Uncertainty Quantification in a Cardiac Arrhythmia Model: Application to Intra-Atrial Reentrant Tachycardia

Maarten Volkaerts [1], Marie Cloet[2], Hans Dierckx[2], Piet Claus[3], Giovanni Samaey[1]

[1] NUMA, Department of Computer Science, KU Leuven, Leuven, Belgium
[2] Department of Mathematics, KU Leuven Kulak Kortrijk Campus, Kortrijk, Belgium
[3] Cardiovascular Imaging and Dynamics, Department of Cardiovascular Sciences, KU Leuven, Leuven, Belgium

Abstract

Patient-specific computational models of cardiac electrophysiology have been proposed as tools to improve the treatment of cardiac arrhythmia in the clinic. To build accurate, trustworthy models for clinical use, quantifying the uncertainties, introduced by fitting the model to noisy and sparse data, is vital. We present a methodology for Bayesian inference of model parameters at the tissue level. We apply this to a toy problem, motivated by Intra-Atrial Reentrant Tachycardia (IART): we infer the size and orientation of a scar in a two-dimensional tissue slab based on synthetic sparse catheter data. We define a likelihood function based on a low-dimensional representation of the data in the form of summary statistics. For a good exploration of the parameter space, we use an adaptive Markov chain Monte Carlo algorithm. As a result, our methodology notably decreases the required number of forward simulations, which greatly benefits the time feasibility of Bayesian inference. The resulting posterior both provides an estimate of the scar parameters and quantifies the uncertainty due to the sparse and noisy measurements. Finally, we highlight the need to adapt forward solvers to avoid artificial discontinuities in the likelihood.

1. Introduction

Personalized heart models or cardiac digital twins integrate biophysics with patient data to offer a virtual counterpart of a patient's heart. These models are receiving growing interest to aid clinical decision-making thanks to their potential to improve the treatment of cardiac arrhythmia through an increased personalization of care [1].

However, the absence of uncertainty quantification in state-of-the-art cardiac digital twins is identified as an important barrier to their clinical introduction [2]. For accurate patient-specific descriptions, we need to quantify the uncertainties introduced by fitting a model to sparse and noisy data. Moreover, including uncertainties also allows assessment of the credibility of predictive model outputs.

Recent work started to include uncertainties in cardiac electrophysiology (EP) models at different scales by using Bayesian inference [3–5]. A posterior probability distribution on the parameters complements an estimate with an uncertainty and is sampled with Markov chain Monte Carlo (MCMC) methods. However, we need many samples to estimate the posterior well and each sample requires a simulation of the EP model. This quickly becomes costly, especially at the tissue scale. As a result, the posterior is often approximated by replacing the EP model with a surrogate during inference [4, 5].

In this work, we present a methodology that enables Bayesian inference at the tissue level directly using the EP model. We infer an elliptical scar in two-dimensional tissue based on a few synthetic intracardiac electrograms (EGM), motivated by atrial arrhythmia in which an excitation wave attaches to a surgical scar. Our method strongly reduces the required number of samples of the posterior and, thus, the number of forward simulations. First, we define a likelihood based on summary statistics (the local activation times (LAT) and the period of the reentry) since low-dimensional data facilitate efficient exploration of the parameter space. We will show that reducing the information content does not hurt the identifiability of the parameters. Second, fast exploration of the parameter space is achieved by using an adaptive MCMC algorithm.

2. Methodology

2.1. Personalized EP model

We consider a monodomain EP cardiac tissue model:

$$\frac{\partial \mathbf{u}}{\partial t} = D\nabla^2 \mathbf{u} + \mathbf{f}(\mathbf{u}). \tag{1}$$

We use the smoothed Karma model [6] as reaction term $f(u)$ to integrate cell dynamics in the reaction-diffusion

equation. This model involves two state variables: the transmembrane potential and a slow recovery variable. We use the pseudo-EGM formulation [7] to compute an approximation of the electrical potential at the endocardium from the transmembrane potential.

In this work, we personalize the geometry of a region of scarred, non-conducting tissue. We consider a two-dimensional tissue slab, representing the atrium, with dimension $L_x = L_y = 108\ mm$. An elliptical scar, with its center at $\left(\frac{L_x}{2}, \frac{L_x}{2}\right)$, blocks all conduction ($D = 0$). During the inference, we estimate the scar geometry parameter $\theta = [r_x, r_y, \phi]$. We specify the scar shape with its long and short axes r_x and r_y and its orientation with the inclination angle ϕ between its long axis and the positive x-axis.

We solve the reaction-diffusion model from Equation (1) using the software package ithildin [8]. This solver uses a second-order finite difference approximation of the Laplacian and a Forward Euler time-stepper. Time and space discretization are fixed: $dt = 0.25\ ms$ and $dx = 0.3\ mm$. We perform domain splitting across 72 cores.

2.2. Data

We perform inference based on synthetic EGM data, generated as if the measurements were made by a PENTARAY catheter (Biosense Webster, Diegem, Belgium). Five flexible arms each contain four electrodes at a distance of 4 mm from each other. In this work, we consider the catheter to be at a fixed location at a distance from the scar. All electrodes touch the heart wall directly and measure the potential at the endocardium. For the exact electrode locations, we refer to the code (see also Section 3)

The tissue is observed during sustained reentry. To generate such behavior, the initial condition attaches a spiral wave to the scar. The EGM data y consist of 20 time traces with a resolution of $1\ ms$ over a time interval of length $T = 1700\ ms$, chosen such that four periods of the reentry are captured. Starting from the third reentry, steady-state behavior is assumed. We add independent noise with distribution $\mathcal{N}(0, 10^{-2})$.

2.3. Bayesian inference

We define the forward map $F(\theta)$ as the modeled endocardial potential at the electrode locations for the scar geometry θ. This map combines solving the monodomain equation for the protocol from Section 2.2 and the pseudo-EGM formalism. Even for the true scar geometry θ_{true}, the modeled potential will not match the measurements exactly due to measurement noise ε:

$$y = F(\theta_{true}) + \varepsilon. \tag{2}$$

Using Bayes' rule, we can identify a posterior probability distribution on θ that captures the resulting uncertainties on the model parameters as follows:

$$p(\theta|y) = \frac{l(y|\theta)p(\theta)}{p(y)}. \tag{3}$$

The prior distribution $p(\theta)$ encodes prior knowledge about the scar. During the experiments, we assume no prior knowledge by defining $p(\theta)$ as a wide uniform distribution. The likelihood $l(y|\theta)$ quantifies for each possible scar geometry parameter θ how likely this parameter would have resulted in the data y (see Section 2.4). The evidence $p(y)$ is an unknown normalization constant.

By using Bayesian inference, we take into account that a whole range of parameter values could have resulted in the data y. In contrast, most cardiac digital twins limit themselves to finding θ^*, that optimally matches the data with the modeled potential, as an approximation for θ_{true}.

Moreover, Bayesian inference offers a rigorous mathematical framework that combines the forward model and our knowledge about the measurement noise to link the scar and the surrounding reentry in a consistent manner. This contrasts with clinical practice: a catheter roves over the endocardium to map the electrical propagation pathway based on LATs and to label low-voltage areas as scar. However, we have no guarantee that the registered scar and activation map are perfectly consistent.

2.4. Likelihood

To quantify how likely a scar geometry θ would have resulted in the data y, the likelihood is defined as $l(y|\theta) = p(y|\theta) = p(y|F(\theta))$. Hence, from Equation (2), it is clear that the likelihood captures the noise model of ε.

In this work, we reduce the EGM time traces to a low-dimensional set of summary statistics to increase the efficiency of the inference. Equation (2) can be rewritten in terms of the summary statistics as follows:

$$s(y) = s(F(\theta)) + \varepsilon'. \tag{4}$$

The likelihood is now defined as $l(y|\theta) = p(s(y)|s(F(\theta)))$ and now models the measurement noise of the statistics ε'.

The dimension reduction of the data improves the efficiency of the algorithms proposed in Section 2.5 as it facilitates faster exploration of the parameter space. We need fewer samples and, thus, fewer forward simulations to find the posterior probability distribution. In Section 3, we verify whether the loss of information, caused by reducing the data to a set of statistics, impacts the identifiability of the scar parameters.

We define the following summary statistics: the LAT at each electrode during the fourth reentry and the period of the steady-state reentry. This can be extended by including more relevant quantities later on. In this work, we assume, $\varepsilon' \sim \mathcal{N}(0, 0.1I)$, while a more thorough modeling of this distribution is postponed.

2.5. Adaptive MCMC

MCMC methods sample the posterior probability distribution from Equation (3). The right-hand side of this equation is not known in closed form but it can evaluated for θ up to an unknown constant value. During every iteration, MCMC proposes a new sample based on the previous sample and uses the posterior probability of the proposed sample to decide upon its acceptance.

For efficiency, we need a proposal that explores the parameter space well. This is made easier thanks to our low-dimensional scar parameter and our likelihood definition. Moreover, we propose an adaptive MCMC algorithm to sample the posterior distribution of strongly correlated parameters, r_x and r_y in particular, efficiently [9]. In contrast with Metropolis-Hastings [10, 11], we learn the covariance structure of the posterior distribution online to improve the exploration of the parameter space. We propose larger moves in directions with higher variance.

For the first 100 samples, we fix the proposal distribution to a multivariate Gaussian with $\mathcal{N}(0, C_0)$ with $C_0 = diag(1.69 * 10^{-4}, 1.69 * 10^{-4}, \pi^2/147456)$. After, we update the covariance of the proposal after every 10 samples such that $\mathcal{N}(0, C_n/10)$ with C_n the sample covariance of the first n samples. We construct chains of length 10000, starting from the ground truth value $\theta = [10.0, 4.0, 0.0]$. The first 1000 samples are discarded as burn-in.

3. Results & Discussion

We obtain histograms as an approximation for the posterior distribution of the parameters [1] (see Figure 1). This posterior distribution quantifies the scar parameter uncertainties due to measurement noise. For the assumed noise model, we find a narrow posterior that captures the ground truth value of the scar geometry. The marginal distributions of r_y and ϕ are centered around their ground truth value and even have a Gaussian-looking shape. Meanwhile, for r_x, the ground truth appears closer to the tail, although the bias between the maximum-a-posteriori estimate and the ground truth is small. Finally, Figure 2 shows the strong posterior correlation between r_x and r_y.

At this point, evaluating the performance of our methodology is more important than conducting a study of the posterior uncertainties. After all, the posterior shape is influenced by the noise model for the summary statistics. To analyze the nature of the parameter uncertainties caused by clinical data, follow-up work should develop and include a more accurate noise model. In the remainder of this Section, we will analyze our algorithm.

First, adaptive MCMC proves to sample the posterior distribution efficiently. Due to the strong correlation between r_x and r_y, the directions with the largest variance are distinct from the axes in the parameter space. Adaptive MCMC learns the posterior covariance structure during sampling and proposes larger moves in directions of high variance. Figure 3 illustrates how this increases the efficiency of the sampling. After a short burn-in time, adaptive MCMC mixes notably better than classic MCMC, meaning that the entire posterior distribution is sampled continuously. As a result, we need fewer samples to estimate the entire posterior distribution well and we reduce the number of forward simulations.

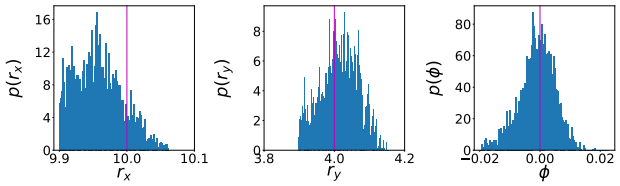

Figure 1: Histograms of the marginal posterior probability distribution are obtained with adaptive MCMC for the long principal axes r_x, the short principal axes r_y, and the inclination angle ϕ of the elliptical scar.

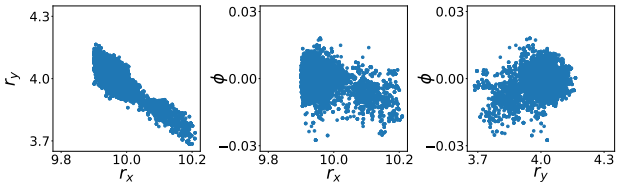

Figure 2: The samples of the posterior probability distribution are projected on the (r_x, r_y)-, the (r_x, ϕ)- and the (r_y, ϕ)-plane. A strong posterior correlation between the two principle axes r_x and r_y is observed.

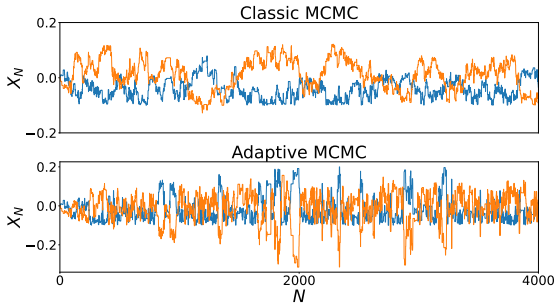

Figure 3: The Markov chains are projected on the r_x- and r_y-axis. Adaptive MCMC improves mixing considerably for parameters with a strongly correlated posterior.

Second, we manage to fully characterize the scar from 21 statistics. While reducing the dimension of the data might induce some additional uncertainties, the statistics

[1] All code to run the experiments is publicly available at https://gitlab.kuleuven.be/numa/public/cinc24-uq-for-iart.git

are informative enough to not affect the identifiability.

Finally, we also notice an important limitation. A numerical artifact in the forward solver introduces an artificial discontinuity in the posterior, very clearly visible in the marginal distribution for r_x (Figure 1). Each mesh point of an equidistant finite difference mesh is labeled as either healthy or scarred (see Figure 4), and the simulation discards unexcitable, scarred tissue. For a small perturbation to a geometry parameter, no mesh points change label. As a result, the solution of the forward problem and the likelihood evolve discontinuously in terms of the scar parameters. In turn, this imposes an artificial discontinuity in the posterior. To improve the accuracy of the posterior, we must adapt the forward solver such that it no longer creates artificial discontinuities in the likelihood.

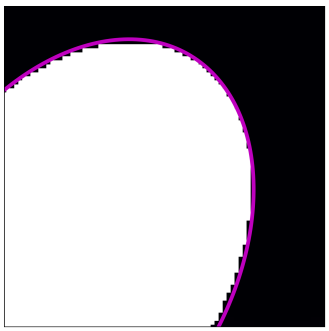

Figure 4: Mesh points are labeled as inside (white) or outside (black) the elliptical scar bordered by the purple curve. The discretized scar only changes for large enough perturbations that change at least the label of one mesh point.

4. Conclusion

Bayesian inference facilitates uncertainty quantification when inferring model parameters from data. We propose a methodology, that uses adaptive MCMC, to estimate the posterior distribution with notably fewer forward simulations to enable inference at the tissue level. While we illustrate this for the identification of three scar parameters, the use of this method can be extended to the inference of more complex scars or other functional parameters. The likelihood based on summary statistics increases the efficiency of the algorithm without affecting the identifiability of the scar parameters. Follow-up work should include a more accurate noise model and investigate how including extra statistics influences the posterior shape. Moreover, we focus on removing numerical artifacts in the likelihood.

Acknowledgments

This project was supported by the KU Leuven internally funded project 'Bayesian Inversion of Cardiac Electrograms for Patientspecific analysis (BICEPS)" IDN/22/006.

References

[1] Niederer SA, Lumens J, Trayanova NA. Computational Models in Cardiology. Nature Reviews Cardiology 2019; 16(2):100–111. ISSN 1759-5002, 1759-5010.

[2] Mirams GR, Pathmanathan P, Gray RA, Challenor P, Clayton RH. Uncertainty and Variability in Computational and Mathematical Models of Cardiac Physiology: Uncertainty and Variability in Cardiac Models. The Journal of Physiology 2016;594(23):6833–6847. ISSN 00223751.

[3] Johnstone RH, Chang ET, Bardenet R, De Boer TP, Gavaghan DJ, Pathmanathan P, Clayton RH, Mirams GR. Uncertainty and Variability in Models of the Cardiac Action Potential: Can we Build Trustworthy Models? Journal of Molecular and Cellular Cardiology 2016;96:49–62. ISSN 00222828.

[4] Konukoglu E, Relan J, Cilingir U, Menze BH, Chinchapatnam P, Jadidi A, Cochet H, Hocini M, Delingette H, Jaïs P, Haïssaguerre M, Ayache N, Sermesant M. Efficient Probabilistic Model Personalization Integrating Uncertainty on Data and Parameters: Application to Eikonal-Diffusion Models in Cardiac Electrophysiology. Progress in Biophysics and Molecular Biology 2011;107(1):134–146. ISSN 00796107.

[5] Coveney S, Corrado C, Oakley JE, Wilkinson RD, Niederer SA, Clayton RH. Bayesian Calibration of Electrophysiology Models Using Restitution Curve Emulators. Frontiers in Physiology 2021;12. ISSN 1664-042X.

[6] Marcotte CD, Grigoriev RO. Dynamical mechanism of atrial fibrillation: A topological approach. Chaos An Interdisciplinary Journal of Nonlinear Science 2017; 27(9):093936. ISSN 1054-1500, 1089-7682.

[7] Bishop MJ, Plank G. Bidomain ECG Simulations Using an Augmented Monodomain Model for the Cardiac Source. IEEE Transactions on Biomedical Engineering 2011;58(8):2297–2307. ISSN 0018-9294, 1558-2531.

[8] Kabus D, Cloet M, Zemlin C, Bernus O, Dierckx H. The Ithildin Library for Efficient Numerical Solution of Anisotropic Reaction-Diffusion Problems in Excitable Media. PLOS ONE 2024;19(9):e0303674. ISSN 1932-6203.

[9] Haario H, Saksman E, Tamminen J. An Adaptive Metropolis Algorithm. Bernoulli 2001;7(2):223. ISSN 13507265.

[10] Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E. Equation of State Calculations by Fast Computing Machines. The Journal of Chemical Physics 1953; 21(6):1087–1092. ISSN 0021-9606, 1089-7690.

[11] Hastings WK. Monte Carlo Sampling Methods Using Markov Chains and their Applications. Biometrika 1970; 57(1):97–109. ISSN 1464-3510, 0006-3444.

Address for correspondence:

Maarten Volkaerts
Celestijnenlaan 200a, 3001 Leuven
maarten.volkaerts@kuleuven.be