



Bayesian inference for fitting cardiac models to experiments: estimating parameter distributions using Hamiltonian Monte Carlo and approximate Bayesian computation

Alejandro Nieto Ramos^{1,2} · Flavio H. Fenton³ · Elizabeth M. Cherry⁴

Received: 15 April 2022 / Accepted: 2 October 2022 / Published online: 2 November 2022
© International Federation for Medical and Biological Engineering 2022

Abstract

Customization of cardiac action potential models has become increasingly important with the recognition of patient-specific models and virtual patient cohorts as valuable predictive tools. Nevertheless, developing customized models by fitting parameters to data poses technical and methodological challenges: despite noise and variability associated with real-world datasets, traditional optimization methods produce a single “best-fit” set of parameter values. Bayesian estimation methods seek distributions of parameter values given the data by obtaining samples from the target distribution, but in practice widely known Bayesian algorithms like Markov chain Monte Carlo tend to be computationally inefficient and scale poorly with the dimensionality of parameter space. In this paper, we consider two computationally efficient Bayesian approaches: the Hamiltonian Monte Carlo (HMC) algorithm and the approximate Bayesian computation sequential Monte Carlo (ABC-SMC) algorithm. We find that both methods successfully identify distributions of model parameters for two cardiac action potential models using model-derived synthetic data and an experimental dataset from a zebrafish heart. Although both methods appear to converge to the same distribution family and are computationally efficient, HMC generally finds narrower marginal distributions, while ABC-SMC is less sensitive to the algorithmic settings including the prior distribution.

Keywords Cardiac action potential · Mitchell-Schaeffer model · Fenton-Karma model · Alternans · Statistical computing

1 Introduction

A broad range of models of cardiac action potentials have been developed [1] to describe the heart’s complex electrical dynamics across multiple species and regions of the heart. These models also vary in complexity and level of detail, from simple two-variable models with a few parameters up to as many as dozens of variables and hundreds of parameters. In most cases, these models are published with a single set of parameter values. Nevertheless, there is often a need or desire to customize the models by obtaining parameter values that can be used to match particular experimental recordings [2], to represent individual patients [3], or to create virtual cohorts [4].

The challenges of finding parameter values to match specific input data or properties are well known. Within cardiac electrophysiology modeling, many approaches have been used, from optimization techniques to heuristic methods. Some examples include least squares variations [5], sequential quadratic programming [6], genetic algorithms [7–9], and a hybrid method combining particle swarm optimization with a local gradient-based algorithm [10]. More recently,

✉ Elizabeth M. Cherry
elizabeth.cherry@gatech.edu

Alejandro Nieto Ramos
axn2780@rit.edu

Flavio H. Fenton
ffenton@physics.gatech.edu

¹ School of Mathematical Sciences, Rochester Institute of Technology, 1 Lomb Memorial Drive, 14623 Rochester, NY, USA

² Present Address: Epilepsy Center, Neurological Institute, Cleveland Clinic, 9500 Euclid Avenue, Cleveland, OH 44195, USA

³ School of Physics, Georgia Institute of Technology, 837 State Street NW, 30332 Atlanta, GA, USA

⁴ School of Computational Science and Engineering, Georgia Institute of Technology, 756 West Peachtree Street, 30308 Atlanta, GA, USA

parameterizations have been obtained using Bayesian approaches, such as history matching [11], Bayesian active learning [12], and a combination of Metropolis-Hastings and Gibbs sampling [13].

Nevertheless, models that include a single set of parameter values, however well fitted, lack a depth of information that can be included when a distribution of parameter values is obtained. For example, when multiple datasets are available, simplification by fitting to the mean may misrepresent properties of the data [14]. In addition, models with a single set of parameter values neglect that many models are not fully identifiable from the necessarily limited input data used for fitting [5, 9, 15].

In contrast, Bayesian inference allows fitting a probability model to a set of data with the result summarized by a probability distribution on the parameters of the model. The distribution not only provides a description of randomness in the observed data, it also can facilitate making predictions of unobserved quantities. In addition, when working with observed data, the distributions provide information about the ranges of values that the different model parameters can take. Traditional full Bayesian Markov chain Monte Carlo (MCMC) methods, such as Metropolis-Hastings or Gibbs sampling, have limitations, including random walk behavior and poor scalability with the dimensionality of parameter space [16], and sensitivity of the desired posterior distribution to the full Bayesian specification [17]. Several methods have been developed to overcome the computational limitations of traditional Bayesian methods. One approach is the use of approximate Bayesian methods [18], where the likelihood, which can be expensive to compute, is not used, resulting in greater efficiency. An example is approximate Bayesian computation sequential Monte Carlo (ABC-SMC) [19, 20], which has been applied to obtain parameter distributions for the Hodgkin-Huxley neural model [15] and for the O'Hara et al. model [21] of cardiac cells [22].

One option to calculate the exact, rather than approximate, distributions without sacrificing efficiency is Hamiltonian Monte Carlo (HMC) [16, 23], a full Bayesian method that uses the gradient of the target distribution to explore the parameter space in a more efficient manner than the methods mentioned above. HMC has been used successfully for model calibration in ecology [24] and pharmacometry [25], but its utility for cardiac action potential models [26] compared to approximate methods has not yet been established.

In this work, we use HMC and ABC-SMC to find parameter probability distributions for two fairly low-dimensional cardiac action potential models: the Mitchell-Schaeffer model [27], which has two variables and five parameters, and the Fenton-Karma model [28],

which has three variables and 13 parameters. We test these two Bayesian methods using synthetic data and experimental recordings from zebrafish hearts and compare their performance.

2 Methods

Below we describe the cardiac action potential models used for fitting the data, the datasets to be fit, and the methods used for Bayesian inference, including details of our implementations.

2.1 Cardiac action potential models

In this work, we seek parameter value distributions for two cardiac action potential models: the three-variable Fenton-Karma model and the reduction of this model to two variables by Mitchell and Schaeffer. Because the model parameters will be referred to frequently, we include the model equations in full.

The Mitchell-Schaeffer (MS) model [27] uses two variables, the voltage u and inactivation gating variable h , along with inward, outward, and stimulus currents (I_{in} , I_{out} , and I_{stim} , respectively) to describe the transmembrane currents that give rise to action potentials.

$$\begin{aligned} \frac{du(t)}{dt} &= I_{in}(u, h) + I_{out}(u) + I_{stim}(t), \\ \frac{dh(t)}{dt} &= \begin{cases} \frac{1-h}{\tau_{open}}, & u < u_{gate} \\ \frac{-h}{\tau_{close}}, & u > u_{gate}, \end{cases} \end{aligned}$$

where

$$\begin{aligned} I_{in} &= h \frac{u^2(1-u)}{\tau_{in}}, \\ I_{out} &= -\frac{u}{\tau_{out}}. \end{aligned}$$

The five parameters include a threshold u_{gate} that determines the dynamics of the gating variable h ; the remaining four parameters are time constants that effectively govern the durations of the depolarization (τ_{in}) and repolarization (τ_{out}) phases as well as the closing (τ_{close}) and opening (τ_{open}) of the gate. Initial values were set to $u = 0$ and $h = 1$ for each cycle length used. The stimulus current was applied periodically according to the specified cycle length (CL) for 1ms with a magnitude of 0.66.

We also used the Fenton-Karma (FK) model [28], which is a phenomenological model that describes cardiac action potentials. It includes three state variables (voltage u and gating variables v and w) and 13 parameters.

$$\frac{du(t)}{dt} = -(I_{fi} + I_{so} + I_{si}) + I_{stim},$$

$$\frac{dv(t)}{dt} = \begin{cases} -\frac{v}{\tau_v^+}, & u \geq u_c \\ \frac{1-v}{\tau_v^-}, & u_c > u \geq u_v, \\ \frac{1-v}{\tau_{v1}}, & u < u_v \end{cases}$$

$$\frac{dw(t)}{dt} = \begin{cases} -\frac{w}{\tau_w^+}, & u \geq u_c \\ \frac{1-w}{\tau_w^-}, & u < u_c \end{cases}$$

where

$$I_{fi} = \begin{cases} -\frac{(1-u)(u-u_c)v}{\tau_d}, & u \geq u_c \\ 0, & u < u_c \end{cases}$$

$$I_{so} = \begin{cases} \frac{1}{\tau_r}, & u \geq u_c \\ \frac{u}{\tau_0}, & u < u_c \end{cases}$$

$$I_{si} = -\frac{w}{2\tau_{si}}(1 + \tanh(k(u - u_c^{si}))).$$

The fast inward current I_{fi} , the slow outward current I_{so} , and the slow inward current I_{si} represent summary sodium, potassium, and calcium transmembrane currents, respectively. The magnitude of the 1ms-long stimulus current was 0.35. Initial values were set to $u = 0$, $v = 1$, and $w = 1$ for each cycle length considered.

For both models, the differential equations were solved using an adaptive forward Euler scheme, with a timestep size of 0.1 ms for the first 4 ms after the beginning of the stimulus followed by an increase to 0.5 ms in all cases except when fitting the MS model to synthetic data, in which case the time step was increased to 0.25 ms.

Because HMC is a gradient-based method, discontinuities should be avoided. For this reason, the Heaviside functions in the MS and FK models were replaced by smooth functions when using HMC. For instance, in the case of the MS model, the equation for the gating variable was modified to

$$\frac{dh(t)}{dt} = \frac{1-h}{\tau_{open}}(1-p) - \frac{h}{\tau_{close}}p,$$

where $p = \frac{1}{2}(1 + \tanh(s(u - u_{gate})))$. The steepness parameter s was set to a large value of 50 for a steep transition resembling the Heaviside function.

Table 1 MS model parameter values and intervals: synthetic; experimental

	Dataset	Intervals for initial ABC-SMC priors	Center values for HMC folded normal priors	Initial values for HMC
τ_{in}	0.3	(0.1, 1); (0.1, 1.5)	0.3; 0.98	0.3; 0.3
τ_{out}	6	(3, 15); (1, 20)	6; 11.7	6; 6
τ_{open}	120	(50, 250); (10, 300)	120; 207	120; 120
τ_{close}	150	(100, 200); (10, 300)	150; 197	150; 150
v_{gate}	0.13	(0.1, 0.5); (0.1, 1)	0.13; 0.31	0.13; 0.13

2.2 Datasets

Both synthetic and experimental datasets were used to test the methods. The datasets include significant changes in action potential shape and durations as a result of rate changes, including a bifurcation to alternans (alternating long and short APs despite a constant CL) at the shortest CLs. Synthetic data were generated for the MS model using the default parameter set [27], with values given in the first column of Table 1. For the FK model, parameter set 4 from [29] was used; see the first column of Table 2. These sets of parameters are subsequently referred to as the true values for each model. Each dataset was derived from three selected CLs from voltage recordings (MS synthetic: 400, 350, and 310 ms; FK synthetic: 400, 350, and 300 ms; experimental: 350, 300, and 276 ms). For each chosen CL, data were generated by applying six stimuli and recording the action potentials induced by the last two stimuli. To form the synthetic datasets, these data were recorded at time intervals of 0.5 ms during the first 4 ms in order to adequately capture the upstroke, and then at intervals of 15 ms during the remainder of the action potential. For the experimental datasets, the temporal resolution was 0.4 ms for the first 4 ms, 1 ms for the next 4 ms, and 15 ms for the remainder of each cycle.

To make the datasets more realistic, Gaussian noise was then added with a mean of 0 and standard deviation of 0.03, which was smaller than the level of noise typically observed in cardiac optical-mapping voltage signals but higher than seen in microelectrode recordings.

The experimental dataset consisted of microelectrode recordings of voltage from zebrafish hearts obtained previously [26]; examples of action potentials are shown in Fig. 1. The original resolution of the data was 0.1ms. To form the dataset, a nonuniform resolution of 0.5ms was used for the first 4ms after applying the stimulus followed by an increase to 15ms until the next stimulus was applied. This approach allowed us to reduce the size of the dataset (and correspondingly the computational time) while retaining good accuracy during the upstroke. Voltage values were normalized to lie in the interval [0,1] using the maximum and minimum values in each voltage trace used.

Table 2 FK model parameter values and intervals: synthetic; experimental

	Dataset	Intervals for initial ABC-SMC priors	Center values for HMC folded normal priors	Initial values for HMC
τ_d	0.407	(0.03, 1); (0.03, 1)	0.407; 0.44	0.407; 0.3
τ_r	34	(1, 209); (1, 209)	34; 84.29	34; 110
τ_0	9	(1, 50); (1, 50)	9; 26.69	9; 20
τ_{si}	26.5	(5, 300); (5, 300)	26.5; 283.38	26.5; 280
τ_v^+	3.33	(1, 100); (1, 100)	3.33; 50.81	3.33; 27
τ_{v1}^-	15.6	(1, 300); (1, 300)	15.6; 84.81	15.6; 80
τ_{v2}^-	5	(1, 2500); (1, 2500)	5; 591.19	5; 350
τ_w^+	350	(1, 800); (1, 800)	350; 199.52	350; 200
τ_w^-	80	(1, 500); (1, 500)	80; 218.63	80; 200
u_c	0.15	(0.01, 0.3); (0.01, 0.3)	0.15; 0.15	0.15; 0.2
u_v	0.04	(0.001, 0.04); (0.001, 0.04)	0.04; 0.01	0.01; 0.01
u_c^{si}	0.45	(0.1, 1.5); (0.1, 1.5)	0.45; 0.36	0.45; 0.45
k	15	(1, 50); (1, 50)	15; 4.25	15; 5

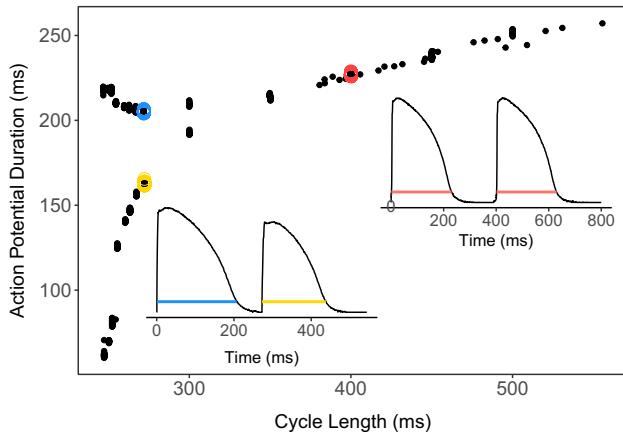


Fig. 1 Action potential duration as a function of cycle length for the full zebrafish dataset. Insets show action potentials at CLs of 400 and 275 ms, with alternans present for the shorter CL. Horizontal lines correspond to APDs

The experimental recordings included data obtained at multiple CLs (see Fig. 1), and synthetic data could be obtained for any CL of interest. For both synthetic and experimental datasets, the last two action potentials in a series of six obtained while pacing at a fixed CL were fitted to minimize transient behavior. After performing a series of initial experiments with different numbers and selections of CLs, we chose to use three CLs close to the bifurcation point, one at a long CL without alternans and two at shorter CLs within the alternans regime. Although the datasets to be fit utilized only three CLs, data from additional CLs were used to generate plots of action potential durations (APDs) as a function of CL for comparisons of results at other unfitted CLs.

When generating bifurcation plots, CLs were decreased until block was reached, and action potential durations (APDs) were measured using a fixed threshold of $u = 0.1$.

2.3 Bayesian inference

For a general system $\mathbf{x}'(t) = f(\mathbf{x}, t, \boldsymbol{\theta})$ where the system state $\mathbf{x} = \mathbf{x}(t) \in \mathbb{R}^n$ depends on time t and $\boldsymbol{\theta} \in \mathbb{R}^m$ is a vector of parameters, we consider noisy synthetic and experimental data of the form $y_i = \mathbf{x}(t_i; \boldsymbol{\theta}) + \epsilon_i$. Here y_i is the i^{th} observation of the state $\mathbf{x}(t_i; \boldsymbol{\theta})$ at time t_i , and ϵ_i is an independent normally distributed error with mean $\mu = 0$ and standard deviation σ . A central idea in Bayesian statistics is that because the vector of parameters $\boldsymbol{\theta}$ is fixed but unknown, it can be considered a multidimensional random variable, with the uncertainty in the parameter values described by a probability model. By Bayes' theorem,

$$p(\boldsymbol{\theta}|\mathbf{y}) \propto p(\mathbf{y}|\boldsymbol{\theta})p(\boldsymbol{\theta}),$$

where $p(\boldsymbol{\theta}|\mathbf{y})$ is the final distribution of the parameters conditioned on the data (also known as the target or posterior distribution), $p(\mathbf{y}|\boldsymbol{\theta})$ is the likelihood of the data given the parameters, and $p(\boldsymbol{\theta})$ is the prior or initial distribution of the parameters.

Although obtaining an analytical form for the final distribution generally is not possible, we can sample from the final distribution using either a full Bayesian method or an approximate approach in which the likelihood is not computed. Examples of full Bayesian approaches are Markov chain Monte Carlo (MCMC) [30] methods like Metropolis-Hastings, Gibbs sampling, and Hamiltonian Monte Carlo. For MCMC methods, the likelihood is given by

$$p(\mathbf{y}|\boldsymbol{\theta}, \sigma) = \prod_{i=1}^T \frac{1}{\sqrt{2\pi} \sigma} e^{-\frac{1}{2} \left(\frac{y_i - u(t_i; \boldsymbol{\theta})}{\sigma} \right)^2},$$

where σ is the standard deviation distribution of the error, which is considered Gaussian centered at 0 with variance σ , and T is the number of time points. An example of an approximate Bayesian computation (ABC) method is rejection sampling [31].

The objective of MCMC methods is to design a Markov chain in such a way that its stationary distribution coincides with that of the target distribution. The Metropolis (or Metropolis-Hastings) algorithm builds an adaptive random walk that converges to the target distribution. The chain is constructed by using a rule that accepts (or rejects) candidates sampled from a known distribution from which it is easy to sample along with a transition probability distribution. For the symmetric proposals used here [32], candidates with a higher probability than the most recent member of the chain are always accepted; to allow the parameter space to be explored, candidates with a lower probability are sometimes accepted based on an acceptance criterion (the Metropolis ratio) that selects more (less) likely candidates more (less) often. After enough iterations, the result is a correlated sample of the target distribution. To obtain a non-correlated sample, only every k elements of the original sample are selected. However, MCMC often is slow to converge and scales poorly as the dimensionality of $\boldsymbol{\theta}$ increases, motivating the development and use of alternative methods.

2.3.1 Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (HMC) is a Metropolis method that uses the gradient of the target distribution to form a Hamiltonian system by taking the parameters given the data as the position variables and adding momentum variables \mathbf{q} [16]. In this way, the properties of Hamiltonian dynamics can be used to sample efficiently from the distribution of the variables of interest, $P(\boldsymbol{\theta}) = p(\boldsymbol{\theta}|\mathbf{y})$. Hamiltonian dynamics have many properties that are crucial in constructing MCMC updates, such as reversibility, invariance, and symplecticness [16].

A Hamiltonian function is formed as

$$H(\boldsymbol{\theta}, \mathbf{q}) = E(\boldsymbol{\theta}) + K(\mathbf{q}),$$

where $E(\boldsymbol{\theta})$ is a canonical distribution (defined by $P(\boldsymbol{\theta}) = \exp(-E(\boldsymbol{\theta}))/Z$) and $K(\mathbf{q}) = \mathbf{q}^T M^{-1} \mathbf{q}$ represents the kinetic energy; here Z is a normalizing constant and M is a positive definite matrix that rotates and rescales the target distribution in the Euclidean space [32]. The joint density has the property that $P_H(\boldsymbol{\theta}, \mathbf{q}) = P(\boldsymbol{\theta})P(\mathbf{q})$, which makes recovery of $P(\boldsymbol{\theta})$ straightforward.

The position and momentum variables evolve according to differential equations that are solved using the leapfrog integrator, a finite-difference method specifically designed to solve dynamical systems in classical mechanics. To implement HMC, the step size and number of steps generally need to be tuned, which typically is a challenging task. However, the No-U-Turn sampler (NUTS) [33] avoids this tuning step by determining the step size and number of steps adaptively in each iteration. Since HMC is also a Metropolis algorithm, candidates are accepted or rejected using the Metropolis ratio, this time to account for numerical error in the leapfrog algorithm. More details can be found in [16].

HMC utilizing NUTS was implemented in R using the statistical platform Stan [34, 35]. To ensure that the Markov chain had time to find the region of interest in parameter space, only data after the first 1000 iterations were considered to be samples from the true posterior; the use of such a burn-in period is standard for MCMC methods. The \hat{R} statistic in Stan was used to verify that the chains were converging to the target distribution [35–37]. A sample size of 500 was used for the posterior distributions; only 100 randomly selected members of the posterior sample were used to generate the figures below involving action potentials and APDs to improve clarity.

Because we found HMC to be quite sensitive to the choice of prior distribution, it was necessary to obtain somewhat informative priors in order to achieve reliable results. For synthetic data using the MS and FK models, the priors used were folded normal distributions, which prevented the distributions from extending into unphysical negative values; the model parameters, which included time constants, thresholds, and a steepness value, could not meaningfully take on negative values. The standard deviations were set to 20% of the values of the means. Initial values needed for HMC were set as the modes of the priors. For experimental data, priors were obtained following a process that involved using ABC-SMC (which is described in the following section), which is more accepting of less informative priors, with wide uniform priors that contained but were not centered around the values used to generate the synthetic data, then using the modes of the resulting marginal posterior distributions as the values the HMC priors are centered around. The initial priors used by ABC-SMC were defined on the intervals given in the third column of Table 1. The resulting priors used for HMC for the MS model with experimental data in the results to be shown were folded normals extending 20% around the values given in the fourth column of Table 1; these values also were used as the initial values for HMC, as indicated in the last column.

For the FK model, the process of obtaining the priors was the same as for the MS model, with ABC-SMC used with wide uniform priors to generate a more informative prior subsequently used by HMC. The initial uniform priors used by ABC-SMC were defined on the intervals given in the third column of Table 2. The priors used for HMC with the FK model were folded normals extending 20% around the values given in the fourth column of Table 2, and the initial values were set to the values in the last column of the table.

2.3.2 Approximate Bayesian computation sequential Monte Carlo

When the likelihood $p(y|\theta)$ is difficult to evaluate or is not available, ABC algorithms can be applied. One type of ABC algorithm is rejection sampling, in which results generated using candidate parameterizations are compared with the data by using a distance function and a tolerance; only those candidates that produce output sufficiently close to the data (within the tolerance) are accepted. The set of accepted candidates $\{\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(N)}\}$ is called a population. ABC rejection's acceptance rate typically is very low, which makes it inefficient.

Using information from the previous population to build a sample iteratively can result in higher efficiency. One such method is ABC sequential Monte Carlo (ABC-SMC) [19]. Sample points from the previous population are weight-sampled and then perturbed using a random walk with a uniform or Gaussian kernel until a certain tolerance is reached. More information can be found in [19].

The ABC-SMC implementation we used was custom written in R and utilized the modifications made to the ABC-SMC algorithm introduced by [19] and also found in [22], including adaptive tolerances and the use of the effective sample size for each population. In addition, to improve convergence and the exploration of the parameter space, we used a decreasing sequence of values for the scale factor for perturbing the populations. As suggested in [22], we used a probability density function as the distance function to generate the first population. Later populations were constrained to produce output closer to the target data using the sum of squared error by choosing a series of smaller tolerances. The tolerance reduction could be modified adaptively if found to be too restrictive and the algorithm stopped when the tolerance reduction was smaller than a specified value; see ref. [22] for more detailed information.

The probability density function used as a distance to obtain the first population was $\rho_{\epsilon_1}(\theta) = \exp \frac{-SE(\theta)/2(\epsilon_1\sigma)^2}{c_{\epsilon_1}}$, where $SE(\theta)$ represents the sum of squared error between the data to fit and the model solution at each time point. Here $\epsilon_1 = 1$ is the first tolerance value, σ is the standard deviation

of the error or noise assumed in the data measurements, and c_{ϵ_1} is a normalizing constant; while the exact constant is unknown, a good approximation can be selected by monitoring the acceptance of candidates [22]. If the value is too high, most samples will be rejected and the computational time will increase, whereas if the value is too low, most samples will be rejected and the result can approach a uniform sample regardless of the true distribution. We found that a value of 1000 achieved a good balance and thus used this value for all cases. The use of a probability density function can help in avoiding an overestimation of the variance in the first population that can occur otherwise. For all subsequent populations, the distance function was simply $SE(\theta)$. Once the first population was calculated, the value of the first tolerance was updated to the sum of squared error between the data and the model solution obtained using the modes of the distributions for each parameter, and the subsequent tolerance was set to be $\frac{3}{4}$ of the updated first tolerance. The value by which the tolerance was reduced for subsequent populations was decreased by a factor of 0.5 but could be modified if too strict, following ref. [22], and the minimum tolerance reduction serving as a stopping criterion was set to 1.5625×10^{-3} .

The population size chosen for ABC-SMC was 500, ensuring a posterior distribution the same size as for HMC. As for HMC, a random selection of only 100 population members was used to generate the action potential and APD figures to improve clarity. For synthetic data using the MS and FK models, the priors used were folded normals extending 20% around the true values. As discussed in the previous section, it was necessary to obtain more informative priors for use with HMC in conjunction with the experimental data, and ABC-SMC was used with an initial uniform prior to generate these priors. However, for the results shown below, ABC-SMC uses the same priors as those generated for use by HMC to allow for a fair comparison.

2.3.3 Assessing accuracy

In the absence of known “correct” distributions, we will assume that HMC, as a full Bayesian method, is more likely to converge to accurate distributions. Within Stan, HMC by default is implemented to utilize four chains, which provides additional checks and accuracy compared to single-chain results, which we use for comparisons with ABC-SMC in the interest of fairness. Thus, we chose to use the default four-chain HMC results from Stan as the “gold standard” for the distributions. We verified that such cases converged according to established metrics in Stan (e.g., \hat{R} values very close to one were used to verify that the chains were converging to the target distribution), produced no or very few warnings within Stan (e.g., no or few divergent transitions,

adequate effective sample size that if too small could indicate unreliable posterior means or variances), and achieved fits to the data with low error. Except where noted, all other cases similarly converge according to the same metrics. For all scenarios considered, the Supplementary Information includes four-chain results, which in all cases are extremely similar to the one-chain results. Note that for the four-chain results, only one chain uses the initial values noted here; the rest use initial values formed as perturbations from these values with Gaussian noise using a 10% deviation.

In addition, for the synthetic data, to show that the parameters could be recovered, we verified that the mean values for the posterior parameter distributions were between quantiles 10 and 90. Although with HMC the noise parameter can be estimated, ABC-SMC makes no assumptions about the error; therefore, we also use the coefficient of determination $R^2 = 1 - \frac{SS_{res}}{SS_{tot}}$, where SS_{res} is the residual sum of squares (the same as $SE(\hat{\theta})$, where the argument is the mean of the posterior distribution) and SS_{tot} is the total sum of squares comparing the data points with their average. For linear systems, this measure explains how well observed outcomes are replicated by

the model, based on the proportion of total variation of outcomes explained by the model [38]. Because the models we are fitting are nonlinear, we will not consider the metric in that context but instead will use it as a type of normalized error metric where the desired value is one. To avoid potential confusion with the use of R^2 in linear regression, we will refer to this normalized error metric as a pseudo- R^2 value.

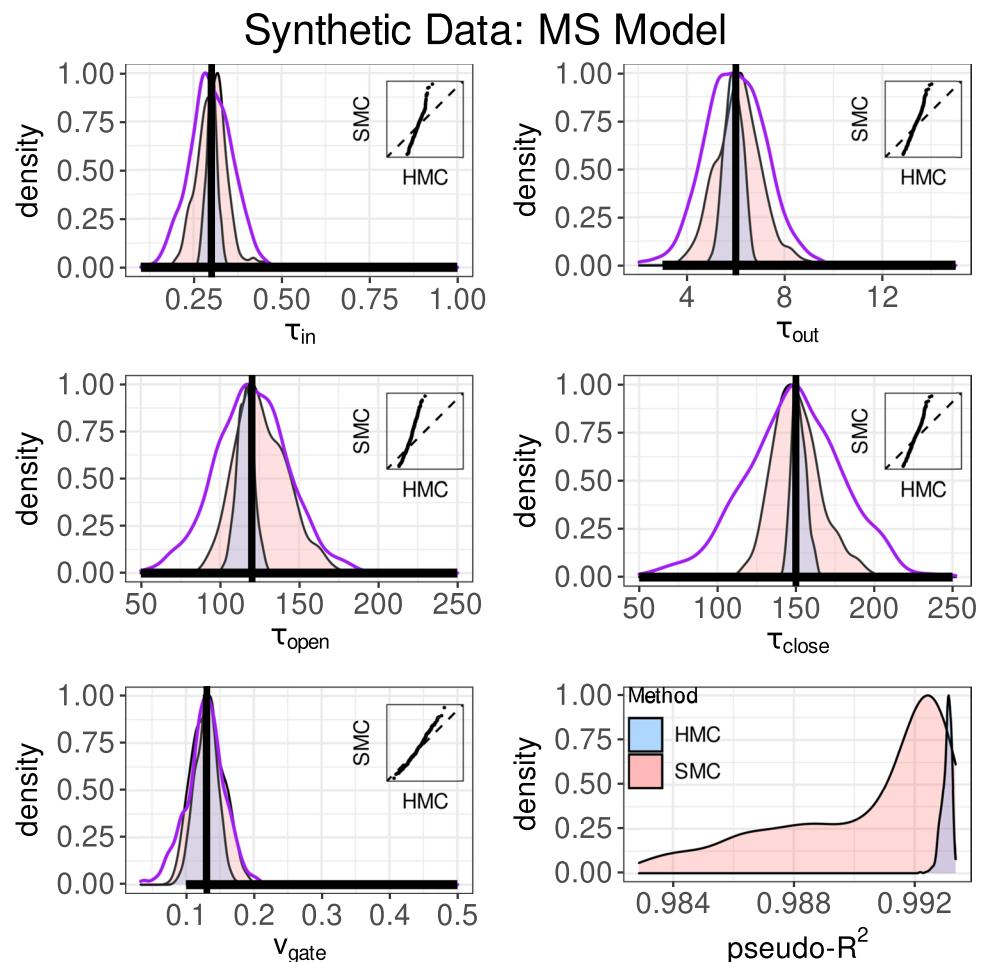
3 Results

Below we present our results using HMC and ABC-SMC to obtain distributions of parameters for the MS and FK models for the two datasets.

3.1 Synthetic dataset

First we consider parameter distributions for the MS model. Figure 2 shows the marginal distributions of the five MS parameters and of the pseudo- R^2 normalized error metric across the final population for HMC (blue) and ABC-SMC

Fig. 2 Marginal distributions for the MS model parameters along with the pseudo- R^2 distribution when fitting to synthetic data (HMC, blue; ABC-SMC, red). The priors used for both methods are shown in magenta. Horizontal lines indicate the support of the non-informative uniform prior used initially by ABC-SMC to generate the informative prior shown in magenta. Vertical lines represent the true values used to generate the synthetic data. Insets show q - q plots for ABC-SMC (vertical) vs. HMC (horizontal); dashed lines have slope one



(red); the true values used to generate the original APs are shown as black vertical lines for comparison. The distributions for all parameters are fairly well centered around the true values for both methods, but the distributions obtained using HMC are narrower than those from ABC-SMC. For comparison, the four-chain HMC results for this scenario align almost exactly with the one-chain HMC results, as shown in Fig. S1, with slight shifts for a few parameters but within a very narrow region of parameter space. Note also that the shapes of the distributions for HMC and ABC-SMC can be quite different; in particular, the distributions from ABC-SMC are less symmetric than those from HMC. The priors used for both cases, shown in magenta, have similar locations but are broader than the resulting posteriors; the non-informative uniform priors used to generate the informative priors in magenta are even broader. Thus, both HMC and ABC-SMC resulted in considerable narrowing of the distributions.

The pseudo- R^2 plots in Fig. 2, which reflect the distributions across the populations for each method, show values that are consistently close to one, above 0.983 for ABC-SMC and above 0.99 for HMC, that indicate low error in the data fittings. In addition, the final population for HMC has less variability in pseudo- R^2 than that for ABC-SMC. The insets correspond to q - q plots of the distributions obtained using ABC-SMC vs. HMC, where the ranges for each axis are given by the corresponding x -axis limits. For most parameters, the q - q plots suggest that the posteriors come from the same distribution family, even though the ABC-SMC distributions in general are wider, indicating that ABC-SMC likely explores more of the parameter space whereas HMC appears to converge to a narrower distribution.

The population members in each case produce action potentials that agree well with those of the dataset for all three CLs used, as shown for a random selection of 100 of the 500 population members in the posterior samples for each method in Fig. 3A–C. Figure 3D shows that the populations obtained from both methods agree well with the data at CLs not used during the fitting, including the location of the bifurcation to alternans. Some parameterizations from the ABC-SMC population can be paced slightly faster than the true parameters before block occurs.

In some cases, the marginal distributions in Fig. 2 appear somewhat broad, but combinations of parameters may be better constrained. Figure S2 shows bivariate scatterplots for all possible parameter pairs for the posterior sample and indicate that some pairs are correlated and may be well constrained when treated together. For example, τ_{in} and τ_{out} appear positively correlated, which is not unexpected given that these represent the reciprocals of inward and outward current conductances, respectively, and increases in one conductance could be offset by increases in the other. Both HMC and ABC-SMC also find negative correlations

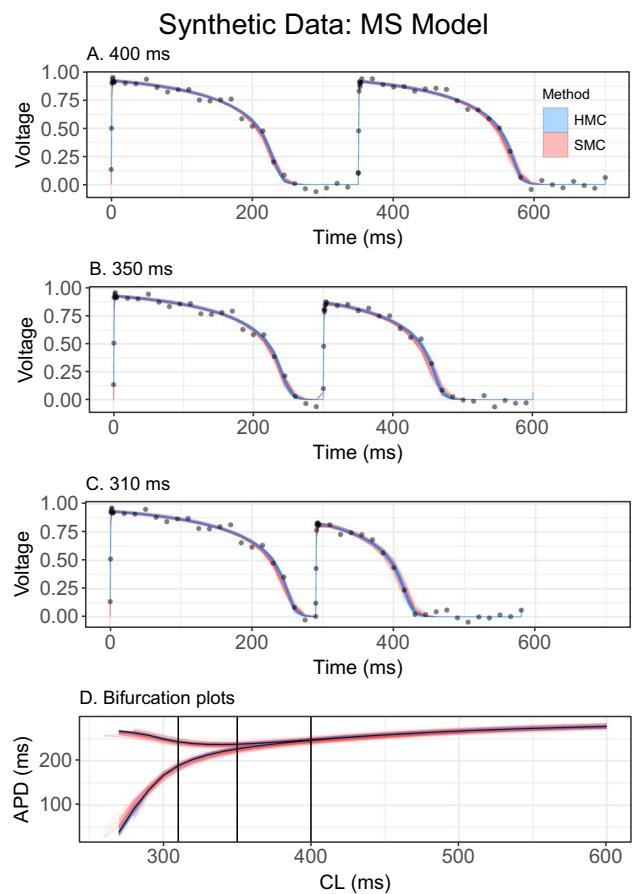


Fig. 3 A–C One hundred action potentials from populations of size 500 using synthetic data for HMC (blue) and ABC-SMC (red) compared with the data points for the MS model. CLs of (A) 400, (B) 350, and (D) 310 ms are included. D One hundred representative plots of APD as a function of CL for synthetic data with the MS model. CLs from 600 ms down to where block occurs are included, with HMC results in blue, ABC-SMC results in red, and the true data points in black. Note that fitted voltage data are output at the same timepoints as the dataset, leading to occasional small plotting artifacts when the points are connected by lines, especially at the beginnings of action potentials

between τ_{in} and τ_{open} (essentially how quickly the inward current turns on and how quickly the gate recovers) and between τ_{out} and τ_{close} (representing how quickly the outward current turns on and how quickly the gate closes, and possibly between τ_{open} and τ_{close} (which control the rate of recovery of the gating variable at different voltages)). HMC also finds a clear positive correlation between τ_{open} and v_{gate} that is not picked up by ABC-SMC, which represents that an increase in the time scale for gating variable recovery can be offset by a decrease in the excitability threshold. ABC-SMC appears to identify a possible weaker positive correlation between τ_{in} and τ_{close} , which suggests that a larger inward current magnitude (smaller τ_{in}) can be compensated by faster closing of the gate. As shown in Fig. S3, the four-chain HMC parameter

correlations agree well with the single-chain HMC results, but with greater definition because the four-chain results include four times as many samples and thus four times as many points in the scatterplots. Overall, although the parameter correlations indicate more structure of the multivariate posterior within the parameter space, the correlated parameters also suggest that identifying the true values of the individual parameters within the correlated pairs may be difficult from the dataset.

For the FK model, the results from the two methods show more variation. Although the marginal distributions for the 13 model parameters largely are centered around the true values, as shown in Fig. 4, the shapes of the distributions for each parameter frequently are visibly different. Most of the marginal distributions obtained using ABC-SMC are fairly broad, again suggesting broader exploration of parameter space, while HMC converges to a narrower distribution. For comparison, the four-chain HMC results again align very closely with the single-chain results shown here; see Fig. S4. For some parameters, including τ_d , u_c^{si} , and τ_w^- , the distributions from HMC are narrow compared to ABC-SMC, whereas for others the distributions are similarly broad. Nevertheless, in nearly all cases, there is considerable narrowing of the marginals with respect to the priors, shown in magenta. The non-informative uniforms used to generate these priors were even broader and for clarity are not shown in Fig. 4, but they are the same as those used for the experimental dataset discussed later. The pseudo- R^2 values also vary much more across the population for HMC than for ABC-SMC, although even for ABC-SMC the values are generally above 0.975 and thus are close to the desired value of one. The q - q inset plots are still close to linear, indicating that the distributions come from the same family, even though they do not always have a slope of one, especially for τ_d , u_c^{si} , and τ_w^- , which feature wider ranges of values from ABC-SMC compared to HMC. In the case of a few parameters like τ_w^+ and τ_0 , the two methods produce distributions centered around different values, with the true value located between the peaks.

The fitted action potentials reflect the apparent increase in variability when using the FK model compared to the MS model. Figure 5 shows a subsampling of the final populations obtained using each method and reflects the greater variation with the samples from ABC-SMC. Nevertheless, good agreement is obtained across a wide range of CLs, even for CLs not close to those used for fitting. In particular, the bifurcation to alternans is well characterized for both methods.

Bivariate scatterplots for parameter pairs, shown in Fig. S5, suggest some correlations across the set of FK model parameters, although in general a smaller fraction of possible pairings are correlated than for the MS model, and HMC and ABC-SMC do not always detect the same

correlations. Both HMC and ABC-SMC identify strong positive correlations between τ_r and τ_{si} , which makes sense as these govern the strength of the repolarizing current and slow inward (calcium) currents, so that an increase in one can compensate for an increase in the other. HMC also identifies strong negative correlations between τ_d and u_c (increasing the excitability by decreasing τ_d can be offset by increasing the excitation threshold), τ_{si} and τ_w^+ (increasing the APD by decreasing τ_{si} can compensate for a slow inward current prolonged via increased τ_w^+), and τ_r and τ_w^+ (speeding up repolarization can be offset by prolonging the slow inward current); ABC-SMC detects only the last of these. Both HMC (strongly) and ABC-SMC (weakly) detect a positive correlation between τ_d and τ_v^+ (decreased fast inward current via larger τ_d can be compensated by prolonging the same current via slowing its inactivation gate).

Figure S5 suggests additional (weaker) correlations that are apparent from only one method. For example, HMC detects additional positive correlations between τ_d and u_c^{si} and between τ_{si} and u_c^{si} as well as negative correlations between τ_w^+ and u_c^{si} , between τ_w^+ and τ_w^- , and between τ_v^+ and u_c . The scatterplot for ABC-SMC suggests another weak correlation between τ_v^+ and τ_{v1}^- that is not detected by HMC.

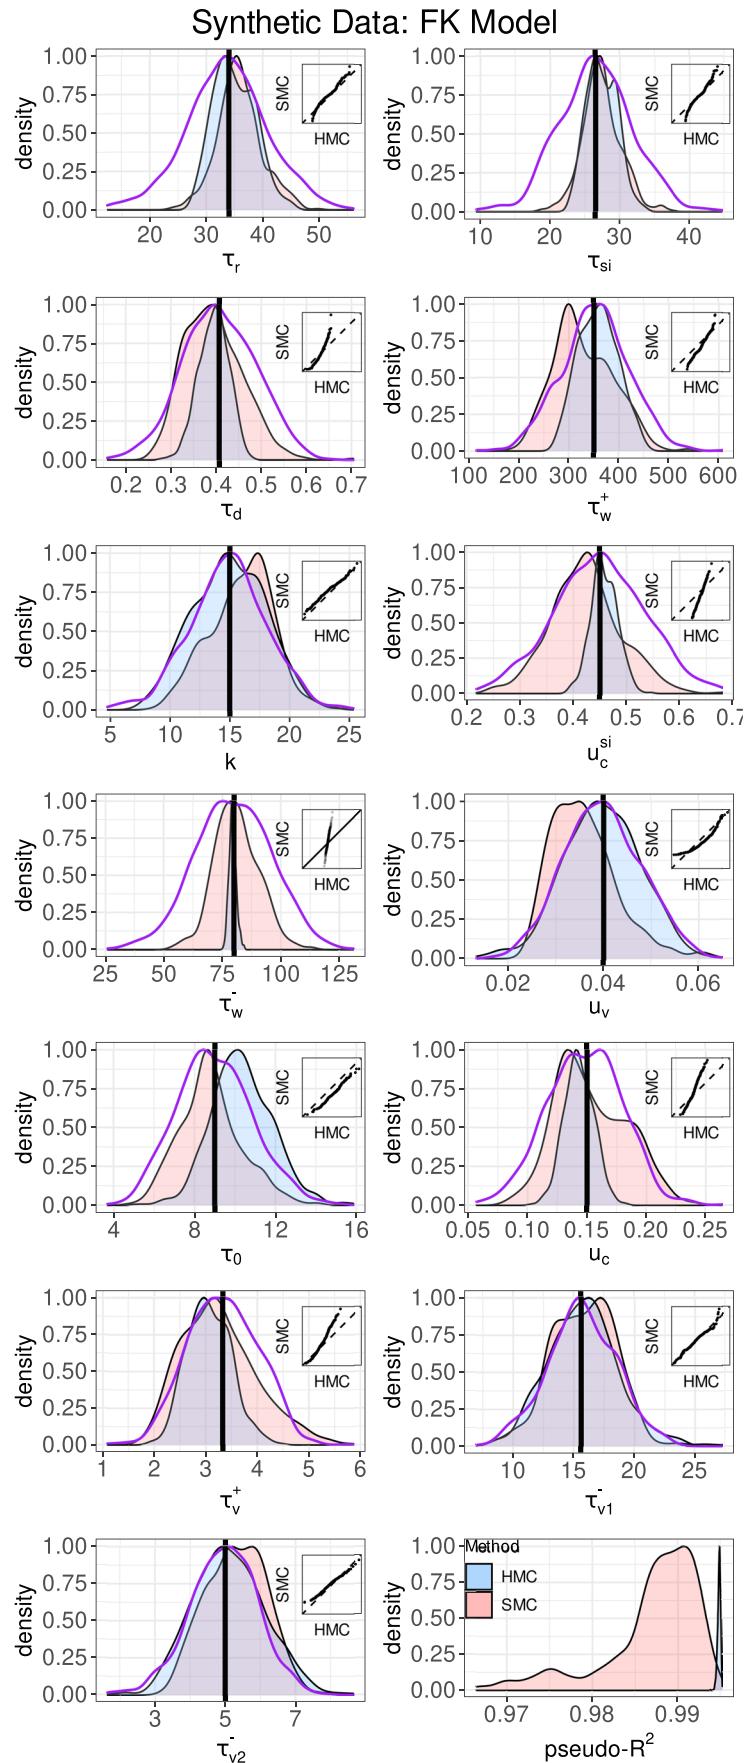
Note that the HMC parameter correlations obtained using four chains and those using one chain agree well, as shown in Fig. S6.

3.2 Experimental data

For the experimental data, the results for the MS model show some differences compared to the synthetic data. As demonstrated in Fig. 6, the marginal distributions from HMC and ABC-SMC are nearly identical for τ_{open} ; for τ_{in} and τ_{out} , the distributions appear slightly offset relative to each other. For the remaining parameters, τ_{close} and u_{gate} , the distributions look similar to the results for the synthetic data, with a broader range of values included for ABC-SMC. The agreement in marginal distribution shapes compared to the synthetic data case leads to linear q - q plots with slope values of nearly one in all cases, consistent with the distributions being from the same families. For comparison, the distributions obtained using four chains vs. one chain for HMC are nearly identical, as shown in Fig. S1. The pseudo- R^2 values are not as close to one when the data are not derived from the model being fit; however, pseudo- R^2 remains above 0.65 for ABC-SMC and above 0.7 for HMC.

Figure 7 shows the fitting to experimental data for 100 of the 500 population members for two of the three CLs, 300 ms and 276 ms. Again, good agreement is seen overall, but with more variability in AP shapes and durations across the population with ABC-SMC. Of particular note, the variability seen in the ABC-SMC posterior sample includes longer and shorter APs, but the HMC sample is

Fig. 4 Marginal distributions for the FK model parameters along with the pseudo- R^2 distribution when fitting to synthetic data (HMC, blue; ABC-SMC, red). The priors used for both methods are shown in magenta. The support of the non-informative uniform prior used initially by ABC-SMC is not shown to improve figure clarity but is the same as that shown in Fig. 8 for the experimental data case. Vertical lines represent the true values used to generate the synthetic data. Insets show q - q plots for ABC-SMC (vertical) vs. HMC (horizontal); dashed lines have slope one



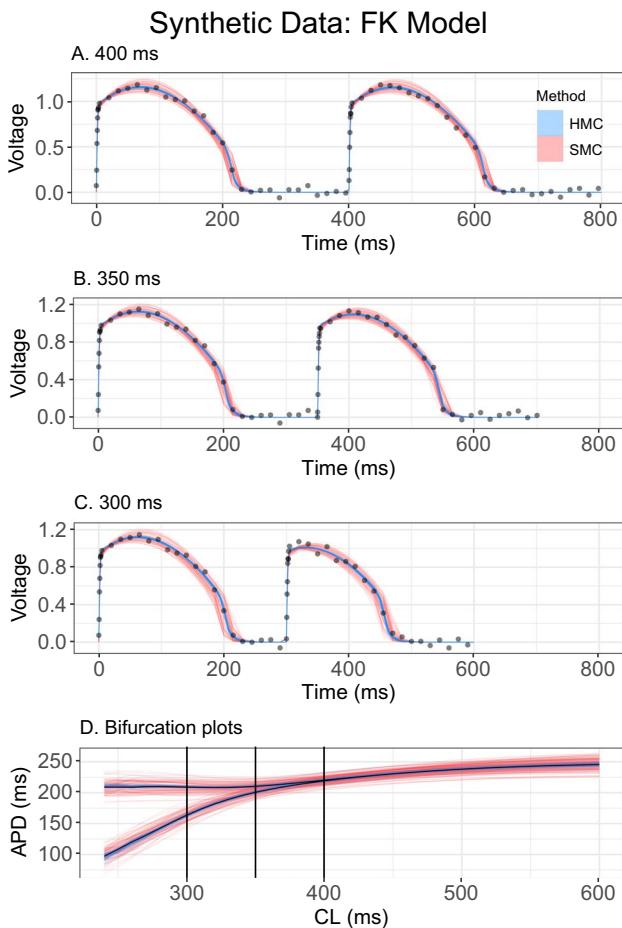


Fig. 5 A–C One hundred action potentials from populations of size 500 using synthetic data for HMC (blue) and ABC-SMC (red) compared with the data points for the FK model. CLs of (A) 400, (B) 350, and (C) 300 ms are included. D One hundred representative plots of APD as a function of CL for synthetic data with the FK model. CLs from 600 ms down to where block occurs are included, with HMC results in blue, ABC-SMC results in red, and the true data points in black

biased toward longer APs for the MS model. This trend can be observed across a broad range of CLs; HMC generally achieves longer APDs than occurred in the experiment. In addition, neither method accurately captures the dynamics associated with alternans, including the bifurcation point and the magnitude of the alternans at the shortest CLs.

Figure S7 depicts bivariate scatterplots for the parameter pairs. A positive correlation between τ_{in} and τ_{out} and a negative correlation between τ_{close} and τ_{out} are indicated by both HMC and ABC; both correlations were observed for the case of synthetic data as well. In addition to these cases, a weak negative correlation between τ_{in} and τ_{open} , also observed for the synthetic data case, is suggested by HMC only, and a weak correlation not observed for synthetic data is suggested by both HMC and ABC between τ_{in} and τ_{close} .

The use of four chains for HMC, with its corresponding increase in the number of points in the scatterplots, only defines correlations more clearly and does not suggest any new correlations; see Fig. S8. Note that no correlations are apparent between v_{gate} and any other parameter when fitting the experimental dataset.

For the FK model, the marginal distributions of the parameters show even more variability, as shown in Fig. 8, with wider ranges of values in many cases obtained using HMC. Although the distributions have similar peak values for some parameters, such as τ_r and τ_w^- , they appear shifted for others, such as τ_d , τ_w^+ , and u_v . The q - q plots for these shifted cases look less than linear, which may suggest that in some cases the distributions may not be from the same family. Nevertheless, the distributions obtained using HMC are nearly identical regardless of the number of chains used, as shown in Fig. S4. In addition, the pseudo- R^2 values are somewhat closer to one than they were for the MS model for the same data. Despite this variability in the values of parameters, the fitted action potentials for the FK model closely fit the data, as shown in Fig. 7 (right column). Both methods show less variability across the sample for the FK model than for the MS model. However, there is a tendency for the upstroke to be fitted poorly for higher voltage values for CLs exhibiting alternans. Despite this fact, the dynamics overall, especially during alternans, are much better captured by the FK model than by the MS model.

When fitting experimental data with the FK model, fewer parameter correlations are apparent than when fitting synthetic data, as shown in Fig. S9. Both HMC and ABC-SMC detect a strong negative correlation between τ_r and τ_v^+ (increasing the repolarizing slow outward current by decreasing τ_r can be offset by increasing τ_v^+ and thus inactivating the fast inward current more slowly). HMC suggests another negative correlation, between τ_d and τ_{v2}^- (a decrease in excitability by increasing τ_d can be offset by decreasing τ_{v2}^- , which makes it easier to initiate a new action potential by accelerating the recovery of the fast inward current inactivation gate). HMC finds the same correlations regardless of the number of chains used (see Fig. S10), and neither correlation is apparent when fitting synthetic data. The much smaller number of correlations compared to when fitting synthetic data may reflect the difficulty the model faces when fitting alternans from the zebrafish experiments.

3.3 Consistency and robustness

Both HMC and ABC-SMC are dependent on randomization within their algorithms; as a result, it is possible that results could be different every time either algorithm is executed if the algorithm is not appropriately sampling the posterior. To assess whether the results are robust

Fig. 6 Marginal distributions for the MS model parameters along with the pseudo- R^2 distribution when fitting to experimental data (HMC, blue; ABC-SMC, red). The priors used for both methods are shown in magenta. Horizontal lines indicate the support of the non-informative uniform prior used initially by ABC-SMC to generate the informative prior shown in magenta. Insets show q - q plots for ABC-SMC (vertical) vs. HMC (experimental); dashed lines have slope one

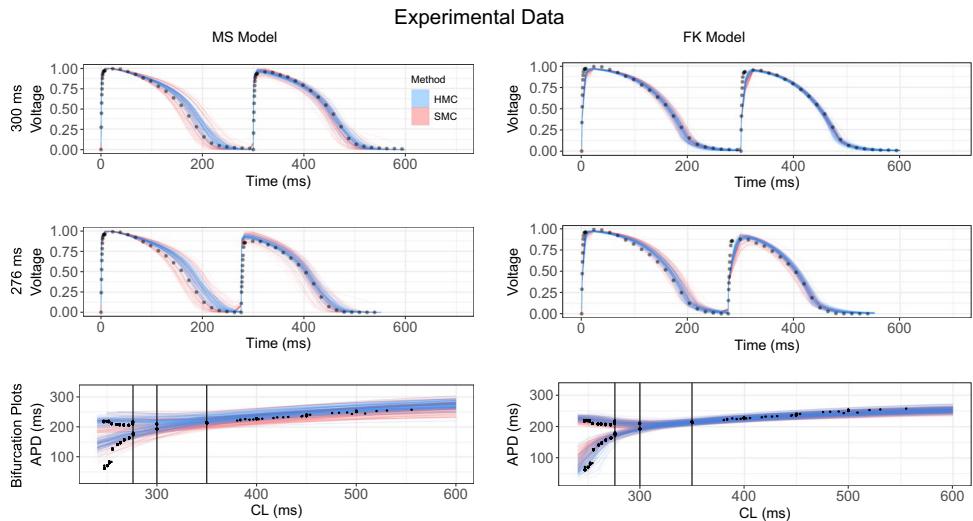
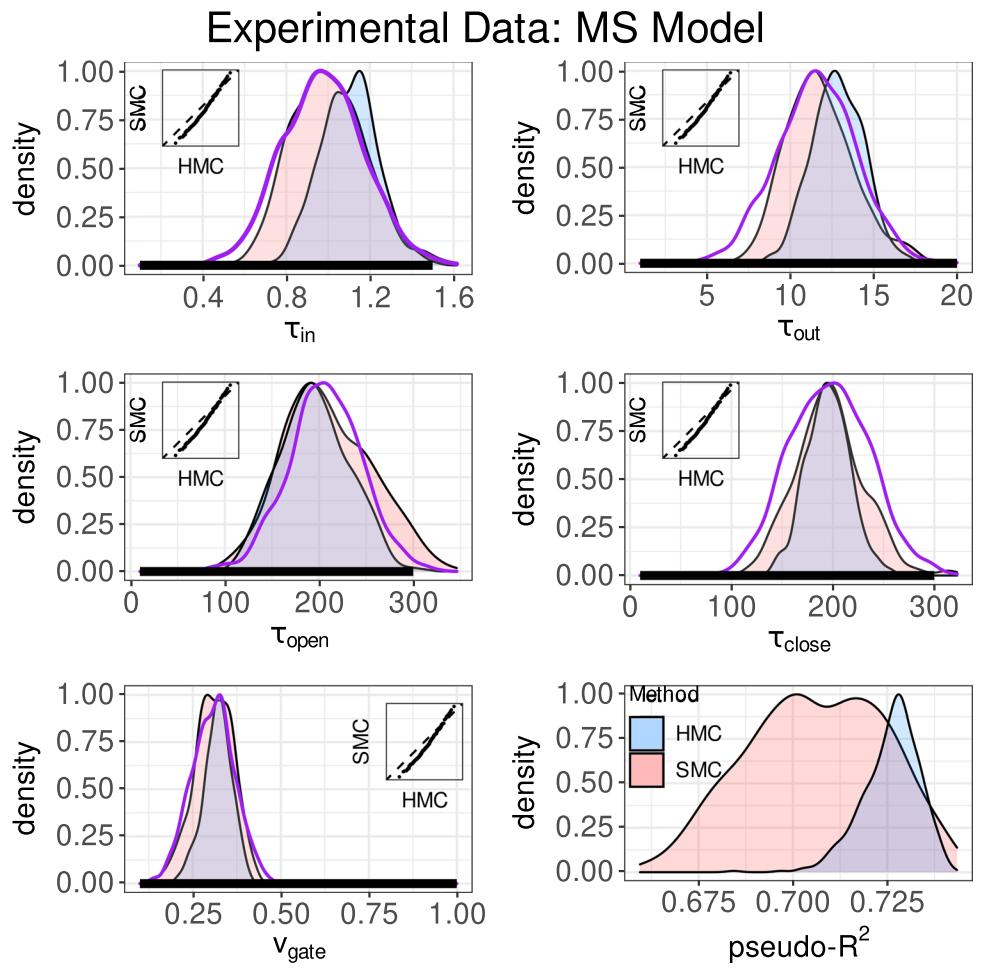
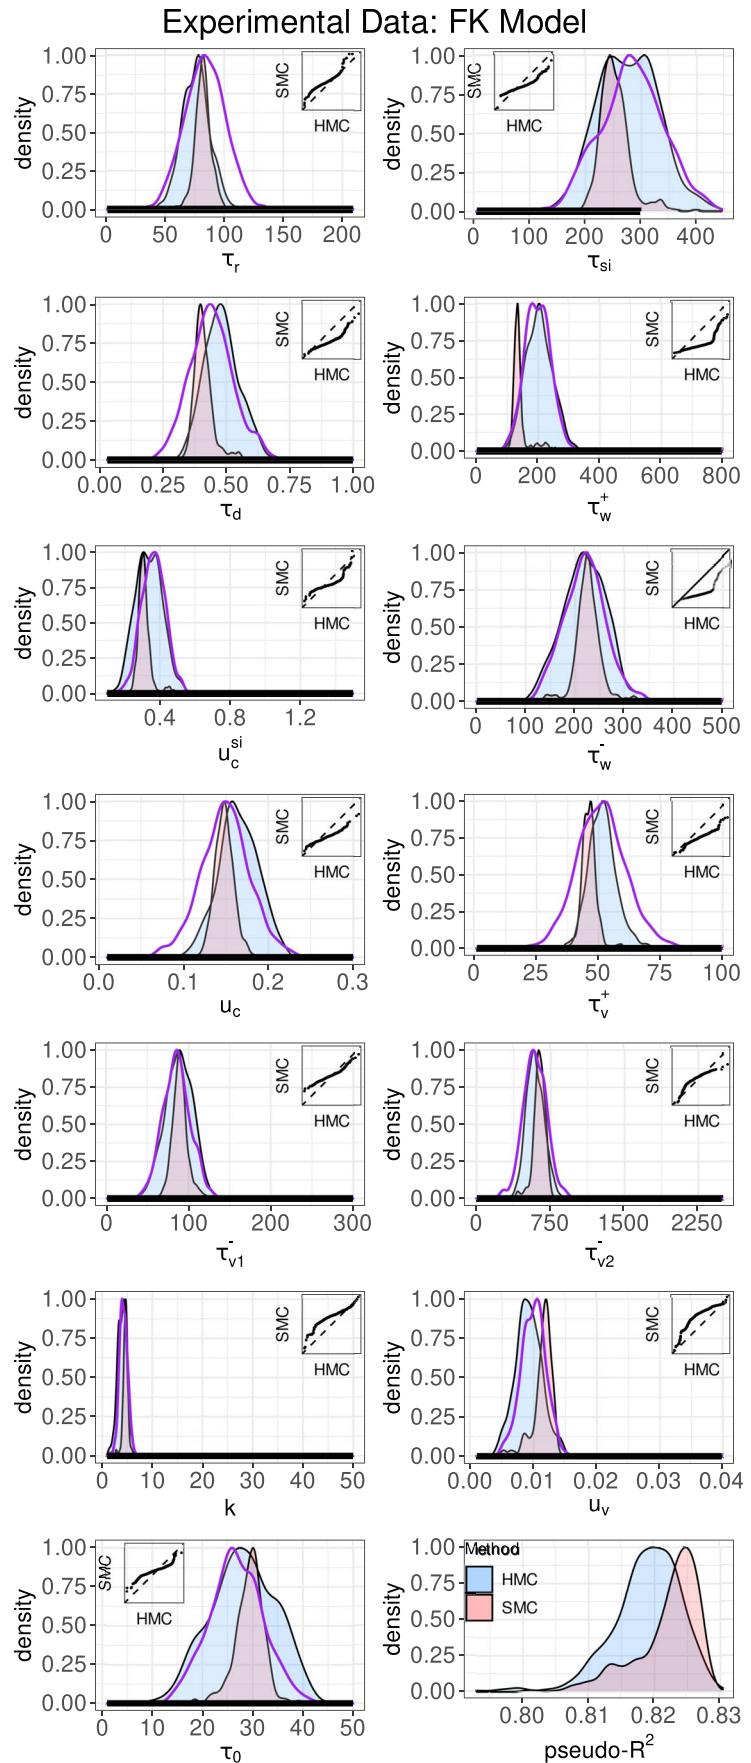


Fig. 7 Fit action potentials and bifurcation plots for the MS model (left) and the FK model (right) using HMC (blue) and ABC-SMC (red) for experimental data (shown in black). Top two rows show action potentials for CLs of 300 and 276 ms from a sample of 100 randomly selected population members. Bottom row shows APD as

a function of CL for the same sample. Note that fitted voltage data are output at the same timepoints as the dataset, leading to occasional small plotting artifacts when the points are connected by lines, especially at the beginnings of action potentials

Fig. 8 Marginal distributions for the FK model parameters along with the pseudo- R^2 distribution when fitting to experimental data (HMC, blue; ABC-SMC, red). The priors used for both methods are shown in magenta. Horizontal lines indicate the support of the non-informative uniform prior used initially by ABC-SMC to generate the informative prior shown in magenta. Insets show q - q plots for ABC-SMC (vertical) vs. HMC (horizontal); dashed lines have slope one



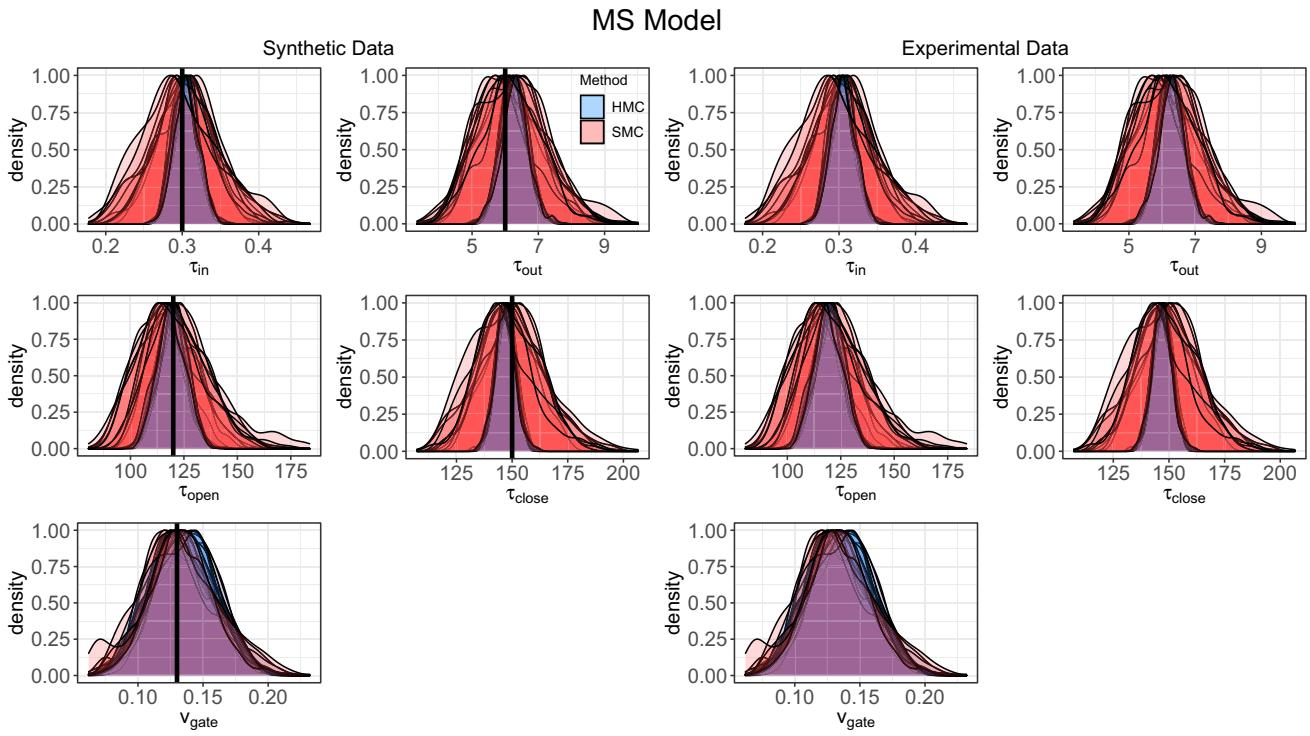


Fig. 9 Superimposed marginal distributions for the MS model parameters when fitting to (left) synthetic and (right) experimental data (HMC, blue; ABC-SMC, red) from 10 runs of each algorithm (HMC,

blue; ABC-SMC, red). Vertical lines represent the true values used to generate the synthetic data

to different randomizations, we ran each algorithm ten times for each dataset and for each model. Figure 9 summarizes the resulting distributions for the MS model for both the synthetic (left) and experimental (right) data. In both cases, the distributions obtained are fairly consistent for each method. For the synthetic data, the marginal distributions are unimodal and centered near the correct values. However, the marginal distributions for all time constants are wider for ABC-SMC than for HMC; for ν_{gate} , the distribution widths are similar for both methods. For the experimental data, there is even greater consistency across the ten runs for each method along with more similarity of the posterior distributions for the two methods. The distributions generally are wider with ABC-SMC, but the difference in widths is less pronounced than with the synthetic data. Some differences in the locations of the peaks for τ_{in} and τ_{out} are evident, as discussed earlier when using experimental data for the MS model in Fig. 6.

For the FK model, more variation occurs for both datasets. As shown in Fig. 10, across the ten runs, the posteriors obtained using HMC when fit to synthetic data are fairly consistent, but ABC-SMC can yield distributions with more variation in shape and location. For example, the peak location for the distributions of τ_r ,

τ_{si} , and τ_w^+ obtained using ABC-SMC vary substantially. For other parameters, such as τ_w^- and u_c^{si} , the results for HMC and ABC-SMC have similar peak locations, but much broader distributions for ABC-SMC. With experimental data, even greater variability across the ten distributions obtained using ABC-SMC is evident, with noticeable peak shifts for nearly all parameters. HMC produces more consistent results, indicating that this method samples the posterior appropriately and thus is less sensitive to randomization effects.

As a final measure of robustness, Fig. 11 shows the distribution of R^2 across the posterior sample for each of the ten runs using each method for the two models and the two datasets. In all scenarios, there is more variability in the pseudo- R^2 value across the posterior sample for ABC-SMC than for HMC, which consistently achieves narrow distributions. In the case of synthetic data and the ABC-SMC algorithm, the distributions are wider for the FK model fit than for the MS model; the FK model also has lower and more variable pseudo- R^2 values. HMC achieves pseudo- R^2 values closer to one for both models. In comparison with the results for synthetic data, the pseudo- R^2 values for the experimental data are lower and the distributions are wider for all cases, but with ABC-SMC distributions are wider than those of HMC

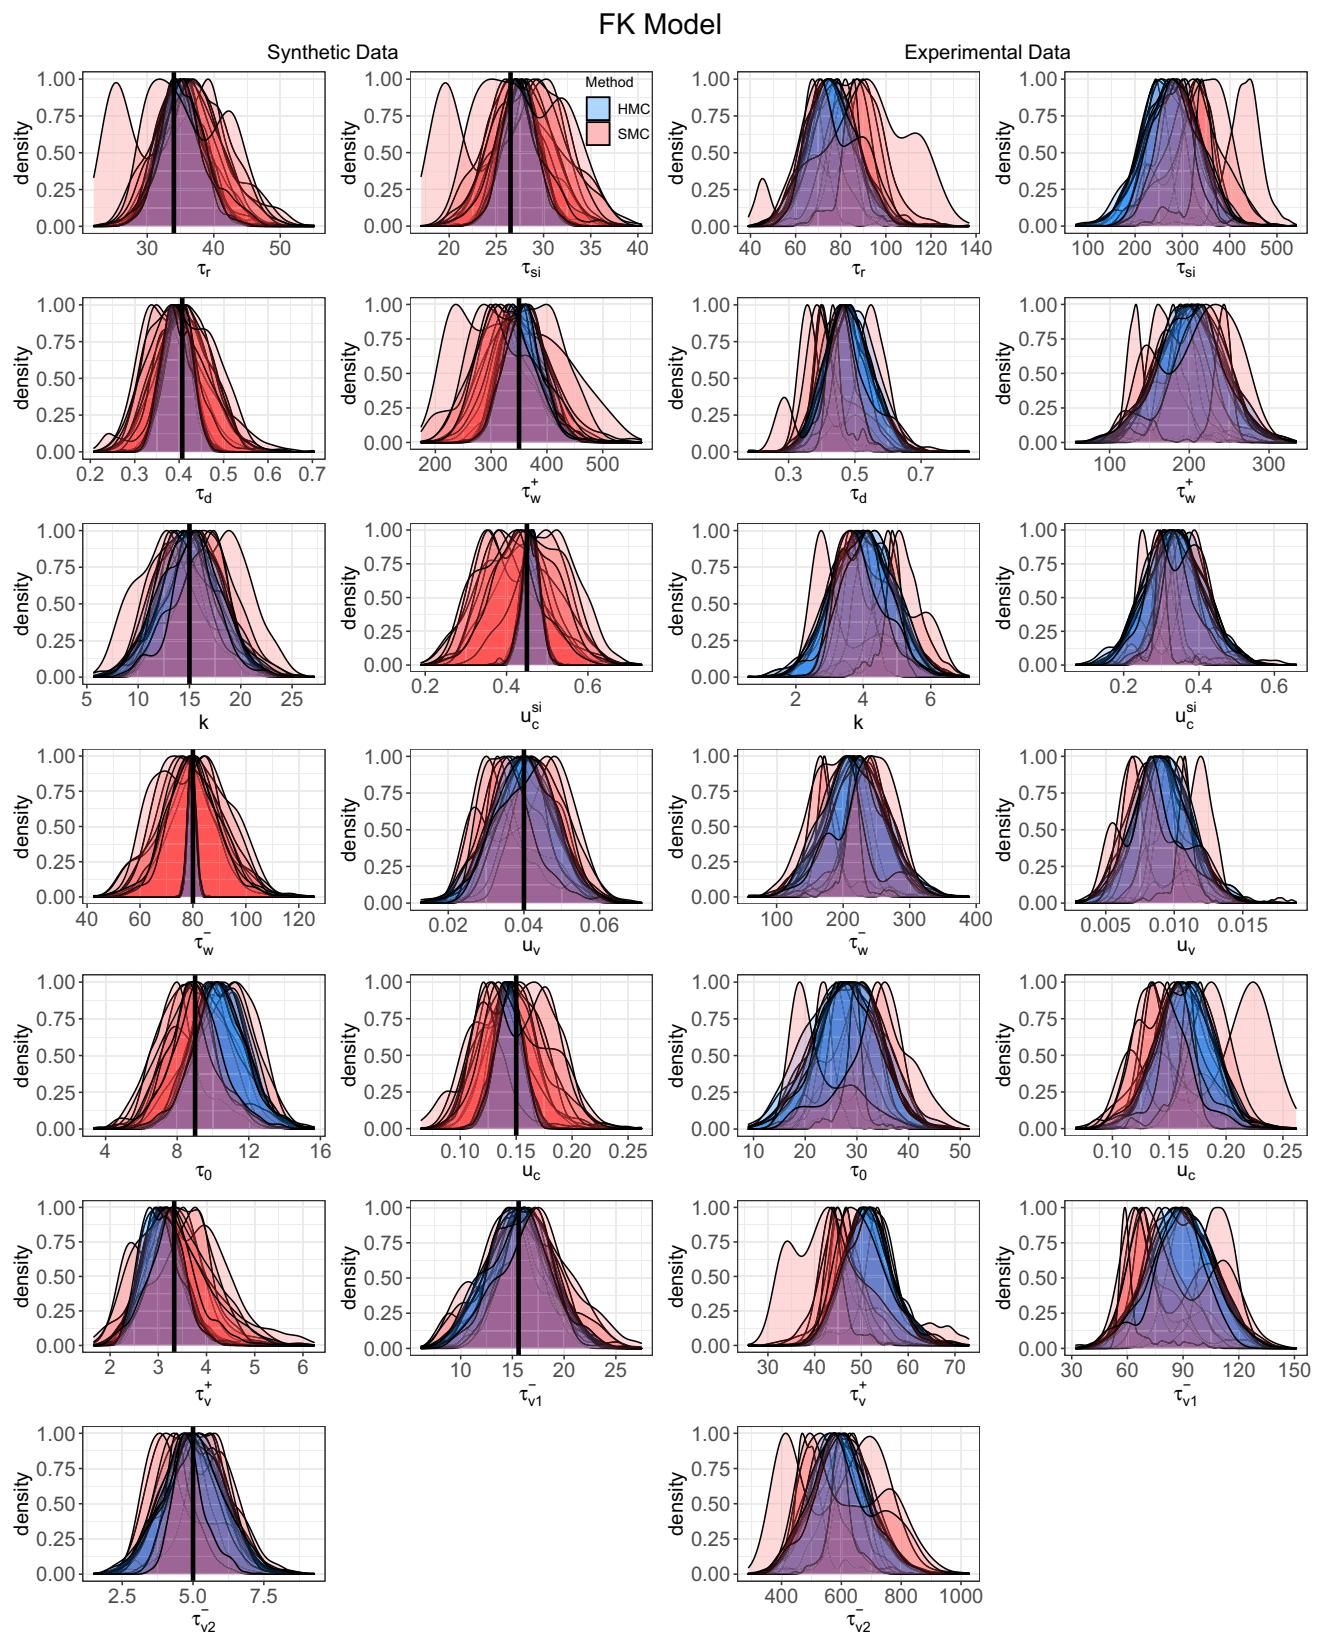


Fig. 10 Superimposed marginal distributions for the FK model parameters when fitting to (A) synthetic and (B) experimental data (HMC, blue; ABC-SMC, red) from 10 runs of each algorithm (HMC, blue; ABC-SMC, red). For synthetic data, the vertical lines represent the true values

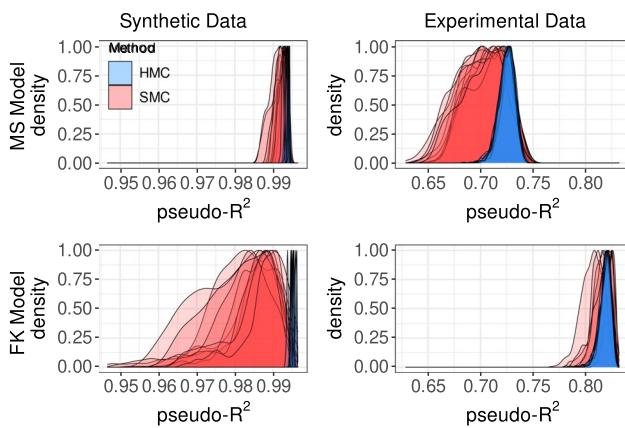


Fig. 11 Superimposed distributions of pseudo- R^2 from 10 runs of each algorithm (HMC, blue; ABC-SMC, red) for the MS (upper) and FK (lower) models and for synthetic (left) and experimental (right) data

for both models. In contrast to the synthetic data case, however, the values of pseudo- R^2 are closer to one for the FK model than for the MS model, a result that suggests, consistent with the results shown in Fig. 7 that the FK model is better able to fit the experimental data than the MS model.

3.4 Effect of dataset time resolution

The time resolution of the datasets used for fitting was downsampled from the available resolution. A physiological reason for the downsampling is that the different time scales of the action potential cause a uniform time resolution to deemphasize the upstroke [39], which is crucial to resolve for capturing behavior. The significant downsampling we perform thus results in a more balanced dataset by increasing the contribution of the upstroke. In addition, the experimental data recording was obtained using a resolution finer than is typical, and we wanted to ensure that good results could be obtained at more reasonable resolutions. Nevertheless, the downsampling eliminates information that potentially could improve the fittings obtained.

To study the effect of time resolution on the distributions and fittings obtained using HMC, we compared our results with those obtained using finer resolution with the MS model. With synthetic data, the resolution used was 0.4 ms for the first 4 ms and 5 ms for the remainder of each cycle (compared to 0.5 ms for the first 4 ms, then increased to 15 ms; see Section 2.2); the same priors and initial values were selected as before and are given in Table 1. Despite this increase in data resolution, which slightly increased computational time, little difference can be seen in either the

marginal distributions, shown in Fig. S11, or in the action potential fittings and APDs over a wide range of CLs, shown in Fig. S12. A decrease in variability can be seen in the bivariate scatterplots of parameter pairs in Fig. S13, which show essentially the same correlations as the lower resolution set but with greater clarity (and possibly a negative correlation between τ_{in} and v_{gate}).

When fitting the MS model to experimental data, increasing the dataset resolution again retained good fittings to the action potentials and similar distributions. In this case, the coarser resolution, as discussed in Section 2.2, was 0.4 ms for the first 4 ms, 1 ms for the next 4 ms, and 15 ms for the rest of each cycle; the finer resolution was the same except that after the first 8 ms the resolution was 5 ms. Priors and initial values were the same as those used earlier and are given in Table 1. Figure S14 shows the distributions, which are well-shaped and much narrower than the priors, very similar to the distributions in Fig. 6. Good fittings are obtained for the action potentials and APD across different CLs, as shown in Fig. S15. As with the synthetic data, the bivariate scatterplots, depicted in Fig. S16, show similar correlations but less variability; in this case, a possible new positive correlation between τ_{open} and v_{gate} appears.

3.5 Effect of prior selection for HMC

As mentioned in Section 2.3.1, HMC appeared to be sensitive to the choice of prior, a finding that motivated our use of ABC-SMC to develop an appropriate informative prior. Although using priors derived from ABC-SMC seems to help HMC obtain good results under a variety of conditions, these priors tend to be relatively narrow, and it would be ideal for HMC to work effectively with different types of priors. In this section, we show the effects of choosing different types of priors on the performance of HMC for several scenarios.

Figure S17 shows an example of HMC fitting the FK model to synthetic data for the broad uniform prior used successfully with ABC-SMC (see Table S2). For nearly all parameters, the marginal distributions are far from the true values. The action potentials obtained are misshapen, with the plateau missed and a bump during repolarization, as shown in Fig. S18. Stan provided a number of warnings to indicate poor performance: \hat{R} values were as large as 1.49, 495 divergent transitions occurred after the warm-up period (nearly the size of the posterior sample, which was 500), and the effective sample size was below 30 for all parameters, indicating that the posterior means and variances may be unreliable. When fitting experimental data, the true values are unknown. However, there is still ample evidence of unreliable distributions and poor fittings. Figures S19 and S20 show the results of HMC fitting the MS model to experimental data using a broad uniform prior (see

Table S1). Although less dramatic than the previous example, some of the distributions are poorly shaped and far from the distributions obtained with more informative priors (see Fig. 2), and the action potentials generally are too long. In addition, Stan provided similar metrics indicating unreliability, including \hat{R} values as large as 1.07, 439 divergent transitions after the warm-up period (again, nearly the size of the posterior sample), and the effective sample size was below 50 for all parameters, suggesting that the posterior means and variances may not be reliable.

However, it is possible to find different priors that are effective for HMC. For example, Fig. S21 shows marginal distributions obtained using gamma distributions (Gamma 1 parameters and initial values in Table S1). In this case, the only warning Stan gives is that the posterior variance may be unreliable, but the effective sample size is still almost always at least half the posterior sample size, and other metrics including \hat{R} are consistent with convergence. The action potentials from the posterior sample, shown in Fig. S22 are generally good fits to the data. Another example of good results obtained using gamma priors, this time fitting the experimental dataset using the FK model (parameters in Table S2), is shown in Figs. S23 and S24. The posterior distributions and action potentials are well-shaped, and Stan gives no warnings.

However, not all gamma distributions lead to good results. Figure S25 shows the marginal distributions that arise from another gamma prior derived from a fitting to the FK model obtained by hand (Gamma 2 parameters and initial values in Table S2). The distributions are extremely narrow and strangely shaped, although the action potentials, shown in Fig. S26, look fairly good. Nevertheless, Stan gives many warnings that the results may not be reliable, including \hat{R} values over 2, effective sample sizes below 10, and 500 problematic transitions. Overall, although HMC may not require priors to be as informative as those used for the results shown in Sections 3.1 and 3.2, it still tends to perform poorly with priors that are substantially different from the true posteriors.

4 Discussion

In this section, we compare results obtained for the different models, datasets, and Bayesian methods. We also describe some limitations of this study.

4.1 Influence of dataset and model choices

In all cases, the distributions found using HMC tended to be consistent, with more variability in the resulting

distributions occurring with ABC-SMC, which suggests that ABC-SMC may not always sample the posterior appropriately. For synthetic data, HMC often, but not always, gave rise to narrower distributions than ABC-SMC, whereas for experimental data, HMC's distributions sometimes were wider in such a way as to cover more of the variability of the ABC-SMC distributions.

With regard to models, we found clear differences in results for the two models considered. With synthetic data, for both methods, the peaks of the parameter distributions for the MS model usually coincided, with minimal shifting between the methods (despite variations in distribution widths). At the same time, the MS model showed a high incidence of correlations between pairs of parameters, especially with synthetic data (see Figs. S2 and S7). Given the model structure, this is not unexpected. The timescales of the main phases of the action potential are governed largely by particular time constants: τ_{in} during the upstroke, τ_{close} during the plateau, τ_{out} during early repolarization, and τ_{open} during late repolarization and rest. Many of the pairs of correlated parameters are time constants governing consecutive action potential phases. Because they scale inward and outward currents affecting the voltage, τ_{in} and τ_{out} also typically were correlated. The excitation threshold v_{gate} in the synthetic data case showed some correlation with the time constants associated with the transition from unexcited to excited states: τ_{open} and, to a lesser extent, τ_{in} . Overall, the MS model was able to obtain good fits in most cases, with less variability, but with more parameter correlations and some limitations in matching the action potential shapes from the zebrafish voltage recordings.

In contrast, for the FK model, greater variability arose, especially for the experimental data. We note that some of the parameters with greatest variability, such as τ_w^+ and u_c^{si} for synthetic data along with τ_w^+ , τ_{v1}^- , and τ_{v2}^- for experimental data, corresponded to parameters with more variability in a previous study fitting model parameters using a genetic algorithm [9]. Physiologically, τ_w^+ helps set the maximum APD, so the limited number of CLs used for fitting may not constrain this value as well as others. Similarly, τ_{v1}^- and τ_{v2}^- help to set the minimum diastolic interval and the slope of the APD restitution curve, respectively, and also may require additional data to be set properly, especially for the experimental setting where the model may not accurately describe the data and thus may not be able to fit all possible scenarios with low error throughout the action potentials. The increased variability and lack of observed bivariate correlations suggest that the FK model may be more difficult to fit, although its greater flexibility also extends its capability to produce different action potential properties like restitution compared to the MS model.

The experimental data scenario also showed fewer FK parameter correlations compared to the synthetic case; see Figs. S5 and S9. For the experimental case, only one strong correlation was visible, between parameters strongly affecting action potential duration (τ_r , which sets the strength of the slow outward current while the voltage is above threshold, and τ_v^+ , which essentially governs the duration of the fast inward current). With synthetic data, many additional correlations were observed, which most likely reflects the fact that the FK model is better able to fit data it generates, even in the presence of noise, than experimental data that are not described perfectly by the model. Some of these correlations, such as between the excitability τ_d and the excitation threshold u_c , as well as between the strength of the slow inward current governed by τ_{si} and the time course of the slow inward current's gating variable set by τ_w^+ , have been observed previously and visualized in the parameter fitness landscape [9].

4.2 Bayesian method considerations

Many aspects of our results across the two methods were fairly consistent. In particular, the q - q plots of the ABC-SMC vs. HMC samples do not support the idea that the posteriors obtained with the two algorithms come from different families of distributions. Nevertheless, each of the two methods considered has advantages and disadvantages. HMC performs exact inference and explores parameter space more efficiently than traditional methods like Metropolis-Hastings. Furthermore, HMC was consistent when running the programs multiple times, giving posteriors centered around the same value, for synthetic or experimental data. However, it can be difficult to use HMC because of the need to choose initial points and a prior distribution. Inappropriate selections for either initial points or the prior can affect convergence and lead to unreliable results, but finding good choices can be time-consuming.

In contrast, although we show results only for the same priors used for HMC, ABC-SMC could produce good results with a variety of different priors, including uniform, gamma, and folded normal distributions. Even with relatively non-informative priors such as wide uniforms, ABC-SMC was able to find a useful approximation to the posterior in all cases we tried, for synthetic or experimental data. However, the lower bound for the pseudo- R^2 distribution when fitting experimental data was smaller using wide uniform priors than folded normal distributions (MS: not lower than 0.48 vs. 0.65; FK: not lower than 0.67 vs. 0.79). While we expect that this result occurs because the uniforms are less informative than the folded normal distributions, more study would be needed to make a fair comparison. In addition to imposing less stringent requirements for the priors, we have found that ABC-SMC could find an approximate population

fit even if the population size was small (e.g., 100; results not shown). Furthermore, we demonstrated that HMC performed well and consistently for different temporal resolutions of the dataset, although higher resolutions require more computational time.

We also note that the lower consistency of the ABC-SMC results might suggest that the ABC-SMC results could be improved with a larger sample size. However, when we have tested output from larger samples, consistency was not meaningfully improved (results not shown). Instead, we believe that the variability is more inherent to the difference in how the ABC-SMC method works: it is an approximation and may not always sample the posterior properly. One possible way to achieve greater consistency may be by adjusting and tightening the ABC-SMC tolerances, but it is difficult to optimize the tolerance schedule; choosing tolerances that are too small can lead to nearly all candidate solutions being rejected, which considerably slows the algorithm and in some cases may compromise convergence.

In terms of computational efficiency of the two methods for our main findings in Sections 3.1 and 3.2, using experimental data with ABC-SMC generally took around 5 min and more than twice as long for synthetic data when fitting the MS model. For the FK model, ABC-SMC took only slightly longer to fit experimental data, around 6 min, and about 1.5 times longer than that to fit synthetic data. For HMC, the experimental data could be fit to the MS model much more quickly, in about 3.5 min, with the fitting to synthetic data taking about eight times longer. However, when fitting synthetic data, HMC took around 12 min to fit the MS model but around 4 h to fit the FK model. A possible explanation for the long times required for HMC to fit synthetic data may be that the likelihood is very flat for some regions of parameter space, limiting choices for acceptable candidate parameterizations. Another possible reason that HMC is much faster for fitting experimental data may be the nature and extent of noise and variability in the dataset. Indeed, HMC is unable to obtain appropriate fittings for datasets with no or very low noise, which, along with the magnitude of noise in the experimental data, influenced our selection of the noise level in creating the synthetic datasets considered.

Overall, we found that ABC-SMC was able to obtain useful approximations to the posterior under a broader range of conditions, whereas HMC imposed more constraints for reasonable performance. However, even ABC-SMC saw benefit from the use of an informative prior. Because ABC-SMC does not need initial points and accepts wide priors from several different types of distributions, it can be used to find feasible priors for HMC, and the initial points for HMC can be selected as the modes of the distributions obtained with the first ABC-SMC pass. It is true that comparing HMC results using a prior derived from ABC-SMC to ABC-SMC

using the same prior philosophically seems like using the same data twice for the latter case; however, we can think of ABC-SMC in that case as getting some help from the maximum likelihood estimate, and in practice we likely would only use such an approach for ABC-SMC when comparing directly with another method like HMC. Our hybrid approach, which was used to obtain the results shown here, can be useful to sidestep the difficulties of working with HMC, especially its sensitivity to the selection of priors, while taking advantage of its ability to generate a population that closely fits the data with limited variability and of its performance of exact inference as a full Bayesian method.

4.3 Limitations

In this study, we considered only a limited number of datasets. In particular, we chose data from three CLs, with one at a longer CL and two at shorter CLs within the alternans regime. It may be possible to optimize the selection of CLs beyond what was chosen. In addition, it is possible that performance may change for noisier data or for data with different dynamics that may not be well captured by the model being fit.

Similarly, we only considered two models, and it is possible that performance could differ for a different selection of models. We also note that we have used a single approximate Bayesian method and that a different choice may result in different findings. In addition, we chose to restrict parameter values to be positive using priors with positive support (folded normals), but parameter transforms [40] could be a useful alternative. We also considered only bivariate correlations across parameters, although higher-dimensional correlations may be present. Different algorithmic choices, such as the use of a different distance function, may be worth considering, although some preliminary trials with an absolute value-based distance did not appreciably change the results.

Within the models, we used a simple square-pulse stimulus. Use of a biphasic stimulus [9] could help prevent selection of large values for the excitability parameters (e.g., τ_{in} for the MS model and τ_d for the FK model) that would not allow propagation in tissue and thus may be unphysiological.

The tolerance reduction approach used for ABC-SMC, while adaptive, nevertheless was fixed in advance, following ref. [22]. It would be interesting to try a more sophisticated way to select the tolerances to improve efficiency and to facilitate working with different datasets.

We also note that to make the comparison between ABC-SMC and HMC fair, we chose to use one chain for HMC in Stan, but we found the results were consistent when using the default number of chains, which was four. In addition, other Bayesian methods could be considered, such as Metropolis-Hastings with an adaptive covariance scheme to help the multivariate normal proposal distribution evolve towards the covariance of the accepted samples in the chain so far [41].

5 Conclusion

In this manuscript, we have used two Bayesian methods, HMC and ABC-SMC, to find populations of cardiac action potential model parameters consistent with data used for fitting. We have shown that the methods can work effectively with both synthetic data derived from the models used and for an experimental dataset taken from a zebrafish heart. In nearly all cases, both methods find well-shaped marginal distributions with clear peaks for each model parameter for both the MS model, which has five parameters, and the FK model, which has 13. We also have shown through the use of q - q plots that the posterior distribution samples obtained by the two methods do not give any strong indication of being from different distribution families; in other words, both methods appear to converge to the same type of distribution. In the case of synthetic data, where the true parameters used to generate the dataset are known, those true values in general are well contained within the posterior distributions found, and across multiple runs of the algorithms the true values coincide well with the distribution peaks.

Given that both methods achieve similar results with no clear computational advantage, other considerations may motivate the choice of method. ABC-SMC is generally easier to work with, as it accepts different kinds of prior distributions, and those distributions may be broad, and it often finds useful approximations of the true posteriors. While HMC requires more effort to find an acceptable prior (and indeed, we suggest that ABC-SMC may be useful in this task), it performs exact inference, such that when it converges, it finds the true posterior.

In the future, it may be useful to optimize the data used for fitting to better constrain certain parameter values. For example, in the FK model, the parameter τ_{v1}^- helps to set the minimum diastolic interval; datasets that do not represent that information may have difficulty adequately constraining that parameter and related parameters like u_v . We also expect Bayesian methods such as these will be useful for ongoing efforts including efficient creation of model populations [42] and virtual patient cohorts [4] as well as addressing nonidentifiability of model parameters [15, 43–45] and uncertainty quantification [14, 46].

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1007/s11517-022-02685-y>.

Author contribution All authors contributed to the study conception and design. Data collection and analysis were performed by Alejandro Nieto Ramos. The first draft of the manuscript was written by Alejandro Nieto Ramos and all authors contributed to and commented on subsequent versions of the manuscript. All authors read and approved the final manuscript.

Funding This study was supported by NSF grants CNS-2028677 and CMMI-1762553 and by NIH grant 1R01HL143450.

Availability of data and material The data used for this study are available upon request.

Code availability The code used for this study is available upon request.

Declarations

Conflict of interest The authors declare no competing interests.

References

1. Fenton FH, Cherry EM (2008) Models of cardiac cell. Scholarpedia 3(8):1868
2. Groenendaal W, Ortega FA, Kherlopian AR, Zygmunt AC, Krogh-Madsen T, Christini DJ (2015) Cell-specific cardiac electrophysiology models. *PLoS Comput Biol* 11(4):1004242. <https://doi.org/10.1371/journal.pcbi.1004242>
3. Boyle PM, Zghaib T, Zahid S, Ali RL, Deng D, Franceschi WH, Hakim JB, Murphy MJ, Prakosa A, Zimmerman SL, Ashikaga H, Marine JE, Kolandaivelu A, Nazarian S, Spragg DD, Calkins H, Trayanova NA (2019) Computationally guided personalized targeted ablation of persistent atrial fibrillation. *Nat Biomed Eng* 3(11):870–879
4. Niederer SA, Aboelkassem Y, Cantwell CD, Corrado C, Coveney S, Cherry EM, Delhaas T, Fenton FH, Panfilov AV, Pathmanathan P, Plank G, Riabiz M, Roney CH, dos Santos RW, Wang L (2020) Creation and application of virtual patient cohorts of heart models. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 378(2173):20190558. <https://doi.org/10.1098/rsta.2019.0558>
5. Dokos S, Lovell NH (2004) Parameter estimation in cardiac ionic models. *Prog Biophys Mol Biol* 85(2):407–431. <https://doi.org/10.1016/j.biophys.2004.02.002>
6. Bueno-Orovio A, Cherry EM, Fenton FH (2008) Minimal model for human ventricular action potentials in tissue. *J Theor Biol* 253(3):544–560. <https://doi.org/10.1016/j.jtbi.2008.03.029>
7. Syed Z, Vigmond E, Nattel S, Leon LJ (2005) Atrial cell action potential parameter fitting using genetic algorithms. *Med Biol Eng Compu* 43(5):561–571
8. Bot CT, Kherlopian AR, Ortega FA, Christini DJ, Krogh-Madsen T (2012) Rapid genetic algorithm optimization of a mouse computational model: benefits for anthropomorphization of neonatal mouse cardiomyocytes. *Front Physiol* 3:421. <https://doi.org/10.3389/fphys.2012.00421>
9. Cairns DI, Fenton FH, Cherry EM (2017) Efficient parameterization of cardiac action potential models using a genetic algorithm. *Chaos* 27(9):093922. <https://doi.org/10.1063/1.5000354>
10. Loewe A, Wilhelms M, Schmid J, Krause MJ, Fischer F, Thomas D, Scholz EP, Dössel O, Seemann G (2015) Parameter estimation of ion current formulations requires hybrid optimization approach to be both accurate and reliable. *Front Bioeng Biotechnol* 3:209. <https://doi.org/10.3389/fbioe.2015.00209>
11. Coveney S, Clayton RH (2018) Fitting two human atrial cell models to experimental data using Bayesian history matching. *Prog Biophys Mol Biol* 139:43–58. <https://doi.org/10.1016/j.biophys.2018.08.001>
12. Zaman MS, Dhamala J, Bajracharya P, Sapp JL, Horácek BM, Wu KC, Trayanova NA, Wang L (2021) Fast posterior estimation of cardiac electrophysiological model parameters via Bayesian active learning. *Frontiers in Physiology* 12: 740306. <https://doi.org/10.3389/fphys.2021.740306>. Accessed 2022-04-15
13. Siekmann I, Wagner LE, Yule D, Fox C, Bryant D, Crampin EJ, Sneyd J (2011) MCMC estimation of Markov models for ion channels. *Bioophys J* 100(8):1919–1929. <https://doi.org/10.1016/j.bpj.2011.02.059>
14. Pathmanathan P, Shotwell MS, Gavaghan DJ, Cordeiro JM, Gray RA (2015) Uncertainty quantification of fast sodium current steady-state inactivation for multi-scale models of cardiac electrophysiology. *Prog Biophys Mol Biol* 117(1):4–18. <https://doi.org/10.1016/j.biophys.2015.01.008>
15. Daly AC, Gavaghan DJ, Holmes C, Cooper J (2015) Hodgkin-Huxley revisited: reparametrization and identifiability analysis of the classic action potential model with approximate Bayesian methods. *Royal Society Open Science* 2(12):150499. <https://doi.org/10.1098/rsos.150499>
16. Neal, R.M.: MCMC using Hamiltonian dynamics. In: *Handbook of Markov chain Monte Carlo*. Chapman & Hall/CRC Handb. Mod. Stat. Methods, pp. 113–162 (2011)
17. Vernon I, Liu J, Goldstein M, Rowe J, Topping J, Lindsey K (2018) Bayesian uncertainty analysis for complex systems biology models: emulation, global parameter searches and evaluation of gene functions. *BMC Syst Biol* 12(1):1. <https://doi.org/10.1186/s12918-017-0484-3>
18. Del Moral P, Doucet A, Jasra A (2006) Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 68(3): 411–436
19. Toni T, Welch D, Strelkowa N, Ipsen A, Stumpf MPH (2009) Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *J R Soc Interface* 6(31):187–202. <https://doi.org/10.1098/rsif.2008.0172>
20. Del Moral P, Doucet A, Jasra A (2012) An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Stat Comput* 22(5):1009–1020. <https://doi.org/10.1007/s11222-011-9271-y>
21. O'Hara T, Virág L, Varró A, Rudy Y (2011) Simulation of the undiseased human cardiac ventricular action potential: model formulation and experimental validation. *PLoS Comput Biol* 7(5):1002061
22. Daly AC, Cooper J, Gavaghan DJ, Holmes C (2017) Comparing two sequential Monte Carlo samplers for exact and approximate Bayesian inference on biological models. *J R Soc Interface* 14(134):20170340. <https://doi.org/10.1098/rsif.2017.0340>
23. Duane S, Kennedy AD, Pendleton BJ, Roweth D (1987) Hybrid Monte Carlo. *Phys Lett B* 195(2):216–222. [https://doi.org/10.1016/0370-2693\(87\)91197-X](https://doi.org/10.1016/0370-2693(87)91197-X)
24. Monnahan CC, Thorson JT, Branch TA (2017) Faster estimation of Bayesian models in ecology using Hamiltonian Monte Carlo. *Methods Ecol Evol* 8(3):339–348. <https://doi.org/10.1111/2041-210X.12681>
25. Margossian CC, Zhang Y, Gillespie WR (2021) Flexible and efficient Bayesian pharmacometrics modeling using Stan and Torsten, part I. [arXiv:2109.10184](https://arxiv.org/abs/2109.10184) [stat]
26. Nieto Ramos A, Herndon CJ, Fenton FH, Cherry EM (2021) Quantifying distributions of parameters for cardiac action potential models using the Hamiltonian Monte Carlo method. *Computing in Cardiology* 48:9662836–196628364
27. Mitchell CC, Schaeffer DG (2003) A two-current model for the dynamics of cardiac membrane. *Bull Math Biol* 65(5):767–793. [https://doi.org/10.1016/S0092-8240\(03\)00041-7](https://doi.org/10.1016/S0092-8240(03)00041-7)
28. Fenton F, Karma A (1998) Vortex dynamics in three-dimensional continuous myocardium with fiber rotation: filament instability and fibrillation. *Chaos* 8(1):20–47. <https://doi.org/10.1063/1.166311>
29. Fenton FH, Cherry EM, Hastings HM, Evans SJ (2002) Multiple mechanisms of spiral wave breakup in a model of cardiac electrical activity. *Chaos* 12(3):852–892. <https://doi.org/10.1063/1.1504242>
30. Gamerman D, Lopes HF (2006) *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*, Second Edition
31. Marin J-M, Pudlo P, Robert CP, Ryder RJ (2012) Approximate Bayesian computational methods. *Stat Comput* 22(6):1167–1180. <https://doi.org/10.1007/s11222-011-9288-2>

32. Betancourt M (2018) A conceptual introduction to Hamiltonian Monte Carlo. [arXiv:1701.02434](https://arxiv.org/abs/1701.02434) [stat]
33. Hoffman MD, Gelman A (2014) The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J Mach Learn Res* 15(1):1593–1623
34. Carpenter B, Gelman A, Hoffman MD, Lee D, Goodrich B, Betancourt M, Brubaker M, Guo J, Li P, Riddell A (2017) Stan: a probabilistic programming language. *J Stat Softw* 76:1–32. <https://doi.org/10.18637/jss.v076.i01>
35. Stan Development Team (2022) Stan modeling language users guide and reference manual, version 2.29. [//mc-stan.org/](http://mc-stan.org/)
36. Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB (2013) Bayesian data analysis, 3rd edition edn. Chapman and Hall/CRC
37. Vehtari A, Gelman A, Simpson D, Carpenter B, Bürkner P-C (2021) Rank-normalization, folding, and localization: an improved \hat{R} for assessing convergence of MCMC. *Bayesian Analysis* 16(2). <https://doi.org/10.1214/20-BA1221>
38. Chicco D, Warrens MJ, Jurman G (2021) The coefficient of determination R-squared is more informative than SMAPE, MAE, MAPE, MSE and RMSE in regression analysis evaluation. *PeerJ Computer Science* 7:623. <https://doi.org/10.7717/peerj.cs.623>
39. Shahi S, Marcotte CD, Herndon CJ, Fenton FH, Shiferaw Y, Cherry EM (2021) Long-time prediction of arrhythmic cardiac action potentials using recurrent neural networks and reservoir computing. *Frontiers in physiology* 12
40. Whittaker DG, Clerx M, Lei CL, Christini DJ, Mirams GR (2020) Calibration of ionic and cellular cardiac electrophysiology models. *Wiley Interdisciplinary Reviews: Systems Biology and Medicine* 12(4):1482
41. Johnstone RH, Chang ET, Bardenet R, De Boer TP, Gavaghan DJ, Pathmanathan P, Clayton RH, Mirams GR (2016) Uncertainty and variability in models of the cardiac action potential: can we build trustworthy models? *J Mol Cell Cardiol* 96:49–62
42. Britton OJ, Bueno-Orovio A, Van Ammel K, Lu HR, Towart R, Gallacher DJ, Rodriguez B (2013) Experimentally calibrated population of models predicts and explains intersubject variability in cardiac cellular electrophysiology. *Proc Natl Acad Sci USA* 110(23):2098–2105. <https://doi.org/10.1073/pnas.1304382110>
43. Csercsik D, Hangos KM, Szederkényi G (2012) Identifiability analysis and parameter estimation of a single Hodgkin-Huxley type voltage dependent ion channel under voltage step measurement conditions. *Neurocomputing* 77(1):178–188. <https://doi.org/10.1016/j.neucom.2011.09.006>
44. Shotwell MS, Gray RA (2016) Estimability analysis and optimal design in dynamic multi-scale models of cardiac electrophysiology. *J Agric Biol Environ Stat* 21(2):261–276. <https://doi.org/10.1007/s13253-016-0244-7>
45. Daly AC, Gavaghan D, Cooper J, Tavener S (2018) Inference-based assessment of parameter identifiability in nonlinear biological models. *J R Soc Interface* 15(144):20180318. <https://doi.org/10.1098/rsif.2018.0318>
46. Chang KC, Dutta S, Mirams GR, Beattie KA, Sheng J, Tran PN, Wu M, Wu WW, Colatsky T, Strauss DG, Li Z (2017) Uncertainty quantification reveals the importance of data variability and experimental design considerations for in silico proarrhythmia risk assessment. *Front Physiol* 8:917. <https://doi.org/10.3389/fphys.2017.00917>

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.



Alejandro Nieto Ramos received his bachelor's and master's degrees in applied mathematics in Mexico City. He recently received his PhD from the Rochester Institute of Technology and is currently a postdoctoral researcher at the Cleveland Clinic. His research has included analysis of a work by Palestrina using the mathematical and musical counterpoint of the first species, application of a Bayesian approach to fit experimental data of tumor growth in mice with compartmental pharmacokinetics models, and Bayesian inference and prediction in cardiac electrophysiology models with an application to representing

variability.



Flavio H. Fenton is a Professor in the School of Physics at the Georgia Institute of Technology. He received a BS degree in theoretical physics from the Universidad Nacional Autónoma de México, and MS and PhD degrees in physics from Northeastern University. His research is centered on excitable media, complex systems, and pattern formation using an integrated approach combining theory, experiments, and high-performance computer simulations.



Elizabeth M. Cherry is an Associate Professor in the School of Computational Science and Engineering at the Georgia Institute of Technology. Her research involves modeling and simulation, high-performance computing, and computational methods, including techniques like data assimilation, Bayesian estimation, and machine learning, applied primarily to cardiac arrhythmias. She received a BS in Mathematics and American Studies from Georgetown University and a PhD in Computer Science from Duke University focusing on efficient computational methods for solving partial-differential-equations models of electrical signals in the heart.