33	$n = 435, 365, 298, 45$

POD-Galerkin ROM and DL-ROM computational times along with FOM and reduced trial manifold(s) dimensions. DL-ROM provides a more efficient ROM with respect the POD-Galerkin ROMs.

$n_\mu + 1 = 2$; for the case at hand, we report the results for $n = 5$ (very close to the intrinsic dimension $n_\mu + 1 = 2$ of the problem, and much smaller than the POD-Galerkin ROM dimension), providing slightly smaller values of the error indicator (15) than in the case $n = 2$. Regarding instead the POD-Galerkin ROM, in Fig 14 we report results obtained for different tolerances $\varepsilon_{POD} = 10^{-4}, 5 \cdot 10^{-4}, 10^{-3}, 5 \cdot 10^{-3}, 10^{-2}$. In the cases $N_c = 2$ and $N_c = 4$ we only report the results related to the smallest POD tolerances, which indeed allow us to meet the prescribed accuracy on the approximation of the gating variable, which would otherwise impact dramatically on the overall accuracy of the POD-Galerkin ROM. Moreover, we do not consider more than $N_c = 4$ local reduced bases in order not to generate too small local linear subspaces. As shown in Fig 14, the proposed DL-ROM outperforms the POD-Galerkin ROM in terms of both efficiency and accuracy.




Fig 14. Test 2: trend of the error indicator versus the CPU testing computational time. Error indicator ε_{rel} vs. CPU testing computational time for different values of N_c and ε_{POD} . The DL-ROM outperforms the POD-Galerkin ROM in terms of both efficiency and accuracy.

In Fig 15 we show the solutions obtained through the POD-Galerkin ROM with $N_c = 2$ (top) and $N_c = 4$ (bottom) local reduced bases, along with the relative error defined in (20), for the testing-parameter instance $\mu_{test} = 0.9625$ cm at $\hat{t} = 157.2$ ms. In both cases, we have considered the setting yielding the most efficient POD-Galerkin ROM approximation, which require about 30 (40, respectively) seconds to be evaluated. By comparing Fig 15 and Fig 11 (bottom), we observe that the DL-ROM outperforms

the POD-Galerkin ROM in terms of accuracy.




Fig 15. Test 2: POD-Galerkin ROM solutions for different testing-parameter instances. POD-Galerkin ROM solution (left) and relative error ϵ_k^s (right) for $N_c = 2$ (top) and $N_c = 4$ (bottom) at $\tilde{t} = 157.2$ ms, for $\mu_{test} = 0.9625$ cm.

In Fig 16 we show the action potentials obtained through the FOM, the DL-ROM and the POD-Galerkin ROM (with $N_c = 4$ local reduced bases), for the testing-parameter instance $\mu_{test} = 0.9625$ cm, and evaluated at $P_1 = (0.64, 1.11)$ cm and $P_2 = (0.69, 1.03)$ cm. These two points are close to the left core of the figure of eight re-entry, where a shorter action potential duration, and lower values of AP due to the meandering of the cores, are observed. The AP dynamics at those points is accurately captured by the DL-ROM, while the POD-Galerkin ROM might fail in this respect.




Fig 16. Test 2: FOM, POD-Galerkin ROM and DL-ROM APs at P_1 and P_2 . AP obtained through the FOM, the DL-ROM and the POD-Galerkin ROM with $N_c = 4$, for the testing-parameter instance $\mu_{test} = 0.9625$ cm, at $P_1 = (0.64, 1.11)$ cm and $P_2 = (0.69, 1.03)$ cm. The POD-Galerkin ROM approximations are obtained by imposing a POD tolerance $\epsilon_{POD} = 10^{-4}$ and 10^{-3} , resulting in error indicator (15) values equal to 5.5×10^{-3} and 7.6×10^{-3} , respectively.

Test 3: Three-dimensional ventricle geometry

We finally consider the solution of the coupled system (1) in a three-dimensional left ventricle (LV) geometry, obtained from the 3D Human Heart Model provided by Zygote [46]. Here, we consider a single ($n_\mu = 1$) parameter, given by the longitudinal conductivity in the fibers direction. The conductivity tensor takes the form

$$\mathbf{D}(\mathbf{x}; \mu) = \sigma_t \mathbf{I} + (\mu - \sigma_t) \mathbf{f}_0 \otimes \mathbf{f}_0, \quad (21)$$

where $\sigma_t = 12.9 \cdot 0.02$ mm²/ms; \mathbf{f}_0 is determined at each mesh point through a *rule-based* approach, by solving a suitable Laplace problem [47]. The resulting fibers field is reported in Fig 17. The applied current is defined as




Fig 17. Test 3: fibers distribution. Fibers field on the Zygote LV geometry.

$$I_{app}(\mathbf{x}, t) = \frac{C}{(2\pi)^{3/2} \alpha} \exp\left(-\frac{\|\mathbf{x} - \bar{\mathbf{x}}\|^2}{2\beta}\right) \mathbf{1}_{[0, \bar{t}]}(\tilde{t}),$$

where $\bar{t} = 2$ ms, $C = 1000$ mA, $\alpha = 50$, $\beta = 50$ mm², $\bar{\mathbf{x}} = [44.02, 1349.61, 63.28]^T$ mm.

In order to build the snapshot matrix \mathbf{S} , we solve problem (1) completed with the conductivity tensor (21) by means of linear finite elements, on a mesh made by $N = 16365$ vertices, and a semi-implicit scheme in time over a uniform partition of $(0, T)$ with $T = 300$ ms and time step $\Delta t = 0.1/12.9$. We uniformly sample $N_t = 1000$ time instances in $(0, T)$ and we zero-padded [34] the snapshot matrix to reshape each column in a 2D square matrix. The parameter space is provided by $\mathcal{P} = 12.9 \cdot [0.04, 0.4]$ mm²/ms; here we consider $N_{train} = 25$ training-parameter instances and $N_{test} = 24$ testing-parameter instances computed as in Test 2. In this case, the maximum number of epochs is set to $N_{epochs} = 30000$, the batch size is $N_b = 40$ and the training is stopped if the loss does not decrease over 4000 epochs.

In Fig 18 we report the FOM solution for two testing-parameter instances, i.e. $\mu = 12.9 \cdot 0.0739$ mm²/ms and $\mu = 12.9 \cdot 0.1991$ mm²/ms, at $\tilde{t} = 276$ ms, to show the variability of the FOM solution over the parameter space. As expected, front propagation is faster for larger values of the parameter μ .

In Fig 19-20 we report the FOM and DL-ROM solutions, the latter with $n = 10$, at $\tilde{t} = 42.1$ ms and $\tilde{t} = 222.1$ ms, for two testing-parameter instances, $\mu_{test} = 12.9 \cdot 0.1435$ mm²/ms and $\mu_{test} = 12.9 \cdot 0.3243$ mm²/ms. The DL-ROM approximation is essentially as accurate as the FOM solution.




Fig 18. Test 3: FOM solutions for different testing-parameter instances. FOM solutions for $\mu = 12.9 \cdot 0.0739 \text{ mm}^2/\text{ms}$ (left) and $\mu = 12.9 \cdot 0.1991 \text{ mm}^2/\text{ms}$ (right) at $\tilde{t} = 276 \text{ ms}$. By increasing the value of μ the wavefront propagates faster.




Fig 19. Test 3: comparison between FOM and DL-ROM solutions for a testing-parameter instance at different time instances. FOM solution (left) and DL-ROM one (right), with $n = 10$, at $\tilde{t} = 42.1 \text{ ms}$ (top) and $\tilde{t} = 276 \text{ ms}$ (bottom), for the testing-parameter instance $\mu_{test} = 12.9 \cdot 0.1435 \text{ mm}^2/\text{ms}$.

Also for this test case, it is possible to build a reduced nonlinear trial manifold of dimension very close to the intrinsic one $-n_\mu + 1 = 2$ – as long as the maximum number of epochs N_{epochs} is increased; the choice $n = 10$ is obtained as the best trade-off between accuracy and efficiency of the DL-ROM approximation in this case.

The DL-ROM approximation can also replace the FOM solution when evaluating outputs of interest. For instance, in Fig 21 and 22 we show the FOM and DL-ROM activation maps, the latter obtained by choosing $n = 10$ as DL-ROM dimension. Given the electric potential $u = u(\mathbf{x}, t; \boldsymbol{\mu})$, the (unipolar) activation map at a point $\mathbf{x} \in \Omega$ is evaluated as the minimum time at which the AP peak reaches \mathbf{x} , that is,

$$AC(\mathbf{x}; \boldsymbol{\mu}) = \arg \min_{t \in (0, T)} \left(u(\mathbf{x}, t; \boldsymbol{\mu}) = \max_{t \in (0, T)} u(\mathbf{x}, t; \boldsymbol{\mu}) \right).$$

Here we compare the activation maps AC_{FOM} and AC_{DL-ROM} obtained through the




Fig 20. Test 3: comparison between FOM and DL-ROM solutions for a testing-parameter instance at different time instances. FOM solution (left) and DL-ROM one (right), with $n = 10$, at $\tilde{t} = 42.1$ ms (top) and $\tilde{t} = 276$ ms (bottom), for the testing-parameter instance $\mu_{test} = 12.9 \cdot 0.3243$ mm²/ms.

FOM and the DL-ROM, respectively, by evaluating the maximum of the relative error

$$\epsilon_{AC}(\mathbf{x}; \boldsymbol{\mu}) = \frac{|AC_{FOM}(\mathbf{x}; \boldsymbol{\mu}) - AC_{DL-ROM}(\mathbf{x}; \boldsymbol{\mu})|}{|AC_{FOM}(\mathbf{x}; \boldsymbol{\mu})|}$$

over the N mesh points; in the case $\boldsymbol{\mu} = \mu_{test} = 12.9 \cdot 0.31$, the maximum relative error is equal to 4.32×10^{-5} .




Fig 21. Test 3: FOM activation map. FOM activation map for the testing-parameter instance $\mu_{test} = 12.9 \cdot 0.31$ mm²/ms.

In Fig 23 (left) we report the action potentials obtained with the FOM and the DL-ROM, this latter with $n = 20$, computed at point $P = [36.56, 1329.59, 28.82]$ mm for the testing-parameter instance $\mu_{test} = 12.9 \cdot 0.31$ mm²/ms. Moreover, we also report the best approximation of the FOM action potential over a POD space of same dimension $n = 20$ for the sake of comparison. Clearly, in dimension $n = 20$ the DL-ROM approximation is much more accurate than the POD best approximation; to reach the same accuracy (about $\epsilon_{rel} = 5.7 \times 10^{-3}$, measured through the error indicator (15)) achieved by the DL-ROM with $n = 20$, $n = 120$ POD modes would be required.




Fig 22. Test 3. DL-ROM activation map. DL-ROM activation map for the testing-parameter instance $\mu_{test} = 12.9 \cdot 0.31 \text{ mm}^2/\text{ms}$ with $n = 20$.

In Fig 23 (right) we highlight instead the improvements, in terms of efficiency, enabled by the use of the DL-ROM technique. In particular, we point out the CPU time required to solve the FOM for a testing parameter instance, and the one required by DL-ROM (of dimension $n = 10$) at testing time, by using the time resolution in each solution computation requires and by varying the FOM dimension N on a 6-core platform⁷ the FOM solution with $N = 16365$ degrees of freedom requires about 40 minutes to be computed, against 57 seconds required by the DL-ROM approximation, thus implying a speed-up almost equal to 41 times.




Fig 23. Test 3: FOM, DL-ROM and optimal-POD APs for a testing-parameter instance. FOM and DL-ROM CPU computational times. FOM, DL-ROM and optimal-POD APs for the testing-parameter instance $\mu_{test} = 12.9 \cdot 0.31 \text{ mm}^2/\text{ms}$ (left). For the same n , the DL-ROM is able to provide more accurate results than the optimal-POD. CPU time required to solve the FOM and by DL-ROM at testing time with $n = 10$ vs N (right). The DL-ROM is able to provide a speed-up equal to 41.

Conclusion

In this work we have proposed a new efficient reduced order model obtained using deep learning algorithms to boost the solution of parametrized problems in cardiac electrophysiology. Numerical results show that the resulting DL-ROM technique, formerly

⁷Numerical tests have been performed on a MacBook Pro Intel Core i7 6-core with 16 GB RAM.

introduced in [16], allows one to accurately capture complex wave propagation processes, both in physiological and pathological scenarios.

The proposed DL-ROM technique provides ROMs that are orders of magnitude more efficient than the ones provided by common linear (projection-based) ROMs, built for instance through a POD-Galerkin reduced basis method, for a prescribed level of accuracy. Through the use of DL-ROM, it is possible to overcome the main computational bottlenecks shown by POD-Galerkin ROMs, when addressing parametrized problems in cardiac electrophysiology. The most critical points related to (i) the linear superimposition of modes which linear ROMs are based on; (ii) the need to account for the gating variables when solving the reduced dynamics, even if not required; and (iii) the necessity to use (very often, expensive) hyper-reduction techniques to deal with terms that depend nonlinearly on either the transmembrane potential or the input parameters, are all addressed by the DL-ROM technique, which finally yields more efficient and accurate approximation than POD-Galerkin ROMs. Moreover, larger time resolutions can be employed when using a DL-ROM, compared to the ones required by the numerical solution of a dynamical systems through a FOM or a POD-Galerkin ROM. Indeed, the DL-ROM approximation can be queried at any desired time, without requiring to solve a dynamical system until that time, thus drastically decreasing the computational time required to compute the approximated solution at any given time.

Outputs of clinical interest, such as activation maps and action potentials, can be more efficiently evaluated by the DL-ROM technique than by a FOM built through the finite element method, while maintaining a high level of accuracy. This work is a proof-of-concept of the DL-ROM technique ability to investigate intra- and inter- subjects variability, towards performing multi-scenario analyses in real time and, ultimately, supporting decisions in clinical practice. In this respect, the use of DL-ROM techniques can foster assimilation of clinical data with physics-driven computational models.

Supporting information

SI Code. Code and data. The code used in this work can be downloaded from: <https://github.com/stefaniafresca/DL-ROM>. The training and testing datasets will be made available upon request to the authors.

SI Appendix. DL-ROM neural network architecture. Here we report the configuration of the DL-ROM neural network used for our numerical tests. We employ a 12-layers DFNN equipped with 50 neurons per hidden layer and n neurons in the output layer, where n corresponds to the dimension of the reduced nonlinear trial manifold. The architectures of the encoder and decoder functions are instead reported in Table 4 and 5.

Layer	Input dimension	Output dimension	Kernel size	# of filters	Stride	Padding
1			[5, 5]	8	1	SAME
2			[5, 5]	16	2	SAME
3			[5, 5]	32	2	SAME
4			[5, 5]	64	2	SAME
5	N	256				
6	256	n				

Table 4. Attributes of convolutional and dense layers in the encoder \mathbf{f}_n^E .

Layer	Input dimension	Output dimension	Kernel size	# of filters	Stride	Padding
1	n	256				
2	256	N				
3			[5, 5]	64	2	SAME
4			[5, 5]	32	2	SAME
5			[5, 5]	16	2	SAME
6			[5, 5]	1	1	SAME

Table 5. Attributes of dense and transposed convolutional layers in the decoder \mathbf{f}^D .

Acknowledgments

The authors have been partially supported by the ERC Advanced Grant iHEART, “An integrated heart model for the simulation of the cardiac function”, 2017-2022, P.I. A. Quarteroni (ERC2016AdG, project ID: 740132). Moreover, the authors acknowledge Dr. S. Pagani (MOX, Politecnico di Milano) for the kind help in the FOM software development, and Dr. M. Fedele for providing us the computational mesh of the Zygote Solid 3D heart model.

References

1. Quarteroni A, Manzoni A, Vergara C. The cardiovascular system: Mathematical modeling, numerical algorithms, clinical applications. *Acta Numerica*. 2017;26:365–590.
2. Quarteroni A, Dedè L, Manzoni A, Vergara C. Mathematical modelling of the human cardiovascular system: Data, numerical approximation, clinical applications. *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press; 2019.
3. Colli Franzone P, Pavarino LF, Scacchi S. *Mathematical cardiac electrophysiology*. vol. 13 of MS&A. Springer; 2014.
4. Sundnes J, Lines GT, Cai X, Nielsen BF, Mardal KA, Tveito A. *Computing the electrical activity in the heart*. vol. 1. Springer Science & Business Media; 2007.
5. Colli Franzone P, Pavarino LF. A parallel solver for reaction–diffusion systems in computational electrocardiology. *Mathematical Models and Methods in Applied Sciences*. 2004;14(06):883–911.
6. Mirams GR, Pathmanathan P, Gray RA, Challenor P, Clayton RH. Uncertainty and variability in computational and mathematical models of cardiac physiology. *The Journal of Physiology*. 2016;594.23:6833–6847.
7. Johnstone RH, Chang ETY, Bardenet R, de Boer TP, Gavaghan DJ, Pathmanathan P, et al. Uncertainty and variability in models of the cardiac action potential: Can we build trustworthy models? *Journal of Molecular and Cellular Cardiology*. 2016;96:49–62.
8. Hurtado DE, Castro S, Madrid P. Uncertainty quantification of two models of cardiac electromechanics. *International Journal for Numerical Methods in Biomedical Engineering*. 2017; 33(12):e2894.

9. Dhamala J, Arevalo HJ, Sapp J, Horáček BM, Wu KC, Trayanova NA, et al. Quantifying the uncertainty in model parameters using Gaussian process-based Markov chain Monte Carlo in cardiac electrophysiology. *Medical Image Analysis*. 2018;48:43–57.
10. Quaglino A, Pezzuto S, Koutsourelakis PS, Auricchio A, Krause R. Fast uncertainty quantification of activation sequences in patient-specific cardiac electrophysiology meeting clinical time constraints. *International Journal for Numerical Methods in Biomedical Engineering*. 2018;34(7):e2985.
11. Johnston BM, Coveney S, Chang ET, Johnston PR, Clayton RH. Quantifying the effect of uncertainty in input parameters in a simplified bidomain model of partial thickness ischaemia. *Medical & Biological Engineering & Computing*. 2018;56(5):761–780.
12. Pathmanathan P, Cordeiro JM, Gray RA. Comprehensive uncertainty quantification and sensitivity analysis for cardiac action potential models. *Frontiers in Physiology*. 2019;10.
13. Levrero-Florencio F, Margara F, Zacur E, Bueno-Orovio A, Wang Z, Santiago A, et al. Sensitivity analysis of a strongly-coupled human-based electromechanical cardiac model: Effect of mechanical parameters on physiologically relevant biomarkers. *Computer Methods in Applied Mechanics and Engineering*. 2020;361:112762.
14. Quarteroni A, Manzoni A, Negri F. *Reduced basis methods for partial differential equations: An introduction*. vol. 92. Springer; 2016.
15. Pagani S, Manzoni A, Quarteroni A. Numerical approximation of parametrized problems in cardiac electrophysiology by a local reduced basis method. *Computer Methods in Applied Mechanics and Engineering*. 2018;340:530–558.
16. Fresca S, Dedé L, Manzoni A. A comprehensive deep learning-based approach to reduced order modeling of nonlinear time-dependent parametrized PDEs. *arXiv preprint arXiv:200104001*. 2020.
17. Guo M, Hesthaven JS. Reduced order modeling for nonlinear structural analysis using Gaussian process regression. *Computer Methods in Applied Mechanics and Engineering*. 2018;341:807–826.
18. Guo M, Hesthaven JS. Data-driven reduced order modeling for time-dependent problems. *Computer Methods in Applied Mechanics and Engineering*. 2019;345:75–99.
19. Hesthaven J, Ubbiali S. Non-intrusive reduced order modeling of nonlinear problems using neural networks. *Journal of Computational Physics*. 2018;363:55–78.
20. San O, Maulik R. Neural network closures for nonlinear model order reduction. *Advances in Computational Mathematics*. 2018;44(6):1717–1750.
21. Kani JN, Elsheikh AH. DR-RNN: A deep residual recurrent neural network for model reduction. *arXiv preprint arXiv:170900939*. 2017.
22. Mohan A, Gaitonde DV. A deep learning based approach to reduced order modeling for turbulent flow control using LSTM neural networks. *arXiv preprint arXiv:18040926*. 2018.

23. Wan Z, Vlachas P, Koumoutsakos P, Sapsis T. Data-assisted reduced-order modeling of extreme events in complex dynamical systems. *PLoS one*. 2018;13.
24. González FJ, Balajewicz M. Deep convolutional recurrent autoencoders for learning low-dimensional feature dynamics of fluid systems. *arXiv preprint arXiv:180801346*. 2018.
25. Lee K, Carlberg K. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *arXiv preprint arXiv:181208373*. 2018.
26. Klabunde R. *Cardiovascular Physiology Concepts*. Lippincott Williams & Wilkins; 2011.
27. Aliev RR, Panfilov AV. A simple two-variable model of cardiac excitation. *Chaos Solitons Fractals*. 1996;7(3):293–301.
28. FitzHugh R. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*. 1961;1(6):445–466.
29. Nagumo J, Arimoto S, Yoshizawa S. An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*. 1962;50(10):2061–2070.
30. Nash MP, Panfilov AV. Electromechanical model of excitable tissue to study reentrant cardiac arrhythmias. *Progress in Biophysics and Molecular Biology*. 2004;85:501–522.
31. Mitchell CC, Schaeffer DG. A two-current model for the dynamics of cardiac membrane. *Bulletin of Mathematical Biology*. 2003;65(5):767–793.
32. Clayton RH, Bernus O, Cherry EM, Dierckx H, Fenton FH, Mirabella L, et al. Models of cardiac tissue electrophysiology: Progress, challenges and open questions. *Progress in Biophysics and Molecular Biology*. 2011;104(1):22–48.
33. Quarteroni A, Valli A. *Numerical approximation of partial differential equations*. vol. 23. Springer; 1994.
34. Goodfellow I, Bengio Y, Courville A. *Deep Learning*. MIT Press; 2016. Available from: <http://www.deeplearningbook.org>.
35. LeCun Y, Bottou L, Bengio Y, Haffner P. Gradient based learning applied to document recognition. *Proceedings of the IEEE*. 1998; p. 533–536.
36. Hinton GE, Zemel RS. Autoencoders, minimum description length, and Helmholtz free energy. *Proceedings of the 6th International Conference on Neural Information Processing Systems (NIPS'1993)*. 1994;.
37. Kingma DP, Ba J. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*; 2015.
38. Abadi M, Barham P, Chen J, Chen Z, Davis A, Dean J, et al. TensorFlow: A system for large-scale machine learning; 2016. Available from: <https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf>.
39. Göktepe S, Wong J, Kuhl E. Atrial and ventricular fibrillation: Computational simulation of spiral waves in cardiac tissue. *Archive of Applied Mechanics*. 2010;80:569–580.

40. ten Tusscher KHWJ, Panfilov AV. Alternans and spiral breakup in a human ventricular tissue model. *American Journal of Physiology-Heart and Circulatory Physiology*. 2006;291(3):H1088–H1100.
41. Trayanova NA. Whole-heart modeling applications to cardiac electrophysiology and electromechanics. *Circulation Research*. 2011;108:113–28.
42. Plank G, Zhou L, Greenstein J, Cortassa S, Winslow R, O'Rourke B, et al. From mitochondrial ion channels to arrhythmias in the heart: Computational techniques to bridge the spatio-temporal scales. *Philosophical Transactions Series A, Mathematical, Physical, and Engineering Sciences*. 2008;366:3381–409.
43. Nattel S. New ideas about atrial fibrillation 50 years on. *Nature*. 2002;415:219–226.
44. Nagaiah C, Kunisch K, Plank G. Optimal control approach to termination of re-entry waves in cardiac electrophysiology. *Journal of Mathematical Biology*. 2012;67.
45. ten Tusscher K. Spiral wave dynamics and ventricular arrhythmias. PhD Thesis. 2004.
46. Zygote solid 3D heart generation II development report. Zygote Media Group Inc.; 2014.
47. Rossi S, Lassila T, Ruiz Baier R, Sequeira A, Quarteroni A. Thermodynamically consistent orthotropic activation model capturing ventricular systolic wall thickening in cardiac electromechanics. *European Journal of Mechanics - A/Solids*. 2014;48:129–142.

MOX Technical Reports, last issues

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

- 39/2020** Martinolli, M.; Biasetti, J.; Zonca, S.; Polverelli, L.; Vergara, C.
Extended Finite Element Method for Fluid-Structure Interaction in Wave Membrane Blood Pumps
- 38/2020** Sollini, M.; Kirienco, M.; Cavinato, L.; Ricci, F.; Biroli, M.; Ieva, F.; Calderoni, L.; Tabacchi,
Methodological framework for radiomics applications in Hodgkin's lymphoma
- 34/2020** Antonietti, P.F.; Mazzieri, I.; Nati Poltri, S.
A high-order discontinuous Galerkin method for the poro-elasto-acoustic problem on polygonal and polyhedral grids
- 35/2020** Morbiducci, U.; Mazzi, V.; Domanin, M.; De Nisco, G.; Vergara, C.; Steinman, D.A.; Gallo, D.
Wall shear stress topological skeleton independently predicts long-term restenosis after carotid bifurcation endarterectomy
- 36/2020** Pellagatti, M.; Masci, C.; Ieva, F.; Paganoni A.M.
Generalized Mixed-Effects Random Forest: a flexible approach to predict university student dropout
- 37/2020** Fumagalli, A.; Scotti, A.
A mathematical model for thermal single-phase flow and reactive transport in fractured porous media
- 31/2020** Bernardi, M.S.; Africa, P.C.; de Falco, C.; Formaggia, L.; Menafoglio, A.; Vantini, S.
On the Use of Interferometric Synthetic Aperture Radar Data for Monitoring and Forecasting Natural Hazards
- 32/2020** Menafoglio, A.; Sgobba, S.; Lanzano, G.; Pacor, F.
Simulation of seismic ground motion fields via object-oriented spatial statistics: a case study in Northern Italy
- 33/2020** Centofanti, F.; Fontana, M.; Lepore, A.; Vantini, S.
Smooth LASSO Estimator for the Function-on-Function Linear Regression Model
- 29/2020** Piersanti, R.; Africa, P.C.; Fedele, M.; Vergara, C.; Dede', L.; Corno, A.F.; Quarteroni, A.
Modeling cardiac muscle fibers in ventricular and atrial electrophysiology simulations