



Universiteit
Leiden
The Netherlands

Fast creation of data-driven low-order predictive cardiac tissue excitation models from recorded activation patterns

Kabus, D.; Coster, T. de; Vries, A.A.F. de; Pijnappels, D.A.; Dierckx, H.

Citation

Kabus, D., Coster, T. de, Vries, A. A. F. de, Pijnappels, D. A., & Dierckx, H. (2024). Fast creation of data-driven low-order predictive cardiac tissue excitation models from recorded activation patterns. *Computers In Biology And Medicine*, 169.
doi:10.1016/j.combiomed.2024.107949

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/4246781>

Note: To cite this publication please use the final published version (if applicable).



Fast creation of data-driven low-order predictive cardiac tissue excitation models from recorded activation patterns

Desmond Kabus^{a,b}, Tim De Coster^b, Antoine A.F. de Vries^b, Daniël A. Pijnappels^b, Hans Dierckx^{a,*}

^a Department of Mathematics, KU Leuven Campus Kortrijk (KULAK), Etienne Sabbelaan 53, 8500, Kortrijk, Belgium

^b Laboratory of Experimental Cardiology, Leiden University Medical Center (LUMC), Albinusdreef 2, 2333 ZA, Leiden, The Netherlands

ARTICLE INFO

Dataset link: [GitLab.com/HeartKOR, distephym 1.0.0, sparsealgs 1.0.0, sappho 1.0.0, ithildin 3.2.2, ithilmin 1.0.0, py_ithildin 0.6.1](https://gitlab.com/HeartKOR/distephym), DOI: 10.5281/zenodo.8183722

Keywords:

Data-driven modelling
Surrogate modelling
Machine learning
Digital twin
Excitable media
Cardiac electrophysiology
Optical voltage mapping

ABSTRACT

Excitable systems give rise to important phenomena such as heat waves, epidemics and cardiac arrhythmias. Understanding, forecasting and controlling such systems requires reliable mathematical representations. For cardiac tissue, computational models are commonly generated in a reaction–diffusion framework based on detailed measurements of ionic currents in dedicated single-cell experiments.

Here, we show that recorded movies at the tissue-level of stochastic pacing in a single variable are sufficient to generate a mathematical model. Via exponentially weighed moving averages, we create additional state variables, and use simple polynomial regression in the augmented state space to quantify excitation wave dynamics. A spatial gradient-sensing term replaces the classical diffusion as it is more robust to noise. Our pipeline for model creation is demonstrated for an in-silico model and optical voltage mapping recordings of cultured human atrial myocytes and only takes a few minutes.

Our findings have the potential for widespread generation, use and on-the-fly refinement of personalised computer models for non-linear phenomena in biology and medicine, such as predictive cardiac digital twins.

1. Introduction

A wide variety of phenomena qualify as excitable systems, such as epidemics [1] or electrical excitation in the brain [2] or heart [3]. Reliable and accurate mathematical modelling of these systems is necessary to forecast their evolution in time. These predictions are useful for decision making to steer these systems to a desired state.

To diagnose and treat heart rhythm disorders, a common health burden, the scientific community is creating sophisticated models for different processes in the heart with the goal of building a digital twin of each patient's heart, i.e., a full recreation of their heart in the computer [4,5]. One of them is the contraction of the heart, which is controlled by excitation waves travelling through the myocardium [3]. An in-silico model of cardiac tissue electrophysiology is a mathematical model that is able to describe the dynamics of the excitation patterns [3]. Typically, in-silico models for cardiac electrophysiology fall into one of two categories. On the one hand, there are the mathematical models that focus mostly on the overall dynamics in the tissue [6–12]. While most of these models have continuous state variables, there are also cellular automata where each cell can only be in one of finitely many states [13]. On the other hand there are detailed models that aim to model each current relevant for the specific use case of the

model across the cell membrane [14–16]. Typically, these models are fit using in-vitro data obtained from patch-clamping of single cells, or from optical voltage or Ca^{2+} mapping of monolayers or of the epicardium. Capturing cell memory is also crucial in the modelling of cardiac tissue. This may for instance be done via ionic current gates opening and closing or exponentially decaying memory variables [17,18]. A common property in the models focusing on the overall dynamics rather than the individual ion currents is that the state variables used in them often do not have an equivalent in the physical tissue. These internal, hidden states are not fit to actual data, instead they are often chosen as a means to an end.

A so-called surrogate model is a mathematically simpler model that is fit to an existing, more detailed ionic model. For instance, Bueno-Orovio et al. [11] provide parameter sets for their model to act as a surrogate for the models by Priebe and Beuckelmann [19] and ten Tusscher et al. [20]. Surrogate modelling can be used to create patient-specific models much more quickly, such that they can be used in a clinical setting after a relatively short amount of time, e.g. during a hospital visit [21]. The current models of cardiac tissue electrophysiology are typically not fit to actual recordings of the wave dynamics.

* Corresponding author.

E-mail address: h.dierckx@kuleuven.be (H. Dierckx).

<https://doi.org/10.1016/j.combiomed.2024.107949>

Received 20 September 2023; Received in revised form 7 December 2023; Accepted 1 January 2024

Available online 3 January 2024

0010-4825/© 2024 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

We are in this work fitting models to such data in a controlled case, i.e., dense data from isotropic monolayers, which also enables us to directly compare forward simulation to experiment.

Machine learning has become a useful tool in cardiac electrophysiology. The eikonal equation has been solved using physics-informed neural networks, which means that the local arrival time of the excitation wave can be recovered well for the entire domain based on relatively few data points [22]. This method is limited, however, to only model excitation, not recovery and re-excitation. Physics-informed neural networks have also been used to solve the reaction–diffusion systems for Aliev and Panfilov [9] dynamics, as well as recovering one or two parameters of this model from data [23]. Reservoir computing approaches, such as echo-state networks, can also be used to predict the temporal evolution of cardiac action potentials [24]. None of these methods are yet used to create full standalone tissue models for cardiac electrophysiology in a data-driven way.

A recent step forward in in-vitro models for cardiac electrophysiology was the conditional immortalisation of human atrial myocytes [25]. This resulted in cell lines allowing the generation of near infinite numbers of fully functional human atrial myocytes called hiAMs. These hiAMs overcome a major problem associated with the use of heart muscle cells of animals for in-vitro modelling, namely the existence of species specific differences in electrophysiological behaviour including differences in the action potential duration (APD) and shape, conduction velocity (CV), and wave propagation [3,16,25–27]. No in-silico version of the electrophysiological behaviour of this cell line has yet been published.

In this work, we present a general method to create data-driven in-silico tissue models from just one recorded variable over two-dimensional space and time. The method is more generally applicable and faster than the conventional creation of new models. Our method requires only minimal assumptions and works by expanding the data to a multi-dimensional state space. Models created with our method are designed directly at the tissue scale, instead of needing extra steps transitioning from the cell kinetics to tissue dynamics. We both re-create an existing in-silico model with our method as a surrogate model, and use this method to build an in-silico tissue model for hiAMs only from optical voltage mapping (OVM) data. Note that our model creation pipeline is not specific to this cell line and may be applied to OVM data in general. The generated models are specialised to the tissues for which data is contained in the data set. However, since model generation is cheap and robust, models for different tissues can easily be created.

The steps of our method are outlined graphically in Fig. 1. In Section 2, we describe how the data are acquired (2.1), how they are processed, expanded with additional state variables (2.2), statistically binned in multi-dimensional state space histograms (2.3), and how the data-driven models are finally fit to them (2.4). In Section 3, we present results from the models that are generated with this method. Finally, in Sections 4–5, we wrap up by summarising the method and discussing its advantages and shortcomings, as well as giving an outlook on how the models fit by our method can be improved.

2. Methods

2.1. Data acquisition

In this work, we study spatio-temporal data from two excitable media: an in-vitro data set from OVM of monolayers of cardiomyogenically differentiated hiAMs and an in-silico data set from finite-differences simulations solving the mono-domain equations [3]. In this section, we outline how these data are generated and processed.

2.1.1. Optical voltage mapping

To obtain a two-dimensional recording of the electrical activity of tissue as an in-vitro data set, we performed OVM of hiAM monolayers in 10-cm² wells of six-well culture plates [25]. The recorded data consist of one variable, i.e., the measured difference in intensity of the light emitted by the voltage-sensitive dye, over space and time. We have used a MiCAM05-Ultima camera by SciMedia with a resolution of 100 × 100 pixels. With the particular configuration of lenses, each square pixel has a size of $\Delta x = 250 \mu\text{m}$. For the recordings, we use the lowest supported duration between frames of $\Delta t = 1 \text{ ms}$. We record 61 440 consecutive frames, which corresponds to slightly more than a minute and lines up nicely in the camera's file format. For modelling purposes, we assume that the monolayer is approximately homogeneous and isotropic.

To excite the tissue, we place a bi-polar stimulation electrode with two poles with a diameter of 0.5 mm and inter-electrode spacing of 1 mm at the well's boundary, and stimulate with $N_{\text{pulse}} = 200$ current pulses of random amplitude, random duration, and random timing. We call this procedure the stochastic burst pacing (SBP) protocol, inspired by stochastic pacing experiments done by Krogh-Madsen et al. [28]. The amplitudes A_n are sampled from a uniform distribution $U(-0.8 \text{ mA}, 0.8 \text{ mA})$. The times $T_n = T_u + T_e$ between subsequent pulses are sampled from a combination of a uniform and exponential distribution:

$$T_u \sim U(0 \text{ ms}, 50 \text{ ms}) \quad (1)$$

$$T_e \sim \text{Exp}(1/200 \text{ ms}) \quad (2)$$

where $U(a, b)$ denotes the uniform distribution on the interval $[a, b]$ and $\text{Exp}(\lambda)$ the exponential distribution with decay rate λ :

$$\text{Exp}(\lambda)(t) \propto \begin{cases} e^{-\lambda t} & t \geq 0 \\ 0 & t < 0 \end{cases} \quad (3)$$

We choose this probability distribution of durations such that, on the one hand, it covers a typical range of values, but on the other, occasionally, durations are sampled that are much longer than usual.

The durations $\tau_n = \tau_u + \tau_e$ of stimulation are sampled from a similar probability distribution:

$$\tau_u \sim U(0 \text{ ms}, 10 \text{ ms}) \quad (4)$$

$$\tau_e \sim \text{Exp}(1/40 \text{ ms}) \quad (5)$$

The stimulus current then takes the form:

$$I_{\text{stim}}(t) = \sum_{n=1}^{N_{\text{pulse}}} A_n [\text{H}(t - t_n) - \text{H}(t - t_n - \tau_n)] \quad (6)$$

with the start time of the n th pulse:

$$t_n = \sum_{m=1}^n [T_m + \tau_m] \quad (7)$$

and the Heaviside function:

$$\text{H}(t) = \begin{cases} 0 & t < 0 \\ \frac{1}{2} & t = 0 \\ 1 & t > 0 \end{cases} \quad (8)$$

We preprocess these data in a similar way as in a previous publication [29] by Gaussian blurring with a kernel size of three grid points and rescale the data to zero-mean and unit variance.

To bring the resting state close to zero, we subtract the baseline of the data. As baseline, we take the most common recorded intensity in each 4×4 pixel neighbourhood over a time span of 2 s. This is based on the assumption that this most common value corresponds to the resting state.

The final preprocessing step for this data set is to smooth the data using an exponential moving average (EMA) with $\alpha = 0.1$, which, for




Fig. 1. Outline of the model fitting procedure. In the first step, relevant features are extracted from the spatio-temporal data u . Then the state space is discretised and histogrammed. The value of the time-derivative $\dot{u}$ of the main variable is analysed statistically in this state space. A model can then be fit to the state space and used to predict the evolution of the medium.

time $t \geq 0$ with frame duration Δt and the signal over time $s(t)$ at each position in space, is iteratively defined as:

$$\tilde{s}(t) = \begin{cases} \alpha s(t) + [1 - \alpha] \tilde{s}(t - \Delta t) & t > 0 \\ s(t = 0) & t = 0 \end{cases} \quad (9)$$

which, within first-order time discretisation, is equivalent to:

$$\tilde{s}(t) = e^{-at} s(0) + a \int_0^t e^{a[t'-t]} s(t') dt' \quad (10)$$

with $a = \frac{\alpha}{\Delta t} > 0$.

We refer to this recording as the OVM data set and $u(t, \mathbf{x})$ denotes the unit-less signal after preprocessing scaled such that the resting state is close to $u = 0$ and the excited state around $u = 1$. This signal u corresponds to the electrical activity of the cells, i.e., it is linked to the transmembrane voltage over space and time.

Only for qualitative comparisons after a data-driven model is fit to the OVM data set, a second OVM recording is used. In this second recording, several minutes after rotor formation through burst pacing, the resulting single, stable spiral wave is observed. This experiment is also performed with hiAMs in the same six-well format as before and at the same spatial resolution but at a lower temporal resolution of 6 ms between frames. We preprocess these data in the same way as the main OVM data set.

2.1.2. Numerical simulation

To obtain a fully controllable testing data set, we ran an in-silico experiment similar to the in-vitro experiment in Section 2.1.1. We use the finite-differences method to solve the reaction-diffusion system using a rescaled version of the minimalist two-variable tissue model by Aliev and Panfilov [9], or in short the AP96 model, whose equations we will briefly summarise (Eqs. (14)–(15)).

As the diffusivity is proportional to the square of the CV, we set it such that the medium has a measured CV close to that of hiAMs [3,25]:

$$D = 0.0625 \text{ mm}^2 \text{ ms}^{-1} \quad (11)$$

$$\Rightarrow \text{CV} = 0.17 \text{ m s}^{-1} \quad (12)$$

We rescale the entire model over time such that its APD has a value of roughly the APD of hiAMs [25]:

$$r = \frac{\text{APD}_{\text{AP96}}}{\text{APD}_{\text{hiAM}}} = \frac{25.58}{100 \text{ ms}} = 0.2558 \text{ ms}^{-1} \quad (13)$$

Table 1

Parameters of the rescaled AP96 model.

Parameter	Value
a	0.15
ϵ_0	0.002
k	8
μ_1	0.2
μ_2	0.3
D	$0.0625 \text{ mm}^2 \text{ ms}^{-1}$
r	0.2558 ms

The time derivatives are rescaled by this time scaling factor r [9]:

$$\partial_t u = \nabla \cdot D \nabla u - rku[u - a][u - 1] - ruv + \frac{i_{\text{stim}}}{C} \quad (14)$$

$$\partial_t v = -rev - rku[u - a - 1] \quad (15)$$

$$\epsilon = \epsilon_0 + \frac{\mu_1 v}{u + \mu_2} \quad (16)$$

By doing so, we introduced the millimetre as the spatial unit and the millisecond for time into the AP96 model, which uses arbitrary space and time units. The variables $u(t, \mathbf{x})$ and $v(t, \mathbf{x})$ remain unit-less. The value $u = 0$ corresponds to the resting state and when excited, u approaches the value $u = 1$.

For all other parameters, a , ϵ_0 , k , μ_1 , and μ_2 , which are all unitless, we use the default values, as in the original publication by Aliev and Panfilov [9]. An overview of all parameters can be found in Table 1.

We have run a simulation over 43 104 frames at $\Delta t = 1 \text{ ms}$ per frame at the same resolution as for the in-vitro OVM data and with the same SBP protocol. The length of the simulation was chosen such that the tissue can fully recover after the last of the 200 pulses of the SBP protocol. The electrode was placed off-centre at the side of one edge of the medium. For the time integration, we used a simple forward Euler step with numerical time step 0.1 ms and a second-order central finite-differences five-point stencil for the Laplacian $\nabla \cdot D \nabla$. The square simulated tissue was homogeneous and isotropic with no-flux boundary conditions.

We refer to this synthetic recording as the AP96 data set. We only include the variable $u(t, \mathbf{x})$ in the data set, as in the in-vitro data set, i.e., the OVM data set. For surrogate model fitting, we hence also only use one variable related to the transmembrane voltage of the tissue and not the internal state of the cell, which in the AP96 model is encapsulated in the second variable $v(t, \mathbf{x})$.

In the subsequent steps, only these data will be used, partially as training data and partially as testing data. However, we also generate auxiliary data sets for the AP96 model that are only used for qualitative comparisons with the data-driven models after fitting.

To generate a recording of a rotor using the AP96 model, we ran another two-dimensional simulation at the same resolution in space and time with initial conditions that led to a spiral wave: The transmembrane voltage was set to a value of $u = 0.95$ within the rectangle $[0, 0.5L_x] \times [0.3L_y, 0.5L_y]$, and the second variable to $v = 1$ within a second rectangle $[0, 0.5L_x] \times [0.4L_y, 0.6L_y]$, where $L_x = L_y = 100$ mm. Gaussian blurring with a radius of four grid lengths was applied to avoid steep gradients at the edge of the rectangles.

We also measured restitution curves, i.e., APD and CV as functions of the cycle length (CL) of stimulation. For this, we ran a one-dimensional cable simulation of the model with the same spatial and temporal resolution. At the left end of the cable, we stimulated by instantaneously setting u to a value above the excitation threshold to trigger a wave travelling to the right. We repeatedly stimulated starting at a CL of 300 ms, decreasing in steps of 1 ms until we reached a CL where not every stimulus resulted in a pulse that travelled through the entire cable, i.e., where 2:1 conduction block begins.

2.2. State space augmentation

The multi-dimensional space spanned by all variables describing the state of the modelled medium, $\mathbf{u} = [u, v, \dots]^T$, is called state space in contrast to the physical space spanned by $\mathbf{x}$ and time t .

An in-silico tissue model usually prescribes a unique reaction $\mathbf{R}$ at each position $\mathbf{u}$ in the state space. We write a more generic version of Eqs. (14)–(15), the mono-domain description [3]:

$$\partial_t \mathbf{u} = \mathbf{P} \nabla \cdot \mathbf{D} \nabla \mathbf{u} + \mathbf{R}(\mathbf{u}) \quad (17)$$

with the diffusivity matrix $\mathbf{D}$, where $\mathbf{P}$ is a diagonal matrix describing which variables in $\mathbf{u}$ are to be diffused.

For the two-variable AP96 model, with $\mathbf{u} = [u, v]^T$ and $\mathbf{P} = \text{diag}(1, 0)$, the state space of the AP96 model is shown in Fig. 2, coloured only by the component in u of the reaction term $\mathbf{R}$. It can be seen that the second variable v can be used to distinguish cases with positive and negative reaction at the same value of u . A model that can describe both excitation from a resting state and recovery to it, is only possible with at least two variables, such that the inertial manifold can be a loop in the state space.

In both types of data sets considered in this work, we only have data for one variable available, namely u . To still be able to differentiate between excitation and recovery, we extract additional information from u over space and time to use as additional states for the state space, so-called augmentations. A well-designed augmented state space unfolds the data such that cells in the tissue with different histories are identified by different augmented state vectors.

As we want to find models that can be used as forward models in a similar way as in Eq. (17), we impose two more requirements on the states with which we augment the state space:

1. Augmentations should be local in space, i.e., they should only use states within a neighbourhood around each point.
2. Causality must be respected, i.e., only data from previous points in time and the current point in time may be used to compute each augmentation.

The two augmentations that we found to be most useful for model creation are presented in the remainder of this section.




Fig. 2. Reaction in state space of the AP96 model. The reaction term $\mathbf{R}$ is a vector field defining a reaction at each point $\mathbf{u}$ in the state space. Colouring is according to the u -component of the reaction term $\mathbf{R}_u$. The trajectory through the state space from excitation of one point is shown as an example.

2.2.1. Exponential moving averages

The goal of the state space augmentation is to capture that the model can evolve differently from the same value of u . This value may increase at the front of waves, during excitation, or decrease at the wave backs, during recovery.

One way to sense excitation and recovery is by expanding the state space with augmented states that capture the passage of time. A naïve way to do so may be to just include a time-delayed signal, which is a useful augmentation in phase mapping [30]. This is, however, not very robust to noise, as noise in the past may directly impact the present. Another option may be to use Hilbert transforms, but as their computation involves an integral over all of time, they violate causality and are therefore not meeting our requirements [31].

The EMA (Eq. (9)) is robust to noise and typically lags behind the signal by an amount of time inversely proportional to the constant α . It can intuitively be used to differentiate between excitation and recovery for signals u in-between the resting state and excited state: If an EMA with a suitable scale α is lower than u , it is likely that the medium is currently getting excited; if it is higher, the cells are likely recovering.

Multiple EMAs can be used to be able to capture more advanced behaviour: With an appropriately chosen α , we can differentiate between states where the tissue is activated a lot and where it is mostly resting. With such an EMA, restitution characteristics can be captured.

Here, we augment the AP96 data with a single EMA $\tilde{u}$ with $\alpha = 3.0 \times 10^{-2}$. For $\tilde{u}$ in the OVM data, we use $\alpha = 8.0 \times 10^{-3}$.

2.2.2. Standard deviation as a proxy for the absolute gradient

We also conjecture that information about the slopes and higher derivatives in space of u is useful for a well fitting data-driven model. In the reaction–diffusion description, the Laplacian is used to capture curvature effects in space. Because it is a second-order derivative, it is very sensitive to noise. The absolute value of the gradient in combination with other state space variables can also encode much information about the slopes of u . We use a proxy for the absolute gradient in a neighbourhood with high robustness with respect to noise. Such a proxy is the standard deviation (SD) of the signal u in the local neighbourhood up to a radius R around $\mathbf{x}$.

It can be shown analytically that the SD is proportional to the absolute gradient of a plane $h(\mathbf{x}) = \mathbf{u}^T \mathbf{x} + b$, which evaluates to $\|\nabla h\| =$




Fig. 3. Likelihood of each point in the state space of the AP96 model for the numerical simulation generating the data set used for surrogate model fitting.

$\|\mathbf{w}\|$. The SD in a spherical neighbourhood $\mathcal{N}$ up to a radius R around the origin can be found via integration in polar coordinates r, φ :

$$\sigma_h^2 = \langle [h - b]^2 \rangle = \frac{1}{\int_{\mathcal{N}} d\mathbf{x}} \int_{\mathcal{N}} d\mathbf{x} [\mathbf{w}^T \mathbf{x}]^2 \quad (18)$$

$$\sigma_h^2 = \frac{1}{\pi R^2} \int_0^R dr \int_0^{2\pi} d\varphi [w_x r \cos \varphi + w_y r \sin \varphi]^2 \quad (19)$$

$$\sigma_h = \frac{R}{2} \|\mathbf{w}\| \quad (20)$$

Analogously, a similar relation can be shown in three dimensions:

$$\sigma_h = \frac{R}{\sqrt{5}} \|\mathbf{w}\| \quad (21)$$

Assuming that the signal u varies slow enough that a linear function can approximate it well within the chosen neighbourhood, the SD can be used to compute its absolute gradient. Despite the relatively high time-derivative $\dot{u}$ of the cardiac action potential during its upstroke, and hence ∇u , we can make this assumption if the resolution is sufficiently high.

Note that the SD has the useful property, that it can still be computed at or near the boundary of the medium, just based on fewer samples than for fully interior points.

To capture gradient information, we augment the AP96 data set with an SD variable for the neighbourhood up to a radius $R = 5\Delta x$ around each point $\mathbf{x}$ and denote it by g . For g in the OVM data set, we use a larger neighbourhood with $R = 8\Delta x$ for the SD calculation.

2.3. State space histograms

In the usual dynamics of the excitable systems we study in this work, not all points $\mathbf{u}$ in the state space are equally likely. Some state vectors $\mathbf{u}$ may be quite common as being part of the inertial manifold, while others might be completely absent in the data.

For example, for the AP96 model with $\mathbf{u} = [u, v]^T$, it can be seen in Fig. 3 that a part of the state space is not visited and most of the samples are on the inertial manifold, a closed approximately one-dimensional loop. We conjecture that in higher dimensions the state space is even more sparse, e.g. in case of detailed ionic models, some of which contain more than ten state variables.

With the chosen state space expansion (Section 2.2), the state vector is $\mathbf{u} = [u, \tilde{u}, g]^T$. We discretise the state space to 100 bins in each

dimension and collect the data in a multi-dimensional histogram. This results in three-dimensional state space histograms for the two data sets with 10^6 bins each. For efficient memory usage, these histograms are computationally implemented using multi-dimensional sparse arrays. Regions in the state space containing no data do not take up memory in this way. We can then fit the models to the average value inside each bin instead of to all data points. This efficient memory management enables building even higher-dimensional state spaces, though in this work we require only three variables.

The time-derivative $\dot{u}$ fully defines the reaction of the data-driven models developed in this work. The reason for this is that all other variables in the state vector $\mathbf{u}$ can be derived from the one variable u over space and time.

In a well chosen state space, for each state vector $\mathbf{u}$, the reaction $\dot{u}$ is the same, i.e., the states are well separated. Therefore, we calculate the mean value of $\dot{u}$ computed using finite-differences at each bin in the discretised state space, as well as its SD. This mean value of $\dot{u}$ takes a similar role as R_u for the AP96 model and “colours” the state space in a comparable way to R_u in Fig. 2.

2.4. Model fitting

A data-driven model will be fit to the data in state space as a scalar-valued function f :

$$f : \mathbf{u} \mapsto \dot{u} \quad (22)$$

such that the model consists of the state space expansion $\mathbf{u}$ by calculating the EMA (Eq. (9)), the SD within the neighbourhoods around each point (Section 2.2.2), as well as the forward Euler step:

$$u(t + \Delta t, \mathbf{x}) = u(t, \mathbf{x}) + \Delta t f(u(t, \mathbf{x})) \quad (23)$$

As the model function f , we choose a low-order polynomial which is linear with respect to each of the components:

$$f(\mathbf{u}) = w_{\tilde{u}g} \tilde{u}g + w_{\tilde{u}g} \tilde{u}g + w_{\tilde{u}u} \tilde{u}u + w_{\tilde{u}u} \tilde{u}u + w_{gu} gu + w_g g + w_u u + w \quad (24)$$

Assuming that the SD g approximates the absolute gradient $\|\nabla u\|$ of the main state variable u well, we hence essentially describe the propagation of the activation waves using advection, instead of diffusion (Eq. (17)):

$$\partial_t u = f(\mathbf{u}) = f_0(u, \tilde{u}) + c_1 f_1(u, \tilde{u}) \|\nabla u\| \quad (25)$$

$$f_0(u, \tilde{u}) = w + w_u u + w_{\tilde{u}} \tilde{u} + w_{\tilde{u}u} \tilde{u}u \quad (26)$$

$$f_1(u, \tilde{u}) = w_g + w_{gu} gu + w_{\tilde{u}g} \tilde{u}g + w_{\tilde{u}gu} \tilde{u}gu \quad (27)$$

$$\tilde{u}(t) = e^{-at} u(0) + a \int_0^t e^{a[t'-t]} u(t') dt' \quad (28)$$

with real-valued scalar constants $a, c_1, w, \dots$ and $g \approx c_1 \|\nabla u\|$. For an appropriate choice of parameters, obtained for instance through fitting, Eqs. (25)–(28) describe a model of excitation in an advection-regression framework.

To prevent the fit model from leaving the valid state space, after each forward Euler step, we clip the predicted u to the range $[u_{\min}, u_{\max}]$ for which we have training data, the so-called valid range. Note that this also limits the EMA $\tilde{u}$ to that range, while the SD g can still leave the range observed in the data, i.e., for larger than typical slopes.

In each step forward in time, $t \leftarrow t + \Delta t$, we update the state vector $\mathbf{u} = [u, \tilde{u}, g]^T$ of the data-driven model according to the following pseudocode:

1. Compute the new value of u at each point $\mathbf{x}$ using the polynomial f (Eqs. (23)–(24)):

$$u \leftarrow u + \Delta t f(\mathbf{u}) \quad (29)$$

$$u \leftarrow \max(u_{\min}, \min(u, u_{\max})) \quad (30)$$

2. Update the EMA $\tilde{u}$ based on its old value and the updated u (Eq. (9)):

$$\tilde{u} \leftarrow \alpha u + [1 - \alpha] \tilde{u} \quad (31)$$

3. Calculate the SD g of the updated u (Section 2.2.2):

$$g \leftarrow \left\langle u^2 - \langle u \rangle^2 \right\rangle^{\frac{1}{2}} \quad (32)$$

where $\langle \cdot \rangle$ denotes the average of all points in the spherical neighbourhood $\mathcal{N}$ of each position $\mathbf{x}$ up to a radius R .

To find the optimal model function f , we use the least-squares algorithm to find the polynomial of the given order which minimises the loss:

$$\mathcal{L}(f) = \sum_{u \in \mathcal{U}} N(u) [\dot{u}(u) - f(u)]^2 + \sum_{u \in \mathcal{R}} f^2(u) \quad (33)$$

where $\mathcal{U}$ denotes the centres of all bins in the discretised state space containing at least one sample, i.e., one data point, and $N(u)$ is the number of samples in the bin u . To regularise the model space in favour of more physically realistic models, we add additional co-location points $\mathcal{R}$ outside of $\mathcal{U}$ at which we assume f should be close to zero. We choose these co-location points $\mathcal{R}$ by Latin hypercube sampling and only keep points for which no data is contained in the data set $\mathcal{U}$ [32]. We choose the number $N_{\mathcal{R}}$ of co-location points $\mathcal{R}$ based on the number $N_{\mathcal{U}}$ of non-empty bins in the data set $\mathcal{U}$:

$$N_{\mathcal{R}} = 0.5 N_{\mathcal{U}} \quad (34)$$

As training data spanning the state space $\mathcal{U}$ that the model is fit to, we use all frames up to just before the last pulse of the SBP in each recording. The frames of the last pulse of each recording are then used as testing data, i.e., as a reference to compare with the predicted frames of the data-driven model. We only include points at least 20 grid points away from the boundary to minimise boundary effects on the used data.

Besides this first simulation that can be directly and quantitatively compared to the reference for both data sets, we also run simulations using the data-driven models to compare them qualitatively with the reference model, i.e., the AP96 model for the synthetic data set, and the in-vitro model for the OVM data. As the second simulation, for both data-driven models, we run a similar rectangle-based protocol to the one used to stimulate spirals for the reference AP96 model (Section 2.1.2). Here, we use the same stimulation values for u , but $\tilde{u} = 0.6$ inside the second rectangle and $\tilde{u} = 0.1$ outside of it. Thirdly, we also run a one-dimensional cable simulation to obtain APD and CV restitution curves for the fit surrogate AP96 model with the same restitution stimulation protocol as outlined for the reference AP96 model in Section 2.1.2

3. Results

In this section, we present how the state space expansion and data-driven models as described in Section 2, deal with the two data sets. Since the results for the in-silico synthetic data are easier to interpret, we begin with those data.

3.1. Surrogate model for synthetic data

3.1.1. State space

For the data set from numerical simulation of the AP96 model (cf. Section 2.1.2), the three-dimensional state space spanned by $\mathbf{u} = [u, \tilde{u}, g]^T$ is quite sparse. 95 695 bins in this space contain at least one data point, which corresponds to 9% of all bins. Of these bins, 40% contain at least 10 data points. Non-empty bins contain 230 data points on average.

In Fig. 4, we present a three-dimensional view of the state space and two different two-dimensional projections of the state space; to u and g , and to u and $\tilde{u}$, respectively. Panel A shows that most data points lay

on a loop-like inertial manifold, while a large volume of the state space is still covered by data points (panel B). Colouring by the mean of $\dot{u}$, we can see that the two halves of the loop nicely distinguish between excitation and recovery, see the blue and red shades in panel C. Note that the colouring of the state space in panel D corresponds to the SD σ and not to the SD g . To be more specific, this panel is coloured by the SD σ of the samples of the time-derivative $\dot{u}$ in the bins, instead of the SD g of the signal u in the neighbourhoods around each point $\mathbf{x}$, which is one of the axes spanning the state space in Fig. 4. In large parts of these projections, the SD σ is much lower than the mean absolute value of $\dot{u}$, showing that the model's possible states are well-separated in the chosen state space $\mathbf{u}$.

3.1.2. Forward model

A model in the context of this work consists of a “colouring” of the entire three-dimensional state space representing the $\dot{u}$ at each possible state $\mathbf{u}$.

The data-driven surrogate model f is fit to the three-dimensional state space as outlined in Section 2.4 resulting in:

$$f(\mathbf{u}) = -0.22407 \tilde{u} g u - 0.82589 \tilde{u} g + 0.0054601 \tilde{u} u + 0.021655 \tilde{u} \\ + 0.31082 g u + 0.23077 g - 0.0081478 u - 0.009711 \quad (35)$$

Subsequently, a forward simulation is performed with the model, leading to a prediction of u based on the history until that point which can be compared to the reference evolution of u in Fig. 5. This simple model displays excitation and recovery, as well as propagation of the wave. The predicted evolution is similar to its reference: The values for the plateau-phase and resting state are close to the reference values. The depolarisation at the wave front is a bit too late, leading to predicted values that are too low. Likewise, the repolarisation at the wave back is too late, such that there the predicted values are too high.

Using the method by Bayly et al. [33], CV has been measured to be $(0.1403 \pm 0.0036) \text{ m s}^{-1}$ for the prediction, and $(0.1621 \pm 0.0065) \text{ m s}^{-1}$ for the reference. This lower propagation velocity leads to the observed delay of wave front and back.

We also computed the average APD of the activity in the central 60×60 pixels of the frame with the polarisation threshold of $u = 0.5$: For the prediction, we measure an APD of $(56.68 \pm 0.48) \text{ ms}$, which is slightly longer than the reference of $(47.25 \pm 0.44) \text{ ms}$.

At the boundary of the domain, the method is still able to compute the evolution well despite less neighbouring data points being used for the calculation of the SD g .

3.1.3. Prediction of a spiral wave

Using similar initial conditions for the data-driven surrogate AP96 model and its reference, we have run simulations leading to spiral waves. In Fig. 6, the resulting evolutions of u over space and time are juxtaposed. It can be seen that the overall wave dynamics are similar, despite the data-driven model only being fit to data from SBP focal waves, i.e., spherical waves growing from a small source region with different timing and amplitudes. The surrogate model predicts wave fronts and backs at similar scales as the reference. The timing of the predicted spiral is also similar to the reference. However, the prediction differs from the reference in the APD of the initial rotation, at the rotor core, and in memory effects. Such a memory effect can clearly be seen in the reference at $t = 150 \text{ ms}$: The recovery to the resting state is much faster in the rectangle where the second variable v of the reference model has been set to a higher value in the initial state of the simulation. This causes a slightly different evolution afterwards.

3.1.4. Restitution curves

We have also performed one-dimensional cable simulations with the same stimulation protocol for the data-driven surrogate model




Fig. 4. Visualisation of the three-dimensional state space for the AP96 data set and projections onto two dimensions. (A): By colouring the space by the number N of points in each bin, it can be seen that most points lay on a loop. (B): Colouring by $\log_{10} N$, we can see that still a large volume of the space is covered by data. (C): Looking at the mean of $\dot{u}$ at each bin in this projection, we can see that in one side of the loop, u increases while it decreases at the other side. (D): The standard deviation of $\dot{u}$ is indicative of whether the states have been chosen well.

compared to the rescaled AP96 model for reference. In Fig. 7, the resulting restitution curves of CV (panel A) and APD (panel B) as functions of the CL are shown. The overall shape of the curves for the low-order surrogate model is similar to the reference. However, the predicted CVs are lower than the reference values.

3.2. Data-driven model for optical mapping data

3.2.1. State space

For the OVM data set of hiAMs (cf. Section 2.1.1), the three-dimensional state space spanned by $\mathbf{u} = [u, \tilde{u}, g]^T$ is also sparse:




Fig. 5. Prediction of the surrogate model over time compared to the reference solution for the AP96 data set. The difference between the prediction and reference is displayed in the bottom row.




Fig. 6. Spiral wave resulting from similar initial conditions for the data-driven surrogate model for the AP96 data set compared to its reference model. The data-driven model can predict the shape of the spiral wave despite only focal waves being included in the training data it is fit to.




Fig. 7. CV and APD restitution curves for the reference AP96 model and its data-driven surrogate model.




Fig. 8. Visualisation of the three-dimensional state space for the OVM data set and projections onto two dimensions in the same style as in Fig. 4.

152417 bins in this space contain at least one data point, which corresponds to just 15 % of all bins. Of these bins, 71 % contain at least 10 data points. Non-empty bins on average contain 949 data points.

As can be seen in Fig. 8, for these data, we are also observing a loop-like inertial manifold with two halves as in the other data set.

3.2.2. Forward model

For the OVM data set, the fit polynomial for the data-driven model is:

$$f(\mathbf{u}) = -0.41392 \tilde{u}g u - 0.31842 \tilde{u}g + 0.026308 \tilde{u}u + 0.0010949 \tilde{u} + 0.19341 g u + 0.15352 g - 0.014174 u - 0.0044818 \quad (36)$$

The prediction by the forward model for the OVM data set is depicted in Fig. 9. It models the overall behaviour of the tissue, i.e., excitation, recovery, and wave propagation, at similar time scales: The APD for an excitation threshold of $u = 0.5$ of (140.1 ± 5.1) ms for the

prediction is longer than the APD of (127 ± 36) ms measured for the reference. The CV for the prediction is measured at $(0.148 \pm 0.012) \text{ m s}^{-1}$, and at $(0.21 \pm 0.13) \text{ m s}^{-1}$ for the reference. The large uncertainty in the APD and CV measurements for the reference data is due to small regions of inhomogeneity in the recording. Due to the uncertainty and disparity in APD and CV values between prediction and reference, no reliable restitution curves can be measured for this model.

It can be seen that such artefacts due to inhomogeneity in the monolayer and other more complex effects are difficult to predict for the model, which only gets the state $\mathbf{u}(t_0, \mathbf{x})$ at the initial time as input. Yet, similar structures are found in the prediction.

For the OVM data, the repolarisation at the wave back is much slower than the depolarisation at the wave front. For the synthetic data set AP96 re- and depolarisation happen at roughly the same time scale. Still, the one EMA variable $\tilde{u}$ is enough to capture these processes in both cases.




Fig. 9. Prediction of the data-driven model over time compared to the reference solution for the OVM data set. The difference between the prediction and reference is displayed in the bottom row.




Fig. 10. Comparison of a spiral wave for the data-driven model for the OVM data set for hiAMs to a OVM recording of hiAMs showing a stable rotor. The data-driven model can predict the overall dynamics of the spiral wave despite only focal waves being included in the training data it is fit to.

3.2.3. Prediction of a spiral wave

In Fig. 10, we compare a rotor in another OVM recording of hiAMs and a rotor predicted by the in-silico data-driven model. The model has been fit to data from only focal waves. To stimulate a spiral wave, we use the rectangle-based initial conditions outlined in Section 2.1.2. Like for the synthetic data set (AP96, Section 3.1), the tissue model is able to predict the overall spiral dynamics. Note that the spiral wave in the OVM recording has stabilised for several minutes before $t = 0$ ms, while the rotor is just forming in the simulation with the data-driven model. The APD predicted by the model is much longer than the one observed APD by OVM. This leads to the wave front running into the wave back in the simulation.

4. Discussion

In this paper, we have shown that enough information can be extracted from just one spatio-temporal variable to build a state space, which enables the creation of in-silico models that can predict the further dynamics of the system. In layman's terms, from a single movie showing the behaviour of excitation and recovery, a predictive model can be constructed. Just the information from OVM is already enough to fit a data-driven model to the dynamics. With our approach, it is even possible to fit a stable model to a single excitation wave, which, in turn, does not generalise well to other waves. For this reason, the SBP protocol is designed to ensure that less likely stimulation patterns are observed in the data, too. The resulting models are able to replicate the excitation patterns of focal waves included in the training data and can even predict the dynamics of spiral waves, a pattern that the models are not trained on.

While the variables with which we augment the state space should be seen as preliminary and not optimal choices, the models created with these variables already are able to capture crucial behaviour of the tissue, namely excitation, recovery and wave propagation. Both short and long-term memory can be encoded with the use of EMAs as state variables, which can easily be updated on the fly (Eq. (9)). It remains to be seen if more or other state space augmentations can be used to be able to replicate higher-order effects, such as wave front curvature and to obtain more accurate restitution behaviour.

Through state space expansion with EMAs and SDs, we have found a way to encode useful information to describe the mechanics of excitation waves. It is remarkable, that, after “unfolding” with the proper state space expansion, we are able to capture with simple polynomial regression the kinetics of excitation waves sufficiently in the sense that we observe excitation and recovery with an APD similar to the data, and wave propagation at a similar CV.

The polynomial model f is fit using a basic supervised machine learning method. However, this is just a small part of the method presented in this work. We have shown that using unsupervised learning, i.e., data transformations, much insight can be gained about the structure of the data. The insights obtained from presenting the data sets in the chosen three-dimensional state spaces as done in this work make it almost trivial to fit the data-driven model to the data. Fitting a function in the chosen state space is easier to interpret than black box approaches common in machine learning.

We have chosen to use an advection-based description of the wave propagation. This is due to the computation of the diffusion term $\nabla \cdot D \nabla u$ in Eq. (17) being quite sensitive to small variations in u and hence less robust to noisy data. The absolute gradient $\|\nabla u\|$ as approximated

using the SD g can deal with noise much better by design. Solving the forward equations in the advection-regression framework requires less computational effort in forward simulations than most traditional in-silico models and other novel approaches utilising convolutional and physics-informed neural networks. In the numerical solution of reaction–diffusion systems (Eq. (17)), a sufficiently small numerical time step must be chosen according to a stability criterion. Typically, this leads to numerical time steps of the order of magnitude 0.1 ms in two-dimensional simulations of in-silico tissue models of cardiac electrophysiology. Forward Euler stepping our model in the advection-regression framework can be done at much larger time steps of up to 10 ms, with the trade-off of fitting less to the original data. Also the computation of the polynomial f , EMAs and SDs result in an overall computationally cheaper model than typical tissue models.

While this was not needed for this use case, in further research, more sophisticated general function approximators may be tried for f , such as neural networks. Although the Latin hypercube sampling enables some basic regularisation of the solution, it remains to be seen how to keep the fit functions simple enough such that they generalise well. Another limitation that could be addressed when using more flexible function approximators is that, in some cases, the polynomial model f updates the state u such that it tries to leave the valid range.

While other augmentations were tried in the development of this method, such as pseudo-electrograms, the last APD and diastolic interval at each point over time, or the absolute value of the gradient fit to a neighbourhood of each point using the least-squares method, we are here presenting two augmentations we found most useful, the EMA $\bar{u}$ as well as the SD g .

We have not yet taken into account boundary effects in the model creation procedure. Some information about the gradient of the main variable u is encoded in the SD g in the neighbourhood of each point in the medium. The SD g can still be computed near the boundary by treating those points just like interior points but with fewer samples. How to encode boundary conditions in this model could be investigated in follow-up studies.

The procedure to generate in-silico tissue models is designed to work well in a setting where not much is known about the wave propagation in an in-vitro tissue model. Our method only requires very little assumptions to uncover the dynamics. Fitting is not limited to just in-vitro OVM data of focal waves in cardiac tissue excitation models, but can also be done in general where excitation wave dynamics can be observed in two dimensional space over time. An advantage of our method is that models for any observed behaviour can be created and fit to data which is in contrast to classical phenomenological models [6–12]. Fitting parameters of such existing models can get the model closer to observed data. However, if a behaviour falls outside the regime of these models, such methods will not work [23]. With our method, a data-driven model can be fit flexibly with relative ease within minutes and subsequently used to predict the further evolution of the tissue. This puts our model creation pipeline more in line with recent deep learning approaches [22–24]. However, our approach is typically faster in fitting due to the polynomial model function and faster in forward stepping because of the SD-based spatial coupling. Depending on the complexity of the dynamics, a simple approach to modelling such as ours may be enough. Promising use cases to apply the proposed method with minimal changes are the peristaltic motion of for instance embryonic hearts or the digestive tract [34], or the spread of lichens which has been observed to form spiral waves [35]. For more complex dynamics, more than three state space variables may be needed, or an extension using a neural network as the model function f . For very complex dynamics, the aforementioned, more extensive deep learning methods may be necessary [6–12]. Possible extensions of our data-driven model include an online-learning approach, where the model's predictive power is continuously improved while new data comes in. Such a model could, for instance, be fit to patient recordings, such as intra-cardiac electrograms or electrocardiograms on the body surface, and then be used to predict cardiac arrhythmias before they happen.

5. Conclusion

In this work, we presented a method to create simple data-driven in-silico tissue models from in-vitro OVM data and synthetic data that capture the basic behaviour of the corresponding excitable media. These models can predict the activation patterns included in the training data, i.e., only focal waves, and generalise to unseen activation patterns, such as spiral waves. A useful feature of this model creation method is its ability to extract all relevant information that the models are fit to from just one state variable over space and time, such as the transmembrane voltage of cardiomyocytes which can be obtained from a single experiment, e.g. an OVM recording.

CRediT authorship contribution statement

Desmond Kabus: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Tim De Coster:** Methodology, Writing – review & editing. **Antoine A.F. de Vries:** Conceptualization, Resources, Writing – review & editing. **Daniël A. Pijnappels:** Conceptualization, Funding acquisition, Project administration, Resources, Supervision, Writing – review & editing. **Hans Dierckx:** Conceptualization, Funding acquisition, Project administration, Resources, Supervision, Writing – review & editing.

Declaration of competing interest

None declared.

Data availability

The source code of the software implemented for this paper is publicly available at [GitLab.com/HeartKOR](https://gitlab.com/HeartKOR). The functionality is broken up into several separate packages, some of which depend on each other.

The model creation pipeline is implemented as the Distephym Python module, [distephym 1.0.0](#). It stores the augmented state spaces in sparse histograms, specifically implemented for this publication as another module, [sparsealgos1.0.0](#).

Reading, writing, and analysis of OVM recordings is implemented in the Sappho Python module, [sappho 1.0.0](#).

The Ithildin C++ framework, [ithildin 3.2.2](#), is a finite-differences, parallelised, reaction–diffusion solver written in C++. It will be released in a separate, dedicated publication [36,37]. We have used Ithildin to generate the synthetic (i.e., AP96) data set.

For the other forward simulations in this work, we have written another simple finite-differences reaction–diffusion solver in Python, [ithilmin 1.0.0](#), that can easily be combined with other Python modules. It was used for forward runs of the fit data-driven models.

In conjunction with this paper, we are also releasing a new version of the Ithildin Python module, [py_ithildin 0.6.1](#), [29]. This module was used for the processing of in-vitro/in-silico data sets. Results from both reaction–diffusion solvers, [ithildin](#) and [ithilmin](#), can be analysed with this Python module.

The source code of these projects and the AP96 and OVM data sets used in this work have been archived on Zenodo ([DOI: 10.5281/zenodo.8183722](https://doi.org/10.5281/zenodo.8183722)).

Acknowledgements

We are grateful to our collaborators at the LUMC, especially to Sven O. Dekker and Juan Zhang for providing access to hiAM monolayers and the OVM setup, and to Niels Harlaar for the spiral wave OVM recordings, respectively. Also, we would like to thank all members of Team HeartKOR at KULAK for valuable help during writing, in particular, Marie Cloet.

Funding

DK is supported by KU Leuven, Belgium grant GPUL/20/012. HD is supported by KU Leuven, Belgium grant STG/19/007. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. The hiAMs were developed as part of the research programme “More Knowledge with Fewer Animals” (MKMD, project 114022503, to AAFdV), which was financed by the Netherlands Organisation for Health Research and Development (ZonMw) and by the Dutch Society for the Replacement of Animal Testing (dsRAT).

References

- [1] Julia V. Noble, Geographic and temporal development of plagues, *Nature* 250 (5469) (1974) 726–729, <http://dx.doi.org/10.1038/250726a0>.
- [2] Paul C. Bressloff, Waves in neural media, in: *Lecture Notes on Mathematical Modelling in the Life Sciences*, Springer, 2014, pp. 18–19, <http://dx.doi.org/10.1007/978-1-4614-8866-8>.
- [3] R.H. Clayton, O. Bernus, E.M. Cherry, H. Dierckx, F.H. Fenton, L. Mirabella, A.V. Panfilov, F.B. Sachse, G. Seemann, H. Zhang, Models of cardiac tissue electrophysiology: Progress, challenges and open questions, *Prog. Biophys. Mol. Biol.* 104 (1–3) (2011) 22–48, <http://dx.doi.org/10.1016/j.biophys.2010.05.008>.
- [4] Steven A. Niederer, Joost Lumens, Natalia A. Trayanova, Computational models in cardiology, *Nat. Rev. Cardiol.* 16 (2) (2018) 100–111, <http://dx.doi.org/10.1038/s41569-018-0104-y>.
- [5] Natalia A. Trayanova, Ashish N. Doshi, Adityo Prakosa, How personalized heart modeling can help treatment of lethal arrhythmias: A focus on ventricular tachycardia ablation strategies in post-infarction patients, *WIREs Syst. Biol. Med.* 12 (3) (2020) <http://dx.doi.org/10.1002/wsbm.1477>, (ISSN 1939-5094, 1939-005X). URL <https://onlinelibrary.wiley.com/doi/10.1002/wsbm.1477>.
- [6] Richard FitzHugh, Impulses and physiological states in theoretical models of nerve membrane, *Biophys. J.* 1 (6) (1961) 445–466, [http://dx.doi.org/10.1016/s0006-3495\(61\)86902-6](http://dx.doi.org/10.1016/s0006-3495(61)86902-6).
- [7] J. Nagumo, S. Arimoto, S. Yoshizawa, An active pulse transmission line simulating nerve axon, *Proc. IRE* 50 (10) (1962) 2061–2070, <http://dx.doi.org/10.1109/jrproc.1962.288235>.
- [8] Dwight Barkley, A model for fast computer simulation of waves in excitable media, *Physica D* 49 (1–2) (1991) 61–70, [http://dx.doi.org/10.1016/0167-2789\(91\)90194-e](http://dx.doi.org/10.1016/0167-2789(91)90194-e).
- [9] Rubin R. Aliev, Alexander V. Panfilov, A simple two-variable model of cardiac excitation, *Chaos Solitons Fractals* 7 (3) (1996) 293–301, [http://dx.doi.org/10.1016/0960-0779\(95\)00089-5](http://dx.doi.org/10.1016/0960-0779(95)00089-5).
- [10] Flavio Fenton, Alain Karma, Vortex dynamics in three-dimensional continuous myocardium with fiber rotation: Filament instability and fibrillation, *Chaos* 8 (1) (1998) 20–47, <http://dx.doi.org/10.1063/1.166311>.
- [11] Alfonso Bueno-Orovio, Elizabeth M. Cherry, Flavio H. Fenton, Minimal model for human ventricular action potentials in tissue, *J. Theoret. Biol.* 253 (3) (2008) 544–560, <http://dx.doi.org/10.1016/j.jtbi.2008.03.029>.
- [12] Christopher D. Marcotte, Roman O. Grigoriev, Dynamical mechanism of atrial fibrillation: A topological approach, *Chaos* 27 (9) (2017) <http://dx.doi.org/10.1063/1.5003259>.
- [13] G.K. Moe, W.C. Rheinboldt, J.A. Abildskov, A computer model of atrial fibrillation, *Am. Heart J.* 67 (1964) 200–220.
- [14] Marc Courtemanche, Rafael J. Ramirez, Stanley Nattel, Ionic mechanisms underlying human atrial action potential properties: insights from a mathematical model, *Amer. J. Physiol. Heart Circul. Physiol.* 275 (1) (1998) H301–H321.
- [15] Michelangelo Paci, Jari Hyttinen, Katriina Aalto-Setälä, Stefano Severi, Computational models of ventricular-and atrial-like human induced pluripotent stem cell derived cardiomyocytes, *Ann. Biomed. Eng.* 41 (11) (2013) 2334–2348.
- [16] Rupamanjari Majumder, Wanchana Jangsangthong, Iolanda Feola, Dirk L. Ypey, Daniël A. Pijnappels, Alexander V. Panfilov, A mathematical model of neonatal rat atrial monolayers with constitutively active acetylcholine-mediated K⁺ current, *PLoS Comput. Biol.* 12 (6) (2016) e1004946.
- [17] Jeffrey J. Fox, Eberhard Bodenschatz, Robert F. Gilmour, Period-doubling instability and memory in cardiac tissue, *Phys. Rev. Lett.* 89 (13) (2002) <http://dx.doi.org/10.1103/physrevlett.89.138101>.
- [18] Ning Wei, Yoichiro Mori, Elena G. Tolkacheva, The role of short term memory and conduction velocity restitution in alternans formation, *J. Theoret. Biol.* 367 (2015) 21–28, <http://dx.doi.org/10.1016/j.jtbi.2014.11.014>.
- [19] Leo Priebe, Dirk J. Beuckelmann, Simulation study of cellular electric properties in heart failure, *Circ. Res.* 82 (11) (1998) 1206–1223, <http://dx.doi.org/10.1161/01.res.82.11.1206>.
- [20] Kirsten H.W.J. ten Tusscher, Denis Noble, Peter-John Noble, Alexander V. Panfilov, A model for human ventricular tissue, *Amer. J. Physiol. Heart Circul. Physiol.* 286 (4) (2004) H1573–H1589.
- [21] Karli Gillette, Matthias A.F. Gsell, Anton J. Prassl, Elias Karabelas, Ursula Reiter, Gert Reiter, Thomas Grandits, Christian Payer, Darko Štern, Martin Urschler, Jason D. Bayer, Christoph M. Augustin, Aurel Neic, Thomas Pock, Edward J. Vigmond, Gernot Plank, A framework for the generation of digital twins of cardiac electrophysiology from clinical 12-lead ECGs, *Med. Image Anal.* 71 (2021) 102080, <http://dx.doi.org/10.1016/j.media.2021.102080>.
- [22] Francisco Sahli Costabal, Yibo Yang, Paris Perdikaris, Daniel E. Hurtado, Ellen Kuhl, Physics-informed neural networks for cardiac activation mapping, *Front. Phys.* 8 (2020) <http://dx.doi.org/10.3389/fphy.2020.00042>.
- [23] Clara Herrero Martin, Alon Oved, Rasheda A. Chowdhury, Elisabeth Ullmann, Nicholas S. Peters, Anil A. Bharath, Marta Varela, EP-PINNs: Cardiac electrophysiology characterisation using physics-informed neural networks, *Front. Cardiovasc. Med.* 8 (2022) <http://dx.doi.org/10.3389/fcvm.2021.768419>.
- [24] Shahrokh Shahi, Flavio H. Fenton, Elizabeth M. Cherry, A machine-learning approach for long-term prediction of experimental cardiac action potential time series using an autoencoder and echo state networks, *Chaos* 32 (6) (2022) <http://dx.doi.org/10.1063/5.0087812>.
- [25] Niels Harlaar, Sven O. Dekker, Juan Zhang, Rebecca R. Snabel, Marieke W. Veldkamp, Arie O. Verkerk, Carla Cofino Fabres, Verena Schwach, Lente J.S. Lerink, Mathilde R. Rivaud, Aat A. Mulder, Willem E. Corver, Marie José T.H. Goumans, Dobromir Dobrev, Robert J.M. Klautz, Martin J. Schalij, Gert Jan C. Veenstra, Robert Passier, Thomas J. van Brakel, Daniël A. Pijnappels, Antoine A.F. de Vries, Conditional immortalization of human atrial myocytes for the generation of in vitro models of atrial fibrillation, *Nat. Biomed. Eng.* 6 (4) (2022) 389–402, <http://dx.doi.org/10.1038/s41551-021-00827-5>.
- [26] Siyavash Joukar, A comparative review on heart ion channels, action potentials and electrocardiogram in rodents and human: extrapolation of experimental insights to clinic, *Lab. Anim. Res.* 37 (1) (2021) 1–15, <http://dx.doi.org/10.1186/s42826-021-00102-3>.
- [27] Thomas O'Hara, Yoram Rudy, Quantitative comparison of cardiac ventricular myocyte electrophysiology and response to drugs in human and nonhuman species, *Amer. J. Physiol. Heart Circul. Physiol.* 302 (5) (2012) H1023–H1030, <http://dx.doi.org/10.1152/ajpheart.00785.2011>.
- [28] Trine Krogh-Madsen, Peter Schaffer, Anne D. Skriver, Louise Kold Taylor, Brigitte Pelzmann, Bernd Koidl, Michael R. Guevara, An ionic model for rhythmic activity in small clusters of embryonic chick ventricular cells, *Amer. J. Physiol. Heart Circul. Physiol.* 289 (1) (2005) H398–H413, <http://dx.doi.org/10.1152/ajpheart.00683.2004>.
- [29] Desmond Kabus, Louise Arno, Lore Leenknegt, Alexander V. Panfilov, Hans Dierckx, Numerical methods for the detection of phase defect structures in excitable media, *PLoS One* 17 (7) (2022) e0271351, <http://dx.doi.org/10.1371/journal.pone.0271351>.
- [30] Richard A. Gray, José Jalife, Alexandre V. Panfilov, William T. Baxter, Cándido Cabo, Jorge M. Davidenko, Arkady M. Pertsov, Mechanisms of cardiac fibrillation, *Science* 270 (5239) (1995) 1222–1223, <http://dx.doi.org/10.1126/science.270.5239.1222>.
- [31] Pawel Kuklik, Stef Zeemering, Bart Maesen, Jos Maessen, Harry J. Crijns, Sander Verheule, Anand N. Ganesan, Ulrich Schotten, Reconstruction of instantaneous phase of unipolar atrial contact electrogram using a concept of sinusoidal recombination and Hilbert transform, *IEEE Trans. Biomed. Eng.* 62 (1) (2015) 296–302, <http://dx.doi.org/10.1109/tbme.2014.2350029>.
- [32] Michael D. McKay, Richard J. Beckman, William J. Conover, A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics* 42 (1) (2000) 55–61, Publisher: Taylor & Francis.
- [33] Philip V. Bayly, Bruce H. KenKnight, Jack M. Rogers, Russel E. Hillsley, Raymond E. Ideker, William M. Smith, Estimation of conduction velocity vector fields from epicardial mapping data, *IEEE Trans. Biomed. Eng.* 45 (5) (1998) 563–571.
- [34] Antoon F.M. Moorman, Vincent M. Christoffels, Cardiac chamber formation: development, genes, and evolution, *Physiol. Rev.* (2003) <http://dx.doi.org/10.1152/physrev.00006.2003>.
- [35] Stefan C. Müller, Kinko Tsuji, Appearance in nature, in: *Spirals and Vortices: In Culture, Nature, and Science*, Springer, 2019, pp. 31–66, http://dx.doi.org/10.1007/978-3-030-05798-5_2.
- [36] Marie Cloet, Louise Arno, Desmond Kabus, Joeri Van der Veken, Alexander V. Panfilov, Hans Dierckx, Scroll waves and filaments in excitable media of higher spatial dimension, 2023, <http://dx.doi.org/10.48550/arxiv.2304.14861>, arXiv preprint [arXiv:2304.14861](https://arxiv.org/abs/2304.14861). URL <https://arxiv.org/abs/2304.14861>.
- [37] Desmond Kabus, Marie Cloet, Christian Zemlin, Hans Dierckx, The ithidin anisotropic reaction-diffusion solver for excitable media, 2023, forthcoming.