



OPEN

Calibrating cardiac electrophysiology models using latent Gaussian processes on atrial manifolds

Sam Coveney^{1✉}, Caroline H. Roney², Cesare Corrado³, Richard D. Wilkinson⁴, Jeremy E. Oakley⁵, Steven A. Niederer³ & Richard H. Clayton⁶

Models of electrical excitation and recovery in the heart have become increasingly detailed, but have yet to be used routinely in the clinical setting to guide personalized intervention in patients. One of the main challenges is calibrating models from the limited measurements that can be made in a patient during a standard clinical procedure. In this work, we propose a novel framework for the probabilistic calibration of electrophysiology parameters on the left atrium of the heart using local measurements of cardiac excitability. Parameter fields are represented as Gaussian processes on manifolds and are linked to measurements via surrogate functions that map from local parameter values to measurements. The posterior distribution of parameter fields is then obtained. We show that our method can recover parameter fields used to generate localised synthetic measurements of effective refractory period. Our methodology is applicable to other measurement types collected with clinical protocols, and more generally for calibration where model parameters vary over a manifold.

Mechanical contraction of the heart is initiated and synchronised by a travelling wave of electrical excitation and recovery that arises spontaneously in the natural pacemaker. The heart is made up of four chambers: the ventricles pump blood to the body and lungs, while the atria act as reservoirs and primers for the ventricles. A cardiac arrhythmia is a disturbance of regular heart rhythm resulting in a rapid, slow, or irregular rhythm. Atrial fibrillation (AF) is a common and increasingly prevalent cardiac arrhythmia¹. AF can be sustained by re-entry, where electrical activation continually propagates into recovering tissue, creating a self-sustaining rotating wave². Radio-frequency catheter ablation can be used to disrupt re-entrant circuits that act to sustain AF, but is not always effective³.

Two properties of cardiac tissue are important for the development of sustained re-entry, and these properties vary across atrial tissue. Conduction velocity (CV) describes the speed at which an activation wave spreads. The effective refractory period (ERP) is the minimum time interval between two successive stimuli that allows two activation waves to propagate and is related to action potential duration (APD), which is the interval between local activation (depolarization) and recovery (repolarization). Both CV and ERP decrease at shorter pacing intervals, and this dynamic behaviour and its spatial heterogeneity is important for determining the stability of re-entry^{4,5} as well as the complex paths followed by electrical activation during AF⁶. Natural variability in the speed of the excitation wave and the dynamics of excitation and recovery exist both between individuals and within the heart of a single individual^{7,8}. Cardiac tissue exhibits spatial heterogeneity with differences in ion channel conductances, gap junction distributions, and fibrotic remodelling across the heart⁹. These spatial heterogeneities in structural and functional properties lead to heterogeneity in ERP. The resulting dispersion in repolarisation properties is a mechanism for focal arrhythmia initiation¹⁰, and atrial fibrillation initiation through increasing vulnerability to re-entry^{8,11}.

Electrophysiology (EP) models describe how electrical activation diffuses through cardiac tissue. Local activation and recovery are represented by a set of differential equations describing a reaction-diffusion system

¹Leeds Institute of Cardiac and Metabolic Medicine, University of Leeds, Leeds LS2 9JT, UK. ²School of Engineering and Materials Science, Queen Mary University of London, London E1 4NS, UK. ³Division of Imaging Sciences and Biomedical Engineering, King's College London, London WC2R 2LS, UK. ⁴School of Mathematical Sciences, University of Nottingham, Nottingham NG7 2RD, UK. ⁵School of Mathematics and Statistics, University of Sheffield, Sheffield S10 2TN, UK. ⁶Insigneo Institute for In-Silico Medicine and Department of Computer Science, University of Sheffield, Sheffield S10 2TN, UK. ✉email: coveney.sam@gmail.com

that models tissue-scale propagation of activation and cellular activation and recovery^{12,13}. Models of cardiac electrical activation have become valuable research tools, but are also beginning to be used in the clinical setting to guide interventions in patients^{14,15}. These applications require personalised models of both anatomy and electrophysiology to be constructed. Personalised anatomical models can be assembled from medical images, and statistical shape models enable the assessment of varying shape on electrical behaviour¹⁶. Calibration of EP models is difficult because of the limited measurements that can be made routinely in the clinical setting.

EP model parameters determine model behaviour and for a personalised model should be calibrated to reconstruct the heterogeneity in CV and ERP, as well as their dynamic behaviour, in the heart of a specific patient. Measurements of local activation time (LAT), which measures the time of arrival of the activation wavefront relative to the timing of a pacing stimulus, enable reconstruction of heterogeneous CV for pacing at a fixed rate^{17,18}. Calibration to the *dynamics* of activation and recovery is more challenging. Both the quantity and type of data that can be recorded from patients are constrained by the clinical procedure, so it is difficult to determine spatial heterogeneity of repolarisation. An S1S2 pacing protocol can be used to measure restitution curves. The heart is paced for several beats at an initial pacing cycle length S1, followed by a stimulus with a shorter length S2. This protocol is repeated for different values of the S2 interval, and the shortest S2 that can elicit an activation indicates an upper bound for ERP at the stimulus site. While models can be calibrated to reconstruct CV(S2) restitution and ERP from LAT measurements with an S1S2 protocol^{19,20}, recent work raises doubts over whether model parameters can be identified uniquely from these types of measurement alone²¹. Biophysically detailed models of electrical activation have large numbers of parameters and many of these may be unidentifiable from restitution curve data²². There is a need for robust approaches that can interrogate cardiac tissue properties more thoroughly while at the same time minimising additional interventions.

In this paper, we present a novel method for probabilistic calibration of an electrophysiology simulator from spatially sparse measurements using a probabilistic model of electrophysiology parameters on a manifold representing the left atrium of the heart. We focus on estimating parameter fields that reconstruct heterogeneity in ERP. We chose to use a phenomenological EP model that captures the main features of cardiac activation and recovery. We determine two types of ERP measurements for calibrating EP parameters that determine excitability. EP parameters are modelled as latent Gaussian processes (GPs) on a manifold, and linked to observations via surrogate functions and a likelihood function designed for ERP measurements. We use Markov Chain Monte Carlo (MCMC) to obtain the posterior distribution of EP parameter fields across the atrium. We validate our method quantitatively by generating ground truths and calibrating to sparse data. The principles behind our method generalise to other measurement types, such as CV and APD restitution data, making our approach a step forward in the creation of digital twins capable of reproducing the complex dynamics of electrophysiology.

Results

Workflow. The computational model, or ‘simulator’, that we seek to calibrate is composed of (i) a finite element mesh representing an atrial manifold $\mathbf{x} \in \Omega$; (ii) an electrophysiology model, that maps EP parameters $\theta_l(\mathbf{x})$, $l = 1, 2, \dots$ defined on the computational mesh to observable quantities; and (iii) a numerical solver for running EP simulations. Details on obtaining and processing a mesh for suitability in electrophysiology simulations are given in “Methods”, including details on the example mesh used here. The EP model is the modified Mitchell-Schaeffer (mMS) model^{23,24}, the parameters of which are effectively time-constants representing different phases of the action potential. We parameterize the mMS model with the following 5 parameters: $CV_{max}(\mathbf{x})$, $\tau_{in}(\mathbf{x})$, $\tau_{out}(\mathbf{x})$, $\tau_{open}(\mathbf{x})$, $APD_{max}(\mathbf{x})$. See “Methods” for details on the simulation model, parameterisation, and allowable parameter ranges, and details on the numerical implementation. We use the software *open-CARP*²⁵ to solve the mono-domain model for our simulations.

The main task in this work is to calibrate the simulator by inferring parameter fields $\theta_l(\mathbf{x})$ from ERP measurements. Figure 1 represents our modeling workflow, which we summarize here. Our code is available in a Zenodo repository²⁶.

Surrogate functions. The simulator can be used to map parameter fields to ERP fields $ERP(\mathbf{x}) = \text{sim}(\theta_l(\mathbf{x}), l = 1, 2, \dots)$. Given ERP observations at multiple locations on the atrial mesh, as well as an appropriate likelihood model for these observations, the simulator could be used in an MCMC setting in order to calibrate the parameter fields by obtaining samples from the posterior distribution for the EP fields. However, this is an extremely inefficient approach, since ERP depends only on local (rather than remote) tissue properties. We utilize a surrogate function (also called an “emulator”) solution in which we learn the mapping from parameters to ERP. This surrogate function allows us to predict ERP at location \mathbf{x} as $ERP(\mathbf{x}) = f(\theta_l(\mathbf{x}), l = 1, 2, \dots)$, bypassing the need to run the simulator directly for inference.

Gaussian process priors. The mesh has $\approx 10^5$ vertices for which parameters need to be defined, but ERP measurements are restricted to a subset of these vertices, with number of observations on the order of 10^0 – 10^1 . The electrophysiology parameter fields must be assumed to have low-rank structure, induced by spatial correlation, in order to make inferences about EP parameter values at locations other than ERP observation locations. This is achieved here by modeling the EP parameters using latent Gaussian process (GP) priors $\theta_l(\mathbf{x}) \sim \mathcal{GP}$. We use Gaussian Process Manifold Interpolation (GPMI), a method we proposed for defining Gaussian process (GP) distributions on manifolds²⁷. The approach uses solutions $\{\lambda_k, \phi_k(\mathbf{x})\}$ of the Laplacian (Laplace-Beltrami) eigenproblem on the mesh²⁸.

Bayesian calibration. We perform probabilistic calibration with MCMC to obtain the posterior distribution of latent variables in the GPs. We utilize a likelihood function that we developed specifically for ERP measure-

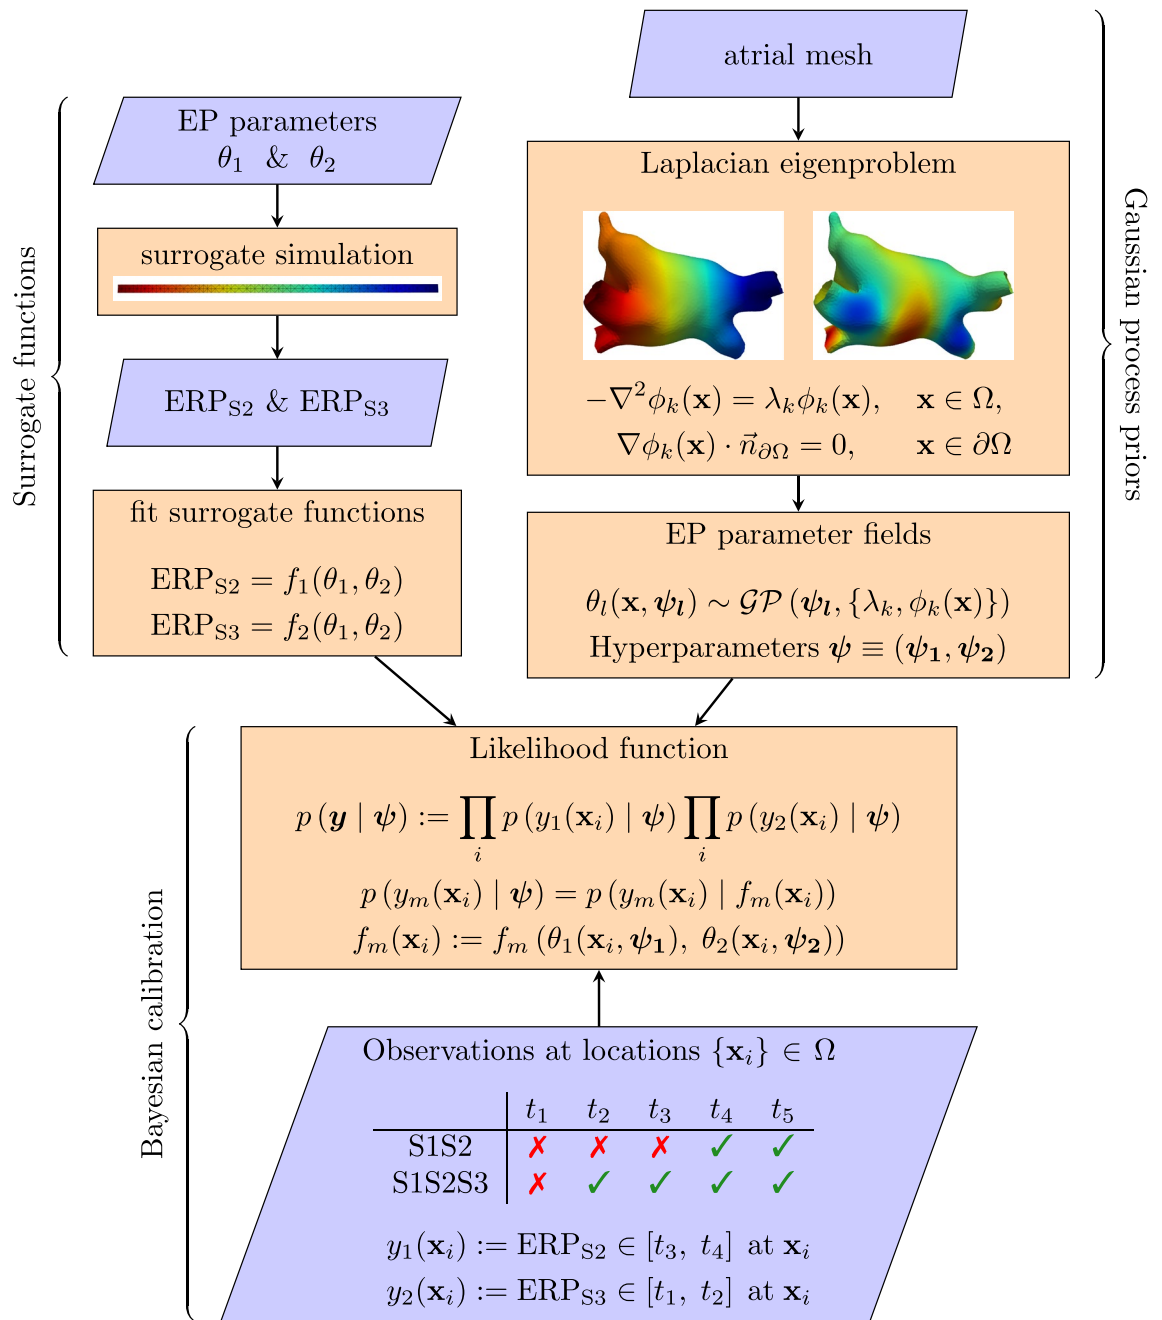
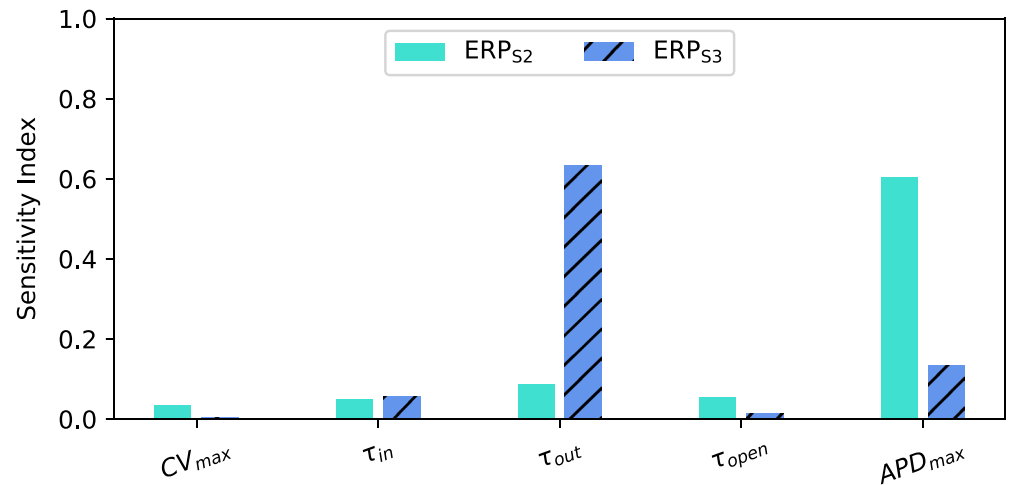


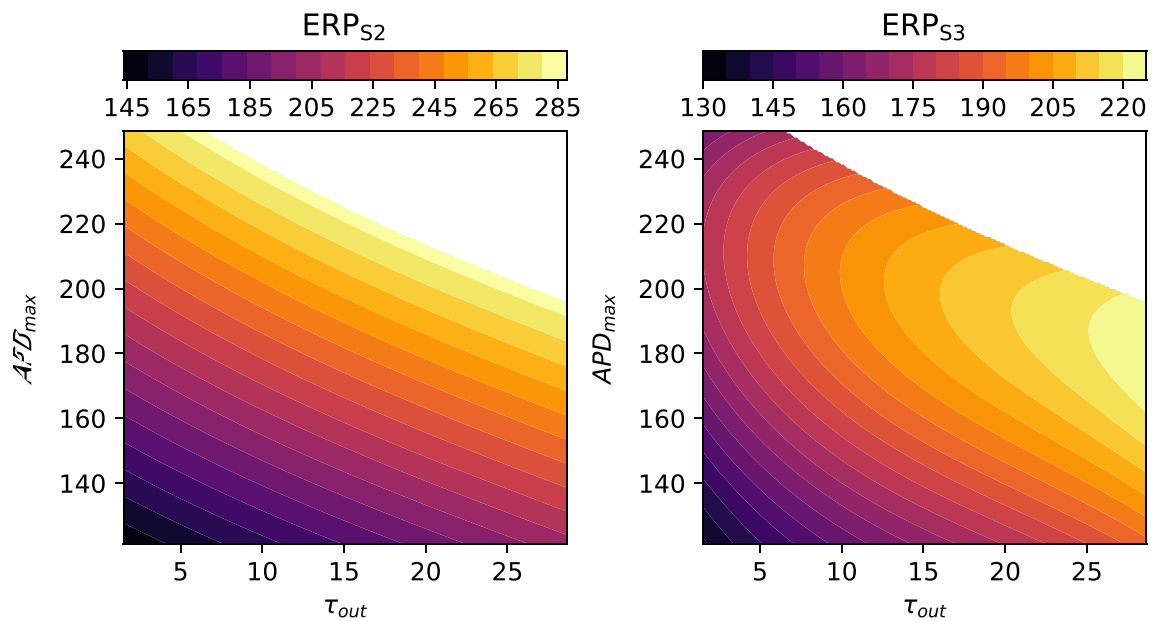
Figure 1. Workflow for using ERP observations to calibrate model parameters. A simplified ‘surrogate simulation’ consisting of a strip of tissue paced from one end is used to determine ERP values for a design of parameters. Surrogate functions are fit to these data, allowing rapid prediction of ERP from model parameters. Electrophysiology parameter fields are modelled as Gaussian processes on the atrial manifold, using a reduced-rank formulation relying on eigenfunctions and eigenvalues of the Laplace-Beltrami operator on the atrial mesh. ERP measurements using an S1S2 (or S1S2S3) protocol measure whether successful activation occurs for S2 (or S3) times $t_1 < t_2 < t_3 \dots$ etc. Using a likelihood function designed for ERP measurements and the surrogate functions, the likelihood can be evaluated given hyperparameters for the GP models. This allows for probabilistic calibration by obtaining the posterior distribution of hyperparameters using MCMC.

ments, which accounts for how an S1S2 pacing protocol to determine ERP effectively measures the S2 interval in which ERP lies, rather than measuring ERP directly.

Sensitivity analysis and surrogate functions. Figure 2a shows sensitivity indices for two types of ERP: an ERP measurement for S1S2 with S1 600 ms, denoted here as ERP_{S2}, and another type of ERP measurement for S1S2S3 pacing for S1 600 ms S2 300 ms, denoted here as ERP_{S3}. The S1S2S3 protocol, consisting of N S1 beats, 1 S2 beat, and 1 S3 beat, is introduced in this paper. We have determined that these measurements can be



(a) Sensitivity analysis of ERP_{S2} and ERP_{S3} for mMS parameters.



(b) Surrogate functions of ERP_{S2} and ERP_{S3} as a function of τ_{out} and APD_{max} .

Figure 2. (a) Sensitivity analysis of ERP for the mMS electrophysiology model for all parameters. τ_{out} and APD_{max} have the largest sensitivity indices for both types of ERP. (b) Surrogate functions for predicting ERP from τ_{out} and APD_{max} . A discontinuity in ERP_{S3} occurs at ERP_{S2} \approx 285 ms, so this region is shown in white.

used to calibrate EP parameters sufficiently to reproduce not only these ERP measurements, but also the time required for the action potential to reach various levels of repolarization recovery (e.g. APD_{20} and APD_{90} , the time required for 20% and 90% recovery). It is a key finding that the S1S2S3 protocol can be used (alongside the standard S1S2 protocol) to disentangle the contributions of separate parts of the action potential to the value of ERP, without needing to measure the action potential directly.

The sensitivity indices in Fig. 2a show that these ERP measurements are mainly determined by τ_{out} and APD_{max} , which approximately correspond to the duration of the repolarization and plateau phases of the action potential respectively. Calibration of other parameters, which determine some aspects of the shape of restitution curves but which do not strongly impact ERP, require both CV and APD restitution curve data from an S1S2 protocol²¹. For this reason, we have determined to use ERP to calibrate $\theta_1 \equiv \tau_{out}$ and $\theta_2 \equiv APD_{max}$. Figure 2b shows contour plots of the surrogate functions for ERP. A discontinuity occurs in the ERP_{S3} surface for parameters combinations resulting in ERP_{S2} > 285 ms, so data for ERP_{S2} > 280 ms were discarded before fitting this function. Note that the majority of clinical ERP_{S2} measurements fall in the range 170–270 ms, so even 280 ms could be considered as an upper limit²⁹.

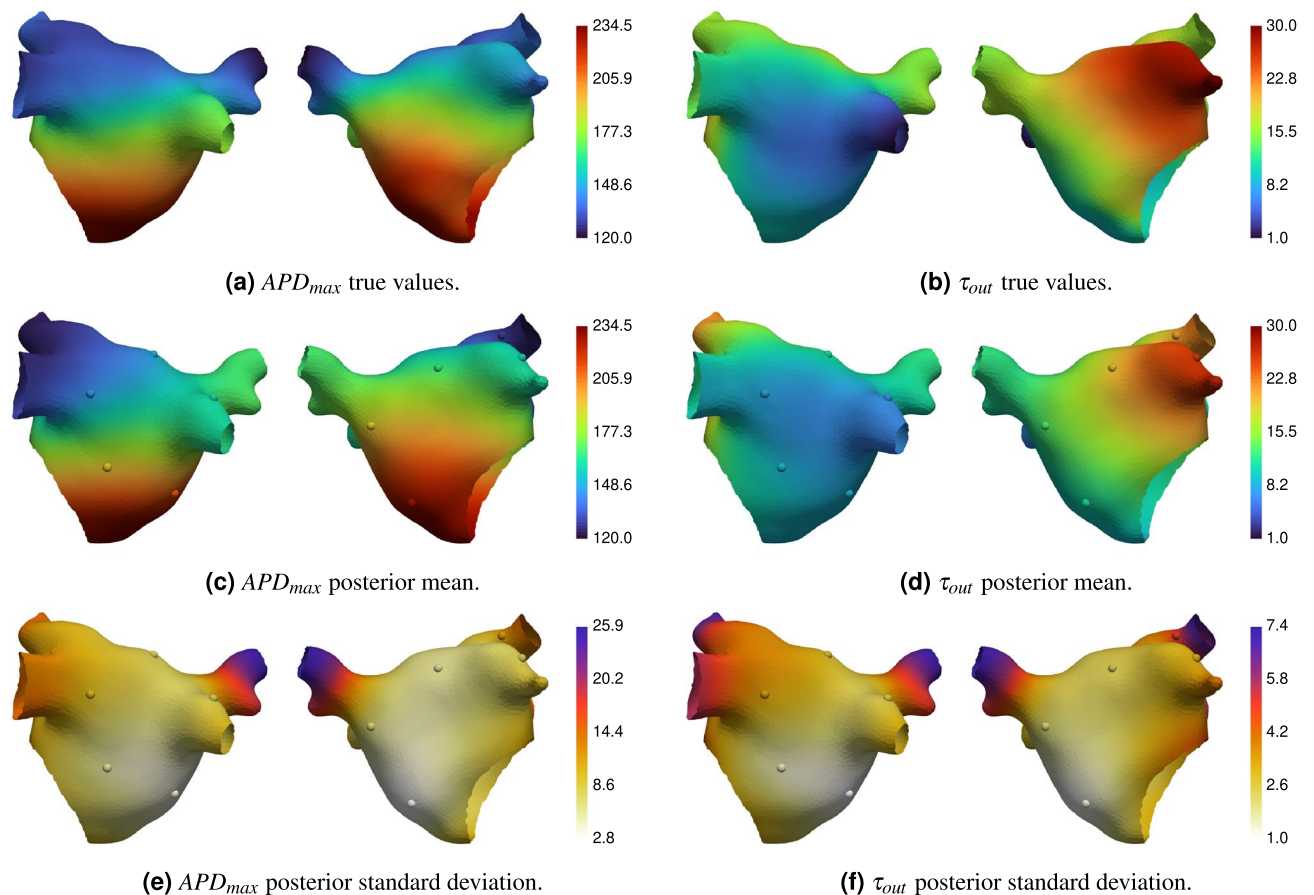


Figure 3. Ground truth and predicted (posterior mean and standard deviation) electrophysiology parameter fields. The spheres show the location of ERP measurements. All units are milliseconds.

Synthetic experiments. To test our methodology, we ran synthetic experiments as detailed in “Methods”. We used a left atrial mesh generated from a scan of an individual performed at St Thomas’ Hospital (see “Methods” for details). We created ground truth parameter fields for τ_{out} and APD_{max} in order to verify our calibration approach. We used 10 measurement locations, placed at random using a maximin design that excluded sites close to the mesh boundaries. The resolution of the S1S2 and S1S2S3 protocol was set to 10 ms. We used 24 eigenfunctions for representing each of the two parameter fields in Eq. (5), which we found to be sufficient to capture spatial variation while allowing good posterior sampling. For MCMC we used 5000 iterations, using 8 chains, discarding the first 50% of the samples as ‘burn-in’, and randomly thinning the remaining samples by a factor of 100 to give 200 posterior samples.

Figure 3 shows the true parameter fields, and the posterior mean and standard deviation of the calibrated parameter fields. Figure 4 shows the true ERP fields, the posterior mean and standard deviation of the ERP fields (calculated from ERP samples, which are calculated from the parameter field posterior samples), and the Independent Standard Errors (ISE) of ERP (the absolute difference between true and posterior mean, divided by the posterior standard deviation). Measurement locations are shown as spheres in Figs. 3 and 4, colored by the corresponding values at each location. Figure 5 shows the APD simulation results from the atrial simulator using the ground truth parameter fields and the posterior mean of the calibrated parameter fields.

The prediction of EP parameter fields τ_{out} and APD_{max} and ERP fields ERP_{S2} and ERP_{S3} captures the ground truth extremely well. Predictions on the pulmonary veins, which are effectively regions of extrapolation, deviate from the ground truth more than other regions on the main body of the atrium. These deviations are on the order of the S2 and S3 resolution, and the posterior variance is higher in these regions. Uncertainty increases with distance from the measurement locations. The ISE scores show that the *distribution* of ERP predicted by the model covers the ground truth well as nearly all values are less than 3. The ISE for ERP_{S2} on the left atrial appendage is above 3, which may be caused by a combination of high ground-truth values for τ_{out} (which are not effectively probed by the measurements) and insufficient basis functions to capture high spatial variation in this region of the mesh. APD from the full atrial simulator using the posterior mean of the parameters (the maximum *a posteriori* estimate could have been used instead) matches the ground truth values very closely, demonstrating that the action potential has been calibrated well using only ERP measurements.

We also performed quantitative validation across a broad range of designs. Figure 6 shows these validation results, for different configurations of the S1S2(S3) pacing protocol (number of ERP observations, resolution of S2 and S3 intervals) and different heterogeneity for ERP, controlled by different correlation lengthscales for

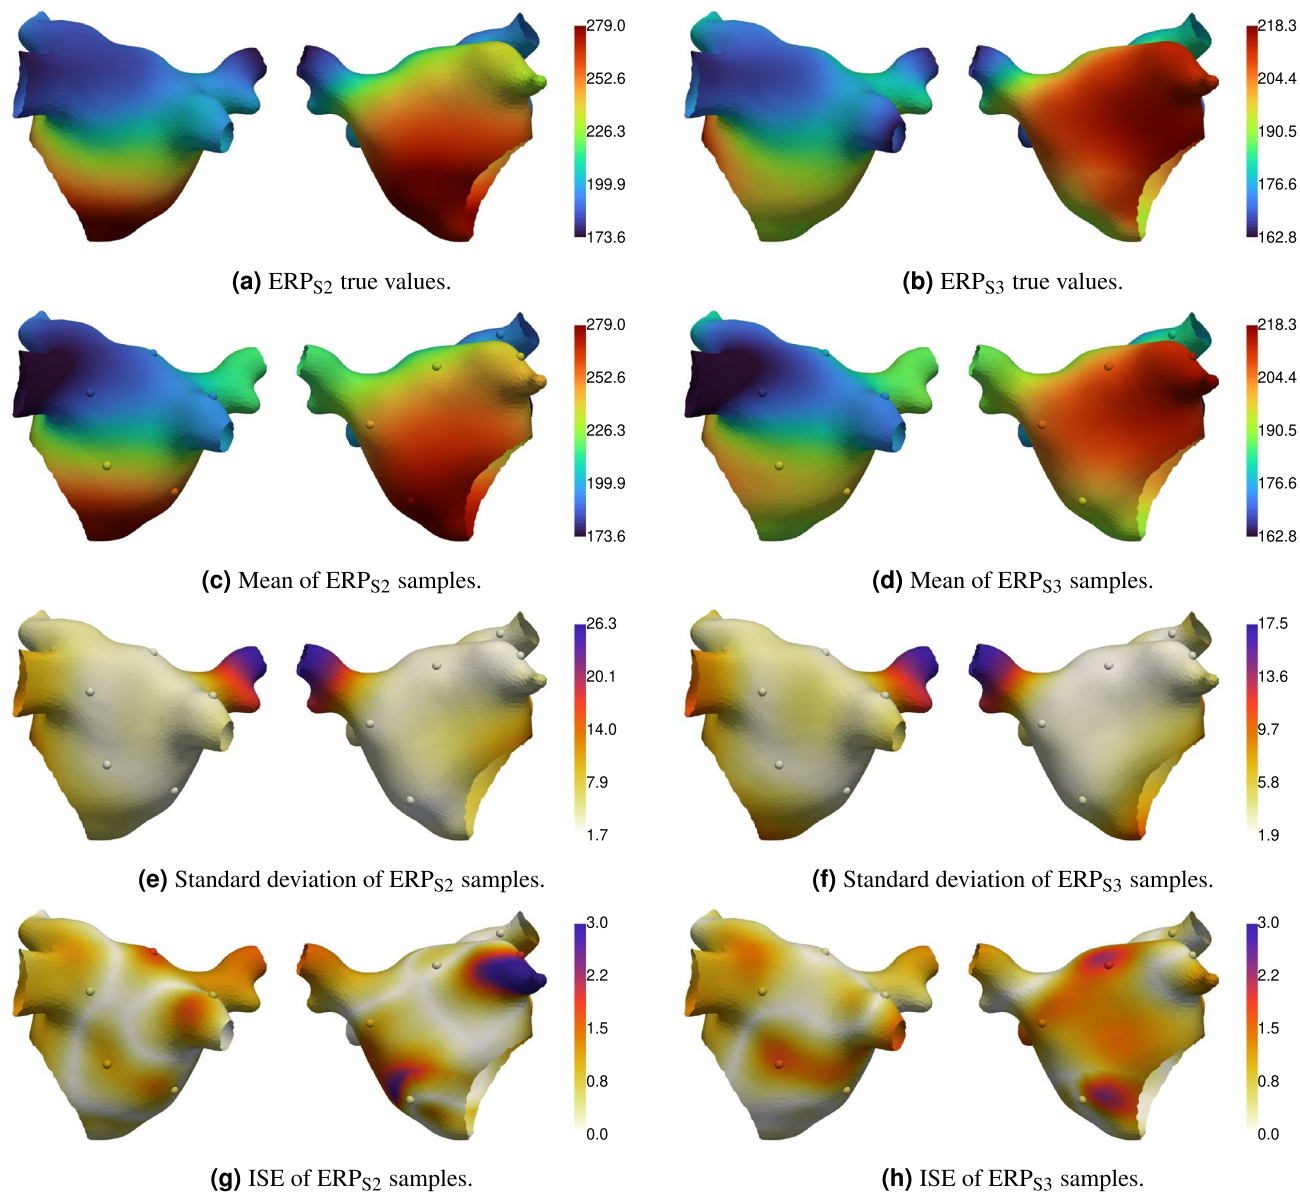


Figure 4. Ground truth and predicted (posterior mean and standard deviation) ERP fields, and independent standard error (ISE) plots comparing predicted and ground truth ERP. The spheres show the location of ERP measurements. All units are milliseconds.

generated APD_{max} and τ_{out} ground-truth fields. A unit of kernel lengthscale is approximately 3.2 mm for this mesh; see “Methods” for details. Prediction values of ERP are based on the maximum *a posteriori* estimate of the parameters, and here we use 32 eigenfunctions per EP parameter field in order to better model fields with more rapid spatial variation. Root Mean Squared Error (RMSE) is reduced with increasing lengthscale (less ERP heterogeneity), decreasing S2 and S3 resolution (more precise measurements), and increasing number of observations. We note that our likelihood function introduces a small amount of bias, discussed below, which for S2 and S3 resolution 10ms causes RMSE to increase slightly from 20 to 40 observations. Overall, the quantitative validation suggests that little is gained above 20 observation locations.

Discussion

In this paper, we have developed a workflow for calibrating an electrophysiology simulator from sparse measurements of excitability. This was done by representing the spatially varying parameter fields as Gaussian processes on a manifold, and linking these parameters to excitability observations through non-linear surrogate functions (emulators). Using a likelihood function for ERP observations, we performed probabilistic calibration to obtain the posterior distribution of the EP parameter fields. Both visual and quantitative comparison demonstrates that this workflow can successfully calibrate a simulator to ERP to a high level of accuracy.

The nature of ERP observations, in which only the interval containing ERP is observed (and the possible brackets around this interval are fixed by the S1S2(S3) protocol), is that the ability to learn more by adding

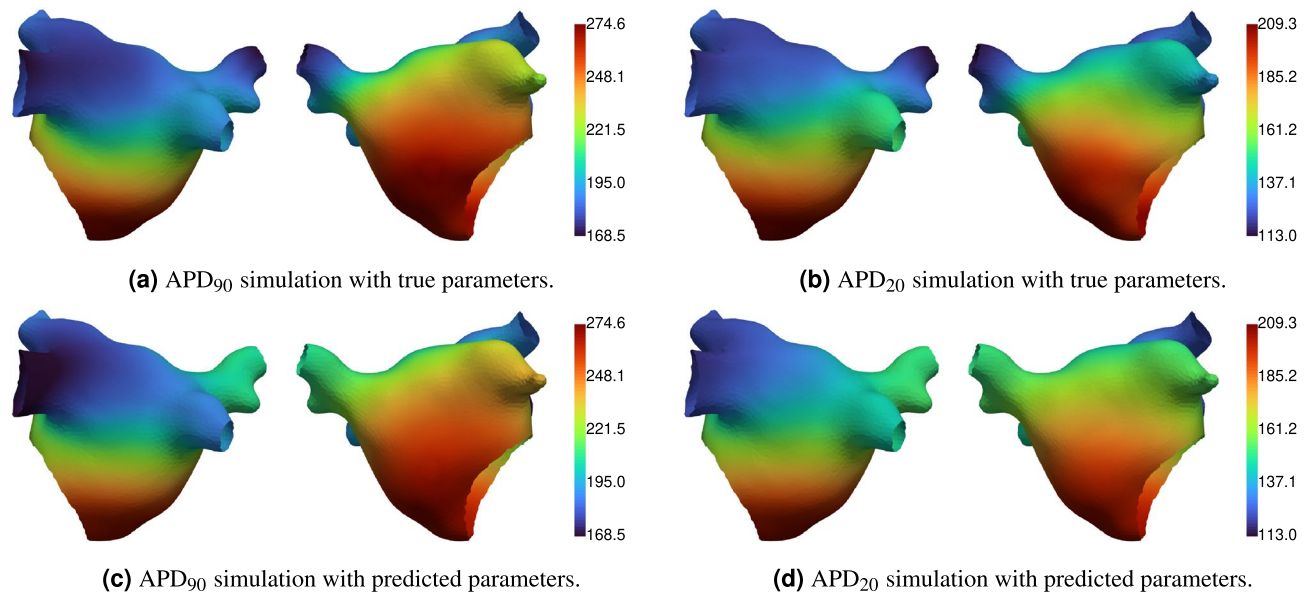


Figure 5. Ground truth and predicted APD values from simulation. All units are milliseconds.

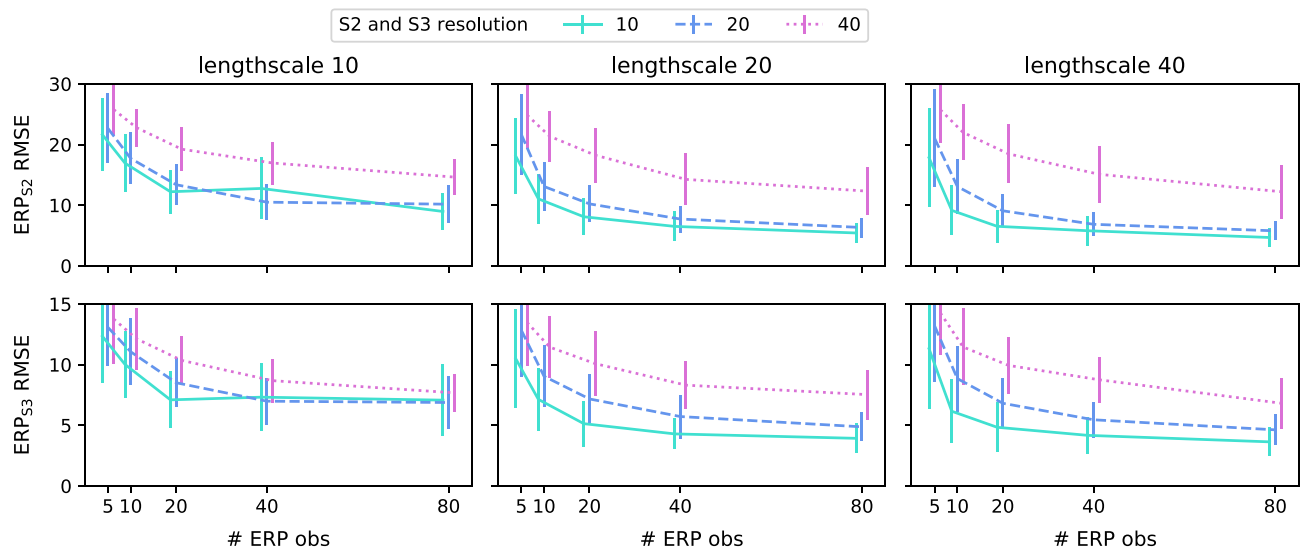


Figure 6. Quantitative validation for ERP prediction, for different lengthscales (1 unit kernel lengthscale ≈ 3.2 mm), number of observations, and S2 and S3 resolution. The model prediction is from the maximum *a posteriori* estimate. 32 basis functions are used for both EP parameter fields. For each combination of lengthscale, observations, and S2 and S3 resolution, 45 different samples of EP parameter fields are generated, and 5 different observation designs are used for each sample. The RMSE scores of ERP for each validation run is calculated over all mesh vertices. The error bars are 1 standard deviation.

observations is strongly limited above a certain point. Figure 6 demonstrates that this limit is reached faster for smaller S2 and S3 resolution. Our likelihood function does introduce a very small amount of bias, since the true likelihood should be constant in the pacing interval, but our approximation decreases on approaching the interval edges. A simple solution would be to pad the ERP observation brackets, which would remove the bias but reduce the precision. Without the assumption that measurements at locations give information about quantities at nearby locations, i.e. spatial correlation, inference about tissue properties beyond measurement sites would not be possible and atrial tissue would need to be sampled everywhere. Such regularization might make it difficult to capture discontinuous changes in tissue properties, although it would be difficult to measure such abrupt changes in tissue behaviour using sparse measurements. It may be possible to utilize other personal data (e.g. scans) or prior information (e.g. a database of clinical measurements) to assist with calibration.

The latent Gaussian process model serves two purposes. Firstly, a run of the electrophysiology model requires specification of parameters at all points on the mesh, and the Gaussian process enables this specification via interpolation between measurement locations. Secondly, we assume that parameter values at neighbouring locations on the mesh are likely to be similar, which means that we need to do joint inference for the parameters at

the measurement locations, rather than inferring parameters at each measurement location independently. In developing our method, we first attempted such an independent inference approach, in which parameters are calibrated at each measurement location independently and then interpolated over the manifold using GPML, but we were not able to obtain satisfactory results. Our current workflow easily allows more complex spatial modeling using multiple latent GPs per EP parameter field, each with independent covariance kernels and hyperparameters that can be freely given suitable priors. It also provides the benefit of being able to constrain the posterior distribution by directly manipulating the posterior samples based on *a priori* knowledge, such that parameter values (or the tissue properties depending on these parameters) should fall within a certain physiological range.

Our proposed workflow for calibration is suitable for other types of data. We have previously shown that Gaussian processes can be used as surrogate functions for CV, APD, and ERP restitution curves²¹. Observations from these restitution curves at different locations over the atrium could be included in calibration simply by including additional contributions to the likelihood function and using “Restitution Curve Emulators” to map from EP parameters to the corresponding restitution curves. Our approach here solves the problem of representing the EP parameter fields on a manifold so as to make probabilistic calibration to sparse measurements into a tractable problem. This allows for propagating uncertainty from measurements through to an ensemble of calibrated models.

Methods

Electrophysiology model. The modified Mitchell-Schaeffer (mMS) cell model^{23,24} for mono-domain tissue simulations with isotropic diffusion is expressed in the following equations:

$$\frac{\partial V_m}{\partial t} = D\nabla^2 V_m + h \frac{V_m(V_m - V_{gate})(1 - V_m)}{\tau_{in}} - (1 - h) \frac{V_m}{\tau_{out}} + J_{stim} \quad (1)$$

$$\frac{\partial h}{\partial t} = \begin{cases} (1 - h)/\tau_{open} & \text{if } V_m \leq V_{gate} \\ -h/\tau_{close} & \text{otherwise} \end{cases} \quad (2)$$

where V_m is a normalised membrane voltage, h is a gating parameter that controls recovery, and J_{stim} is an externally applied stimulus. The 4 cell model parameters $\tau = (\tau_{in}, \tau_{close}, \tau_{out}, \tau_{open})$ are time-constants that approximately characterize stages of the action potential sequence, and D is conductivity. We fixed the excitation threshold V_{gate} to 0.1. As in²¹, we reparameterized the model as follows:

$$CV_{max} = 0.5(1 - 2V_{gate})\sqrt{2D/\tau_{in}} \quad (3)$$

$$APD_{max} = \tau_{close} \log(1 + \tau_{out}(1 - V_{gate})^2/(4\tau_{in})) \quad (4)$$

In this new parameter space, weighted combinations of valid parameters are also valid parameters, which means that spatial interpolation of valid parameters will produce valid parameters. We refer to these *transformed* parameters simply as ‘parameters’. The valid ranges of these parameters are set as CV_{max} 0.1–1.5 m/s, τ_{in} 0.01–0.30 ms, τ_{out} 1–30 ms, τ_{open} 65–215 ms, APD_{max} 120–270 ms.

Atrial mesh. To generate the mesh for the simulator, the left atrial blood pool was segmented from a contrast enhanced magnetic resonance angiogram scan performed at St Thomas’ Hospital³⁰. This segmentation was meshed using a marching cubes algorithm in CEMRGAApp³¹, and the resulting surface was remeshed to a regular edge length of 0.3mm using mmgtools software³², corresponding to around 110,000 vertices, which is sufficient for simulation with the MMS model. This mesh can be found here³³, and is also included with our code²⁶.

Sensitivity analysis. To determine ERP_{S2} , the ERP value under an S1S2 protocol for S1 600 ms, and ERP_{S3} , the ERP value under an S1S2S3 protocol for S1 600 ms and S2 300 ms, we utilized a surrogate simulation: a strip of tissue with homogeneous parameters, paced from one end with the corresponding protocol, with activation measured in the strip centre²¹. The strip simulation is set up to match the atrial simulation as closely as possible (space and time discretization, cell model time-step subdivision, numerical integration, etc). We obtain simulation results with an optimized Latin hyper-cube design of 500 parameter combinations in the parameter range explained above.

Variance-based sensitivity analysis was performed by fitting a General Additive Model (GAM) to model outputs, e.g. ERP_{S2} , as a function of a single model input, e.g. APD_{max} . The expectation of the GAM is then a line through a point-cloud of input-output pairs. The variance of this line (evaluated at the inputs) divided by the variance of the point-cloud gives an approximate sensitivity index of the input on that output^{34,35}. This method can be repeated for all inputs and all outputs. We implement GAMs using the LinearGAM function with 10 splines from the Python module PYGAM³⁶. The sensitivity index of output y for input x can then be calculated as

$$\begin{aligned} \text{gam} &= \text{LinearGAM}(n_splines=10). \text{gridsearch}(x, y) \\ \text{sen} &= \text{gam}. \text{predict}(x). \text{var}() / y. \text{var}() \end{aligned}$$

Surrogate functions. The map from EP parameters (inputs) to high dimensional tissue responses (outputs), such as restitution curves, has been modelled previously using Gaussian processes²¹. Here, cubic polynomials in both $\theta_1 = \tau_{out}$ and $\theta_2 = APD_{max}$ were fit to corresponding values of ERP_{S2} and ERP_{S3}, generated from an optimized Latin hyper-cube design of 100 values of CV_{max} , τ_{out} and APD_{max} , keeping $\tau_m = 0.05$ ms and $\tau_{open} = 120$ ms in order to produce ERP and APD values in a range observed in human atrial tissue²⁹. CV_{max} was varied for robustness, but has negligible effects on ERP, as confirmed by negligible fitting residuals. There is a discontinuity in ERP_{S3} for parameter values producing ERP_{S2} \approx 285 ms, so data for ERP_{S2} > 280 ms were discarded before fitting these functions. We refer to these polynomial fits as ‘surrogate functions’ for ERP_{S2} and ERP_{S3}, denoted as $f_1(\theta_1, \theta_2)$ and $f_2(\theta_1, \theta_2)$ respectively, as they allow for determining ERP without running simulations.

Gaussian process priors. We model the EP parameters fields, $\theta_l(\mathbf{x})$, as spatially correlated random fields defined on the atrial manifold, i.e. $\mathbf{x} \in \Omega$. We use Gaussian Process Manifold Interpolation (GPMI), a method we proposed for defining Gaussian process distributions on manifolds²⁷. The approach uses solutions $\{\lambda_k, \phi_k(\mathbf{x})\}$ of the Laplacian (Laplace-Beltrami) eigenproblem on the mesh²⁸. Using GPMI allows us to represent fields on the atrium using a coordinate system that uses these eigenfunctions as a basis, enabling us to calibrate parameter fields on any given atrial manifold. The prior for each parameter field $\theta_l(\mathbf{x})$ can then be represented using the following probabilistic model, which uses the K smallest eigenvalue solutions to the Laplacian eigenproblem:

$$\theta_l(\mathbf{x}) = m_l + \alpha_l \sum_{k=1}^K (\eta_l)_k \sqrt{S(\sqrt{\lambda_k}, \rho_l)} \phi_k(\mathbf{x}) \tag{5}$$

$$(\eta_l)_k \sim \mathcal{N}(0, 1) \tag{6}$$

with hyperparameters mean m_l , amplitude α_l , and lengthscale ρ_l , for $l = 1, 2$ and $k = 1, \dots, K$. The lengthscales determine the distance over which values are correlated, with larger lengthscale corresponding to smoother parameter fields. The units of the lengthscale hyperparameters are determined by the spatial units of the mesh on which the eigenproblem is solved. See below for details. The hyperparameter vector $\eta_l \in \mathbb{R}^K$ must be given a Gaussian prior in order for this model to approximate a Gaussian process. The function $S(\sqrt{\lambda_k}, \rho_l)$ is the spectral density corresponding to the choice of covariance kernel, with the square root of the eigenvalue $\sqrt{\lambda_k}$ being the ‘frequency’ argument to this function. In this work, we use the spectral density for the radial basis function (exponentiated quadratic) kernel, but other stationary kernels could be used.

It is possible and tractable to perform inference directly on the simulation mesh by solving the Laplacian eigenproblem on this mesh. But for convenience, a lower resolution mesh of 5000 vertices was used, with vertices that are a subset of the higher resolution mesh vertices. The lower resolution mesh is produced using a simulated annealing algorithm to optimally choose a subset of 5000 nodes before meshing these new nodes to form a new surface. The routines for this are found in the quLATi package³⁷, and method details are given in our previous work²⁷. This lower resolution mesh allows for calculation of eigenfunctions with fewer computational resources, less data storage for eigenfunctions, and is convenient for plotting. Values can be easily transferred to the simulation mesh via interpolation (we use the ‘interpolate’ function in the Python software Pyvista³⁸). However, this ‘two-mesh’ approach is entirely optional. The units of lengthscale parameter ρ in Eq. (5) can be empirically related to geodesic distance by drawing many GP samples for a given kernel function on the mesh (using a lengthscale that allows for correlations to approach zero for some pairs of vertices since the mesh is finite), calculating the correlation between these samples at many pairs of vertices, and fitting (via least squares) the kernel function to correlation as a function of geodesic distance between the pairs of vertices. For the mesh in this work, 1 unit of kernel lengthscale corresponds to approximately 3.2 mm.

Bayesian calibration. Given ERP measurements at different locations \mathbf{x}_i over the atrium, it is possible to calibrate the parameter fields $\theta_1(\mathbf{x}) \equiv \tau_{out}(\mathbf{x})$ and $\theta_2(\mathbf{x}) \equiv APD_{max}(\mathbf{x})$ by obtaining the posterior distribution of the hyperparameters in Eq. (5). For convenience, we collect the hyperparameters into the vector $\psi := (m_1, m_2, \alpha_1, \alpha_2, \rho_1, \rho_2, \eta_1, \eta_2)$. Defining the ERP measurements as \mathbf{y} , then we can write the Bayesian inference problem as:

$$p(\psi | \mathbf{y}) \propto p(\mathbf{y} | \psi) p(\psi) \tag{7}$$

$$p(\mathbf{y} | \psi) := \prod_i p(y_1(\mathbf{x}_i) | \psi) \prod_i p(y_2(\mathbf{x}_i) | \psi) \tag{8}$$

where $y_1(\mathbf{x}_i)$ and $y_2(\mathbf{x}_i)$ represent observations of ERP_{S2} and ERP_{S3} respectively. We assume that both types of ERP are measured at each location, but this is not a requirement as terms can just be replaced with 1 if the corresponding measurement is not performed. Clinically, S1S2 protocols are performed by decreasing S2 by $\Delta S2$ until successful activation does not occur on the S2 beat. Therefore, observations of ERP are only observations of an interval in which ERP lies. The observation that each ERP value at a measurement location \mathbf{x}_i lies between two S2 values t_s and t_{s+1} can be expressed in the following way (see Fig. 1 for a graphical representation):

$$y_1(\mathbf{x}_i) := \text{ERP}_{S2} \in [t_s, t_{s+1}] \text{ at } \mathbf{x}_i \tag{9}$$

$$y_2(\mathbf{x}_i) := \text{ERP}_{S3} \in [t_1, t_2] \text{ at } \mathbf{x}_i \quad (10)$$

Observations can be linked to the hyperparameters $\boldsymbol{\psi}$ via the GP fields defined by Eq. (5), which determine the EP parameters at positions on the atrial mesh, and by the surrogate functions, which map these EP parameter values to ERP values:

$$p(y_m(\mathbf{x}_i) | \boldsymbol{\psi}) = p(y_m(\mathbf{x}_i) | f_m(\mathbf{x}_i)) \quad (11)$$

$$f_m(\mathbf{x}_i) := f_m(\theta_1(\mathbf{x}_i, \boldsymbol{\psi}_1), \theta_2(\mathbf{x}_i, \boldsymbol{\psi}_2)) \quad (12)$$

where $\boldsymbol{\psi}_1$ and $\boldsymbol{\psi}_2$ represent partitions of $\boldsymbol{\psi}$ for each EP parameter field.

An SIS2 pacing protocol to determine ERP effectively measures the S2 interval in which ERP lies. Defining the lower bound of this interval by I and the interval width by $\Delta S2$, the true likelihood is given by a truncated uniform, or ‘top-hat’, distribution. In other words, $p(\text{ERP} \in [I, I + \Delta S2] | \boldsymbol{\psi})$ is equal 1 if $\boldsymbol{\psi}$ produces ERP in the specified interval, and 0 otherwise. The surrogate functions $f_m(\theta_1, \theta_2)$ can be used to predict ERP from the EP parameters, which are determined by the GP fields $\theta_l(\mathbf{x})$ depending on the hyperparameters $\boldsymbol{\psi}$.

However, it is more convenient to work with an approximation to this top-hat distribution, which we previously derived for use with ERP measurements²¹. This top-hat likelihood can be approximated by dividing the interval into N sub-intervals, with a normal distribution $\mathcal{N}(c_i, s)$ centered on each sub-interval $c_i = I + (i - 1/2)\Delta S2/N$ with standard deviation equal to the sub-interval width $s = \Delta S2/N$. We choose $N = \Delta S2$, such that $s = 1$. For an observation $y_1(\mathbf{x}^*) := \text{ERP}_{S2} \in [I, I + \Delta S2]$ (and similarly for ERP_{S3}) the likelihood can be approximated as:

$$p(y_1 \in [I, I + \Delta S2] | \boldsymbol{\psi}) = \frac{1}{N} \sum_{i=1}^N \frac{1}{\sqrt{2\pi s^2}} \exp\left(-\frac{(f_1(\boldsymbol{\theta}(\mathbf{x}^*)) - c_i)^2}{2s^2}\right) \quad (13)$$

The shape of this likelihood function is a top-hat with smoothed sides and no discontinuities, such that the likelihood is approximately constant in the interval but rapidly falls to zero near the interval edges. This approximate top-hat distribution has infinite support, allowing gradient-based MCMC to be performed. For the log-likelihood, the readily available `logsumexp` function is used to prevent numerical underflow (this function is available in STAN^{39,40}, which is used for MCMC for this work). Note that this approximate top-hat distribution integrates to 1, rather than having approximately constant value 1 (in the interval), but constant factors do not matter for MCMC so we retain this form for simplicity. Also note that with a small adjustment this likelihood can be used with a Gaussian process surrogate function that predicts mean and variance²¹.

We use STAN (via PyStan)^{39,40} to perform Hamiltonian MCMC, which yields samples from the posterior distribution of the parameters $\boldsymbol{\psi}$. See Results for details. We can use these samples to calculate samples of the EP fields over the entire atrium using Eq. (5), from which ERP samples can be calculated using the surrogate functions. We modify Eq. (5) by replacing α_l with $\alpha_l/|\Phi_1|$ (where $\Phi_1(\mathbf{x})$ is constant), which assists with defining priors for α_l . We used the following priors on the hyperparameters: $\rho_l \sim \text{InvGamma}(1.01, 20)$, $m_l \sim \text{Uniform}(-\infty, +\infty)$, and $\alpha_l \sim \text{InvGamma}(1, 5)$. We found that these priors consistently allow for recoverability of both EP parameter fields. Eq. (6) gives the prior for the remaining hyperparameters.

Parameter samples. To generate ‘ground truth’ parameter fields, we draw samples from a Gaussian process defined by Eq. (5) with Matern 5/2 spectral density function using 256 eigenfunctions. We set parameters $m = 0$ and $\alpha = 1$, and ρ is set to values explained in Results and below. The generated samples for τ_{out} and APD_{max} are then scaled and offset into the full allowable parameters ranges. The same operation is performed for CV_{max} , which is needed for atrial simulations.

Certain combinations of τ_{out} and APD_{max} correspond to regions of parameter space that produce unrealistic ERP. This is handled for both ‘ground truth’ samples and posterior samples by identifying mesh nodes where the parameters produce values $\text{ERP}_{S2} > 280$ ms. The parameter values at these nodes are then replaced by a weighted average of parameter values at other nodes with acceptable ERP values, weighted by $1/d_{BH}^4$ where d_{BH} is biharmonic distance. Biharmonic distance, calculated from the Laplacian eigenvalues and eigenfunctions, is significantly cheaper to calculate than geodesic distance, and avoids topological issues from using Euclidean distance⁴¹ to interpolate values on a manifold. This procedure allows to constrain parameter samples efficiently and effectively, and is far simpler than attempting to encode such constraints into MCMC.

Synthetic experiment. We created ground truth parameter fields for τ_{out} and APD_{max} in order to verify our calibration approach (see above for details). ERP values were calculated using the surrogate functions. For ERP measurements, we generate a design of measurement locations using an optimized ‘maximin’ hypercube design, excluding mesh sites within 0.6 cm of mesh boundaries as potential sites for these measurements since clinical measurements are unlikely to sample these regions. The resolution of the SIS2 and SIS2S3 protocols are set to values specified in Results. A lengthscale of 20 was used for the example shown in Figs. 3, 4 and 5.

APD values were obtained using the atrial simulator (simulations of the mono-domain equation with the mMS model using the software *openCARP*). Tissue was paced for 8 beats from near the coronary sinus, and depolarization and repolarization were measured on the final beat. A spatially varying CV_{max} field was generated for use in this simulation, and τ_{in} and τ_{open} were fixed as described above. Simulations were run either for ground truth parameter fields or for predicted parameter fields resulting from calibration. Note that the parameters described in this manuscript were transformed back into the original parameters for the mMS model for running

simulations in *openCARP*. The diffusion time-step was 0.1 ms and the ionic current time-step was 0.02 ms. The mMS action potential is normalized to have minimum 0 and maximum 1. Activation (depolarization) was measured when V_m reached 0.7 on upstroke, and recovery (repolarization) was measured when V_m fell to 0.8 (APD₂₀), 0.7 (APD₃₀), 0.5 (APD₅₀), and 0.1 (APD₉₀). APD values are the time between activation and recovery. Simulations results for APD₂₀ and APD₉₀ are given in Results.

Data availability

Our code and example mesh are available in a Zenodo repository²⁶.

Received: 6 June 2022; Accepted: 19 September 2022

Published online: 04 October 2022

References

1. Staerk, L., Sherer, J. A., Ko, D., Benjamin, E. J. & Helm, R. H. Atrial fibrillation: Epidemiology, pathophysiology, and clinical outcomes. *Circ. Res.* **120**, 1501–1517. <https://doi.org/10.1161/CIRCRESAHA.117.309732> (2017).
2. Gray, R. A. *et al.* Mechanisms of cardiac fibrillation. *Science* **270**, 1224–1225 (1995).
3. Parameswaran, R., Al-Kaisey, A. M. & Kalman, J. M. Catheter ablation for atrial fibrillation: Current indications and evolving technologies. *Nat. Rev. Cardiol.* **18**, 210–225. <https://doi.org/10.1038/s41569-020-00451-x> (2021).
4. Qu, Z., Weiss, J. N. & Garfinkel, A. Cardiac electrical restitution properties and stability of reentrant spiral waves: a simulation study. *Am. J. Physiol. Heart Circ. Physiol.* **276**, H269–H283. <https://doi.org/10.1152/ajpheart.1999.276.1.H269> (1999).
5. Garfinkel, A. *et al.* Preventing ventricular fibrillation by flattening cardiac restitution. *Proc. Natl. Acad. Sci. U.S.A.* **97**, 6061–6066 (2000).
6. Loewe, A. *et al.* Patient-specific identification of atrial flutter vulnerability: A computational approach to reveal latent reentry pathways. *Front. Physiol.* **9**, 1910. <https://doi.org/10.3389/fphys.2018.01910> (2019).
7. Narayan, S. M., Franz, M. R., Clopton, P., Pruvot, E. J. & Krummen, D. E. Repolarization alternans reveals vulnerability to human atrial fibrillation. *Circulation* **123**, 2922–2930. <https://doi.org/10.1161/CIRCULATIONAHA.110.977827> (2011).
8. Krummen, D. E. *et al.* Mechanisms for human atrial fibrillation initiation: Clinical and computational studies of repolarization restitution and activation latency. *Circ. Arrhythmia Electrophysiol.* **5**. <https://doi.org/10.1161/CIRCEP.111.969022> (2012).
9. Cochet, H. *et al.* Age, atrial fibrillation, and structural heart disease are the main determinants of left atrial fibrosis detected by delayed-enhanced magnetic resonance imaging in a general cardiology population. *J. Cardiovasc. Electrophysiol.* **26**, 484–492. <https://doi.org/10.1111/jce.12651> (2015).
10. Bishop, M. J., Connolly, A. & Plank, G. Structural heterogeneity modulates effective refractory period: A mechanism of focal arrhythmia initiation. *PLoS One* **9**, e109754. <https://doi.org/10.1371/journal.pone.0109754> (2014).
11. Fareh, S., Villemare, C. & Nattel, S. Importance of refractoriness heterogeneity in the enhanced vulnerability to atrial fibrillation induction caused by tachycardia-induced atrial electrical remodeling. *Circulation* **98**, 2202–2209. <https://doi.org/10.1161/01.cir.98.20.2202> (1998).
12. Clayton, R. H. *et al.* Models of cardiac tissue electrophysiology: Progress, challenges and open questions. *Prog. Biophys. Mol. Biol.* **104**, 22–48. <https://doi.org/10.1016/j.pbiomolbio.2010.05.008> (2011).
13. Niederer, S. A., Lumens, J. & Trayanova, N. A. Computational models in cardiology. *Nat. Rev. Cardiol.* **16**, 100–111. <https://doi.org/10.1038/s41569-018-0104-y> (2019).
14. Prakosa, A. *et al.* Personalized virtual-heart technology for guiding the ablation of infarct-related ventricular tachycardia. *Nat. Biomed. Eng.* **2**, 732–740. <https://doi.org/10.1038/s41551-018-0282-2> (2018).
15. Boyle, P. M. *et al.* Computationally guided personalized targeted ablation of persistent atrial fibrillation. *Nat. Biomed. Eng.* **3**, 870–879. <https://doi.org/10.1038/s41551-019-0437-9> (2019).
16. Corrado, C. *et al.* Quantifying atrial anatomy uncertainty from clinical data and its impact on electro-physiology simulation predictions. *Med. Image Anal.* **61**, 101626. <https://doi.org/10.1016/j.media.2019.101626> (2020).
17. Roney, C. H. *et al.* A technique for measuring anisotropy in atrial conduction to estimate conduction velocity and atrial fibre direction. *Comput. Biol. Med.* **1**–13. <https://doi.org/10.1016/j.combiomed.2018.10.019> (2018).
18. Grandits, T. *et al.* PIEMAP: Personalized inverse eikonal model from cardiac electro-anatomical maps. In Puyol Anton, E. *et al.* (eds.) *Statistical Atlases and Computational Models of the Heart. M & Ms and EMIDEC Challenges*, vol. 12592, 76–86. https://doi.org/10.1007/978-3-030-68107-4_8 (Springer International Publishing, Cham, 2021). Series Title: Lecture Notes in Computer Science.
19. Corrado, C. *et al.* Personalized models of human atrial electrophysiology derived from endocardial electrograms. *IEEE Trans. Biomed. Eng.* **64**, 735–742. <https://doi.org/10.1109/TBME.2016.2574619> (2017).
20. Corrado, C. *et al.* A work flow to build and validate patient specific left atrium electrophysiology models from catheter measurements. *Med. Image Anal.* **47**, 153–163. <https://doi.org/10.1016/j.media.2018.04.005> (2018).
21. Coveney, S. *et al.* Bayesian calibration of electrophysiology models using restitution curve emulators. *Front. Physiol.* **12**, 693015. <https://doi.org/10.3389/fphys.2021.693015> (2021).
22. Whittaker, D. G., Clerx, M., Lei, C. L., Christini, D. J. & Mirams, G. R. Calibration of ionic and cellular cardiac electrophysiology models. *WIREs Syst. Biol. Med.* **12**. <https://doi.org/10.1002/wsbm.1482> (2020).
23. Mitchell, C. & Schaeffer, D. A two-current model for the dynamics of cardiac membrane. *Bull. Math. Biol.* **65**, 767–793. [https://doi.org/10.1016/S0092-8240\(03\)00041-7](https://doi.org/10.1016/S0092-8240(03)00041-7) (2003).
24. Corrado, C. & Niederer, S. A. A two-variable model robust to pacemaker behaviour for the dynamics of the cardiac action potential. *Math. Biosci.* **281**, 46–54. <https://doi.org/10.1016/j.mbs.2016.08.010> (2016).
25. Plank, G. *et al.* The openCARP simulation environment for cardiac electrophysiology. *Comput. Methods Programs Biomed.* **208**, 106223. <https://doi.org/10.1016/j.cmpb.2021.106223> (2021).
26. Coveney, S. *samcoveney/mmserp*: v1.0.0. <https://doi.org/10.5281/zenodo.7081857> (2022).
27. Coveney, S. *et al.* Gaussian process manifold interpolation for probabilistic atrial activation maps and uncertain conduction velocity. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **378**, 20190345. <https://doi.org/10.1098/rsta.2019.0345> (2020).
28. Solin, A. & Särkkä, S. Hilbert space methods for reduced-rank Gaussian process regression. *Stat. Comput.* **30**, 419–446. <https://doi.org/10.1007/s11222-019-09886-w> (2020).
29. Bode, F., Kilborn, M., Karasik, P. & Franz, M. R. The repolarization-excitability relationship in the human right atrium is unaffected by cycle length, recording site and prior arrhythmias. *J. Am. Coll. Cardiol.* **37**, 920–925. [https://doi.org/10.1016/S0735-1097\(00\)01189-X](https://doi.org/10.1016/S0735-1097(00)01189-X) (2001).
30. Sim, I. *et al.* Reproducibility of Atrial Fibrosis Assessment Using CMR Imaging and an Open Source Platform. *JACC Cardiovasc. Imaging* **12**, 2076–2077. <https://doi.org/10.1016/j.jcmg.2019.03.027> (2019).
31. Razeghi, O. *et al.* CemrgApp: An interactive medical imaging application with image processing, computer vision, and machine learning toolkits for cardiovascular research. *SoftwareX* **12**, 100570. <https://doi.org/10.1016/j.softx.2020.100570> (2020).

32. Dapogny, C., Dobrzynski, C. & Frey, P. Three-dimensional adaptive domain remeshing, implicit domain meshing, and applications to free and moving boundary problems. *J. Comput. Phys.* **262**, 358–378 (2014).
33. Roney, C. *et al.* Predicting atrial fibrillation recurrence by combining population data and virtual cohorts of patient-specific left atrial models, <https://doi.org/10.5281/zenodo.5801337> (2022).
34. Strong, M., Oakley, J. E. & Brennan, A. Estimating multiparameter partial expected value of perfect information from a probabilistic sensitivity analysis sample. *Med. Decis. Making* **34**, 311–326. <https://doi.org/10.1177/0272989X13505910> (2016).
35. Wood, S. N. *Generalized additive models: An introduction with R* 2nd edn. (Chapman and Hall/CRC, New York, 2017).
36. Servén, D., Brummitt, C., Abedi, H. & hlink. dswah/pyGAM: v0.8.0. <https://doi.org/10.5281/zenodo.1476122> (2018).
37. Coveney, S. samcoveney/quLATI: v2.0, <https://doi.org/10.5281/zenodo.3758042> (2021).
38. Sullivan, C. B. & Kaszynski, A. PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK). *J. Open Sour. Softw.* **4**, 1450. <https://doi.org/10.21105/joss.01450> (2019).
39. Stan Development Team. Stan Modeling Language Users Guide and Reference Manual. (2021).
40. Riddell, A., Hartikainen, A. & Carter, M. PyStan (2.19.1.1). (2021).
41. Lipman, Y., Rustamov, R. M. & Funkhouser, T. A. Biharmonic distance. *ACM Trans. Graph.* **29**, 1–11. <https://doi.org/10.1145/1805964.1805971> (2010).

Acknowledgements

This work was funded by grants from the UK Engineering and Physical Sciences Research Council (EP/P010741/1, EP/P01268X/1). CHR acknowledges a Medical Research Council Skills Development Fellowship (MR/S015086/1).

Author contributions

S.C. conceived and implemented the work. C.H.R. and C.C. contributed to atrial simulation scripts. R.D.W. and J.E.O. advised on sensitivity analysis and calibration. S.D.N. and R.H.C. provided guidance on cardiac modelling and the clinical setting. All authors contributed to the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to S.C.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022