Gaussian process emulation for discontinuous response surfaces with applications for cardiac electrophysiology models

Sanmitra Ghosh[1,*], David J. Gavaghan[1], Gary R. Mirams[2]

1 Computational Biology and Health Informatics, Department of Computer Science, University of Oxford, UK.
2 Centre for Mathematical Medicine & Biology, School of Mathematical Sciences, University of Nottingham, UK.

* sanmitra.ghosh@cs.ox.ac.uk

Abstract

Mathematical models of biological systems are beginning to be used for safety-critical applications, where large numbers of repeated model evaluations are required to perform uncertainty quantification and sensitivity analysis. Most of these models are nonlinear both in variables and parameters/inputs which has two consequences. First, analytic solutions are rarely available so repeated evaluation of these models by numerically solving differential equations incurs a significant computational burden. Second, many models undergo bifurcations in behaviour as parameters are varied. As a result, simulation outputs often contain discontinuities as we change parameter values and move through parameter/input space.

Statistical emulators such as Gaussian processes are frequently used to reduce the computational cost of uncertainty quantification, but discontinuities render a standard Gaussian process emulation approach unsuitable as these emulators assume a smooth and continuous response to changes in parameter values.

In this article, we propose a novel two-step method for building a Gaussian Process emulator for models with discontinuous response surfaces. We first use a Gaussian Process classifier to detect boundaries of discontinuities and then constrain the Gaussian Process emulation of the response surface within these boundaries. We introduce a novel 'certainty metric' to guide active learning for a multi-class probabilistic classifier.

We apply the new classifier to simulations of drug action on a cardiac electrophysiology model, to propagate our uncertainty in a drug's action through

to predictions of changes to the cardiac action potential. The proposed two-step active learning method significantly reduces the computational cost of emulating models that undergo multiple bifurcations.

1 Introduction

Mathematical models are used ubiquitously to develop a mechanistic understanding of complex biological systems. However, the efficacy of these models in safety-critical applications depends on their ability to capture the interactions of several physical variables in detail in order to reproduce biological phenomena accurately [1]. These models are often defined as complex nonlinear dynamical systems of parameterised equations that can become intensive to computationally simulate. Tasks such as uncertainty quantification and sensitivity analysis that require repeated evaluation with different parameter sets thus become computationally burdensome. A computationally cheaper alternative, an *emulator*, that gives a close approximation to the responses (output) of these models is thus extremely useful for the above mentioned tasks (we refer the reader to previous work [2,3] for a detailed introduction to this topic).

Mathematical models are beginning to be used in pre-clinical drug safety and toxicology studies to learn about a compound's action on electrophysiology and associated risk [4–6]. The underlying mathematical models are dynamical systems described as coupled nonlinear ordinary or partial differential equations (ODEs/PDEs) that depict intra- or inter-cellular ionic exchanges and the state of cell electrophysiological components. Such simulations predict the occurrence of changes to the cellular *action potential* (time trace for the cell's transmembrane voltage). Often drug-induced changes are summarised by their influence on action potential biomarkers. One such commonly used marker is the action potential duration (APD) which quantifies the time lag between depolarisation and repolarisation of membrane voltage. Commonly the output of a simulation is summarised by one or more such biomarkers that can be compared with experimental recordings.

Drug effects can be modelled by scaling the conductance parameter of multiple ion channels [7]. The degree of scaling depends on compound concentration, and is deduced from High Throughput Screening (HTS) of multiple ion channels, which is subject to considerable variability that should be taken into account as it has large effects on predictions of drug-induced changes to whole-cell electrophysiology [8,9].

This propagation of uncertainty, from experimental assay results (that form simulation inputs) through to simulation results, is computationally expensive, as repeated simulations involving numerical solution of differential equations have to be carried out for various input parameters. Statistical emulators can be built that model the response surfaces spanned by the simulation outputs. Such an emulator can be trained using a small number of input-output pairs of the simulator; the input being the scalings applied to conductances and output being the action potential biomarkers such as APD. Once trained, the emulator can

then be used as a computationally cheaper alternative to predict the simulation output for a large number of drug blocks. In previous work [9] a simple emulator based on linear interpolation from a multi-dimensional look-up table was used to speed up the uncertainty quantification analysis we had previously performed using a 'brute force' Monte Carlo method in our earliest work on this topic [8].

A more efficient emulator was proposed more recently [10] (in terms of the number of training points required for a given accuracy) for various biomarkers obtained from simulated action potential time courses of the Luo-Rudy cardiac action potential model [11]. This emulator used Gaussian Processes (GPs) to statistically model the output response surfaces of the biomarkers. Despite their computationally attractive properties, designing an emulator for cardiac electrophysiology models is extremely challenging since many of these models undergo bifurcations resulting in discontinuous response surfaces, as we show in Fig. 1 for the widely-used O'Hara (or 'ORd' model) for human ventricular action potentials [12].

In this paper we will present an emulator of APD at 90% repolarisation, APD_{90}, in the O'Hara model, designed to work in spite of the discontinuities in the response surface. Our proposed emulator consists of a two-step method in which we use a boundary detector to segment the response surface along the discontinuities and then apply statistical regression to emulate the responses. We formulate the boundary detection as a classification problem. In both these steps we use Gaussian Processes. The proposed emulator builds upon the work in [10] and [13] with improvements to deal with these bifurcations in model behaviour. Fig. 2 shows the main steps in both simulation and emulation of cardiac biomarkers.

1.1 Concentration-effect curves

Our ODE system for cardiac action potentials has certain parameters that we can modify to simulate drug action. Typically these are the maximal current densities, or maximal conductances, of certain ion currents, denoted G, for various ion currents e.g. G_{Kr}, G_{CaL}, G_{Na} etc.. Simulating the action of drug block typically means scaling a G parameter by multiplying by a factor R which is in the range $[0, 1]$, as described below. In what follows, the vector $\mathbf{R}$ (scaling factors R for each channel j) then defines our parameter set or point in simulator input space.

A concentration-effect curve maps the concentration of a compound to an effect or response. The percentage of the peak ionic current, following a voltage step, is repeatedly recorded and the proportion that remains is the recorded effect or response $R([C])$. Usually such curves are described by a Hill function:

$$R([C]) = 1 - \frac{[C]^n}{[C]^n + [IC50]^n}. \tag{1}$$

This function of concentration $[C]$, has two parameters: $[IC50]$, the half-maximal inhibitory concentration; and the Hill coefficient n. These parameters are

estimated by fitting the concentration-effect curve to screening results. The effect of the conductance block by a specific compound is then studied (in simulation) using a cardiac action potential (AP) model by scaling the maximal channel conductance G of a particular channel j:

$$G_j = G_{j,control} R_j \left([C] \right), \tag{2}$$

where $R_j \left([C] \right) \in [0, 1]$ is the conductance scaling given by the concentration-effect curve (equation (1)) for ion channel j, and $G_{j,control}$ is the maximal conductance of that channel in control (drug free) conditions. This conductance scaling R_j is related to the degree of ion channel block as

$$\% \ block = 100 \times \left(1 - R_j \left([C] \right) \right). \tag{3}$$

1.2 Handling discontinuities

In Chang *et al.* [10] the range of conductances used as input was chosen around the maximal conductances of the Luo-Rudy model in drug free (control) condition. Extending this range, using equation (2), introduces other effects such as absence of depolarisation or repolarisation of the membrane voltage. This is caused due to the dynamical system going through bifurcations as the conductance parameters are set to values beyond a limited range around the maximal points. Examples of such voltage time courses are shown in Fig. 1 for the O'Hara model [12].

For uncertainty propagation, the effect of drug block on APs [9] may span the entire domain of conductance scalings $R_j \in [0, 1]$, where j denotes any specific channel, which are applied following equation (2). It is evident from Fig. 1 that such a task poses a severe challenge to any emulator, as it needs to learn the location of the discontinuity of the emulated surface, as well as the value of the surface, from limited evaluations of an underlying model simulator.

Applying GPs to model discontinuous functions is largely an open problem. Although many advances (see the discussion about non-stationarity in [14] and the references in there) have been made towards solving this problem, a common solution has not yet emerged. In the recent GP literature there are two specific streams of work that have been proposed for modelling non-stationary response surfaces including those with discontinuities. The first approach is based on designing non-stationary processes [15] whereas the other approach attempts to divide the input space into separate regions and build separate GP models for each of the segmented regions. Such input domain segmentation algorithms use a tree based GP model [16,17]. In such a GP model the individual nodes (leaves) of the tree are built using a smaller subset of the inputs. Furthermore, the model is constrained in such a way that inputs between discontinuous regions are not shared among the nodes. Our work is motivated by the latter approach of space partitioning which we turn to next.

We want the emulator to return the APD_{90} for a valid AP only (see Fig 1). Although our emulator is general purpose and can be used with any summary statistic we concentrated on the APD_{90} because its significance in drug induced cardiac toxicity studies – the application for this emulator. Using a two-step

emulator we can use a classifier to label any queried input into one of three categories according to the generated AP trace: 1) no-depolarisation, 2) a valid action potential (AP) and 3) no-repolaristaion. If the input falls under the second category then we can use GP regression to predict the corresponding output APD_{90} value, and associated emulator uncertainty [18]. Augmenting the response surface prediction step with a boundary detection we can circumvent the discontinuity while maintaining all the benefits of emulation as performed by Chang *et al.* [10]. Using Gaussian processes for both classification and surface prediction enables us to use the uncertainty associated with the prediction to carry out sequential design of the input space. Thus we use 'active learning' to choose training data for building the emulator which can further reduce the need for expensive space-filling designs. For the surface regression GP we use conditional entropy [19] as a measure of uncertainty, and for the classifier GP we propose a novel metric to measure the uncertainty associated with its prediction step.

Using a two-step emulator and carrying out sequential design for both these steps enables us to use the proposed emulator as a viable alternative at a fraction of the computational expense of solving ODEs numerically for Monte Carlo samples.

In the following section we will briefly review the fundamentals of GP regression and classification. We will then proceed to explain how we use GP regression and classification to build an emulator of the APD_{90} response surface.

1.3 Brief review of Gaussian processes for regression and classification

In the following sections we will review Gaussian processes for regression and classification. We will also mention briefly the approximation methods required to apply GP to larger datasets.

1.3.1 Gaussian processes regression

Consider the regression problem where we have the dataset $\mathcal{D} = \{(\boldsymbol{X}_i, y_i), i = 1, \ldots, n\}$ consisting of input $\boldsymbol{X}_i \in \mathbb{R}^{\mathbb{D}}$ and noisy scalar observations y_i. In the simplest case we assume that the noise is independent and Gaussian such that the *latent function* $f : \mathbb{R}^{\mathbb{D}} \to \mathbb{R}$ and the (possibly) noisy observations are related, using the notation in [20], as

$$y_i = f(\boldsymbol{X}_i) + \epsilon_i, \tag{4}$$

where $\epsilon_i \sim \mathcal{N}(0, \sigma_{noise}^2)$. In our application the 'latent function' is the simulator output that we wish to emulate. Also, in the context of our application we consider the noise term ϵ to be very small indeed as we are emulating the output of a deterministic ODE system, but it could represent inaccuracies introduced by numerical solution of the cardiac model. In a probabilistic framework we are interested in the probability distribution of function values f_* (or the noisy y_*)

5

at test locations (where we have not run the simulator but wish to infer values from the emulator) that we call $\boldsymbol{X}_*$.

Definition 1. *A Gaussian process is a collection of random variables, any finite number of which have consistent joint Gaussian distributions*

Gaussian process regression is a Bayesian approach where we place a prior over functions. For the regression problem we assume a priori that the function values behave according to

$$p(\boldsymbol{f}|\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_n) = \mathcal{N}(0, K), \tag{5}$$

where $\boldsymbol{f} = (f_1, \ldots, f_n)^T$ is a vector of latent function values $f_i = f(\boldsymbol{X}_i)$ and K is a covariance matrix whose entries are given by the covariance function $K_{i,j} = k(\boldsymbol{X}_i, \boldsymbol{X}_j; \boldsymbol{\phi})$, where $\boldsymbol{\phi}$ is a vector of hyperparameters for the covariance function. A common example of a covariance function is the squared exponential function:

$$k(\boldsymbol{X}_i, \boldsymbol{X}_j) = \nu^2 \exp\left(\frac{\|\boldsymbol{X}_i - \boldsymbol{X}_j\|^2}{l^2}\right), \tag{6}$$

with the hyperparameters $\boldsymbol{\phi} = \{\nu, l\}$, where ν is the prior variance and l is a lengthscale parameter that defines the decay rate in space of the covariance between points on the response surface.

Inference of latent function values for test locations is carried out by first placing a joint prior on the training and test latent function values $\boldsymbol{f}$ and $\boldsymbol{f}_*$

$$\begin{bmatrix} \boldsymbol{f} \\ \boldsymbol{f}_* \end{bmatrix} \sim \mathcal{N}\left(\boldsymbol{0}, \begin{bmatrix} \mathbf{K_{f,f}} & \mathbf{K_{f,*}} \\ \mathbf{K_{*,f}} & \mathbf{K_{*,*}} \end{bmatrix}\right), \tag{7}$$

where $\mathbf{K}$ is subscripted by the variables between which the covariance is computed (and we use the asterisk $*$ as shorthand for $\boldsymbol{f}_*$). We then combine the prior with a likelihood $p(\boldsymbol{y}|\boldsymbol{f}) = \mathcal{N}(0, \sigma_{noise}^2 \mathbb{I})$, $\mathbb{I}$ is the identity matrix, and using Bayes's rule we obtain the joint posterior

$$p(\boldsymbol{f}, \boldsymbol{f}_*|\boldsymbol{y}) = \frac{p(\boldsymbol{f}, \boldsymbol{f}_*)p(\boldsymbol{y}|\boldsymbol{f})}{p(\boldsymbol{y})}, \tag{8}$$

where $\boldsymbol{y} = (y_1, \ldots, y_n)$ is the vector of observations. Note that we have dropped the conditioning on inputs while defining the above probabilities for notational simplicity. However, these probabilities defining a GP model are always conditional on the corresponding inputs. By marginalizing the training set latent function values $\boldsymbol{f}$ we get the desired posterior function values at test locations $\boldsymbol{X}_*$ given by

$$p(\boldsymbol{f}_*|\boldsymbol{y}) = \int p(\boldsymbol{f}, \boldsymbol{f}_*|\boldsymbol{y})d\boldsymbol{f} = \frac{1}{p(\boldsymbol{y})} \int p(\boldsymbol{y}|\boldsymbol{f})p(\boldsymbol{f}, \boldsymbol{f}_*)d\boldsymbol{f}. \tag{9}$$

Since both the joint GP prior and the likelihood are Gaussian we can evaluate the above integral analytically to obtain the posterior latent function at the test locations given by

$$p(\boldsymbol{f}_*|\boldsymbol{y}) = \mathcal{N}(\mathbf{m_f}, \boldsymbol{\Sigma_f}), \tag{10}$$

with the following first and second moment [20]:

$$\mathbf{m_f} = \mathbf{K_{*,f}}(\mathbf{K_{f,f}} + \sigma^2_{\mathbf{noise}}\mathbb{I})^{-1}\boldsymbol{y}$$
$$\boldsymbol{\Sigma_f} = \mathbf{K_{*,*}} - \mathbf{K_{*,f}}(\mathbf{K_{f,f}} + \sigma^2_{\mathbf{noise}}\mathbb{I})^{-1}\mathbf{K_{f,*}} \tag{11}$$

The hyperparameters $\boldsymbol{\phi}$ can be obtained as maximum likelihood estimates by maximizing the log marginal likelihood given by

$$\log P(\boldsymbol{y}|\boldsymbol{\phi}) = -\frac{1}{2}\boldsymbol{y}^T(\mathbf{K_{f,f}} + \sigma^2_{noise}\mathbb{I})^{-1}\boldsymbol{y} - \frac{1}{2}\log\left|(\mathbf{K_{f,f}} + \sigma^2_{noise}\mathbb{I})^{-1}\right| - \frac{1}{2}\log(2\pi). \tag{12}$$

We discuss how GP emulators can be refined in terms of choosing training sites to evaluate the latent function in Section 2.5.2.

1.3.2 Gaussian processes classification

As mentioned previously, we will tackle discontinuities present in the simulator response surface using a boundary detector built using a classifier. Gaussian processes can be used for classification purposes in a discriminative probabilistic [20] framework. Thus we would use a GP classifier to detect the boundaries of discontinuities. Furthermore, we would exploit its probabilistic predictions to propagate uncertainty about the boundaries for carrying out active learning (we discuss this in section 2.5). Next, we briefly review the method for GP classification.

In a classification problem the input remains the same as that of regression but the outputs are discrete class occupancy labels $t_i \in \{-1, +1\}$ ('-1' for not in the class of interest, and '$+1$' for in the class). We are interested in predicting the class membership for a test point $\boldsymbol{X}_*$. This is achieved using a latent function $g(\boldsymbol{X}_*)$ whose value is mapped to the unit interval by means of a sigmoid function $sig : \mathbb{R} \to [0, 1]$ such that [20]

$$\pi := p(t_* = +1|\boldsymbol{X}_*) = sig(g(\boldsymbol{X}_*)). \tag{13}$$

The class membership probability must normalise, $\sum_{t_*} p(t_* = +1|\boldsymbol{X}_*) = 1$, thus we have $p(t_* = -1|\boldsymbol{X}_*) = 1 - p(t_* = +1|\boldsymbol{X}_*)$. The sigmoid function takes the form

$$sig(g(\boldsymbol{X})) = \frac{1}{1 + \exp^{-g(\boldsymbol{X})}}. \tag{14}$$

The likelihood of the class labels $\boldsymbol{t} = (t_1, \ldots, t_n)$ for n data points is assumed to be a Bernoulli distribution given by

$$p(\boldsymbol{t}|\boldsymbol{g}) = \prod_{i=1}^{n} p(t_i|g_i) = \prod_{i=1}^{n} sig(t_i|g_i), \tag{15}$$

where we have assumed that the class labels are i.i.d. and $\boldsymbol{g} = (g_1, \ldots, g_n)$ is the vector of latent function values.

Now just like the regression case, we can put a joint prior $p(\boldsymbol{g}, \boldsymbol{g}_*)$ on the training ($\boldsymbol{g}$) and test ($\boldsymbol{g}_*$) latent function values. This immediately enables the

use of the standard GP machinery to obtain the posterior predictive distribution over the class labels

$$\pi := p(\boldsymbol{t}_* = +1|\boldsymbol{t}) = \int sig(\boldsymbol{t}_*|\boldsymbol{g}_*)p(\boldsymbol{g}_*|\boldsymbol{t})d\boldsymbol{g}_*. \tag{16}$$

Unfortunately the posterior term $p(\boldsymbol{g}_*|\boldsymbol{t})$ is intractable as it involves an integration over the likelihood given by equation (15) with a sigmoid nonlinearity. Approximation schemes can be used to overcome this intractability. While MCMC methods provide the closest approximation [21] such methods are often found to be extremely slow for a datasets of even moderate size. Expectation propagation (EP, [22]) is an iterative deterministic approximation scheme that is widely used for inference in GP classifiers as it provides good accuracy and is much faster than MCMC. See Nickisch & Rasmusson [23] for a review and comparison of different approximations for inference in a binary GP classifier. In EP, the individual (per data point) sigmoidal likelihood terms are approximated by un-normalised Gaussians $\xi_i(g_i)$. We term these local Gaussian approximations as *site* functions. Thus, the likelihood $p(t_i|g_i)$ for i-th data point t_i is approximated as

$$
\begin{aligned}
p(t_i|g_i) &= sig(t_i|g_i), \\
&\approx \xi_i(g_i; \tilde{\mu}_i, \tilde{\sigma}_i, \tilde{Z}_i), \\
&\approx \tilde{Z}_i\mathcal{N}(g_i; \tilde{\mu}_i, \tilde{\sigma}_i),
\end{aligned}
\tag{17}
$$

with site parameters $\{\tilde{\mu}_i, \tilde{\sigma}_i, \tilde{Z}_i\}$. For convenience we can write the product of the local approximate likelihoods as

$$\prod_{i=1}^{n} \xi_i(g_i; \tilde{\mu}_i, \tilde{\sigma}_i, \tilde{Z}_i) = \mathcal{N}(\tilde{\boldsymbol{\mu}}, \tilde{\Sigma}) \prod_{i=1}^{n} \tilde{Z}_i, \tag{18}$$

where $\tilde{\boldsymbol{\mu}} = (\tilde{\mu}_1, \ldots, \tilde{\mu}_n)$ and $\tilde{\Sigma}$ is a diagonal matrix with diagonal elements $\tilde{\sigma}_i$. The posterior latent function $p(\boldsymbol{g}|\boldsymbol{t})$ is then approximated using the site functions as

$$p(\boldsymbol{g}|\boldsymbol{t}) \approx q(\boldsymbol{g}|\boldsymbol{t}) = \frac{1}{Z_{EP}}p(\boldsymbol{g}) \prod_{i=1}^{n} \tilde{Z}_i\mathcal{N}(g_i; \tilde{\mu}_i, \tilde{\sigma}_i) = \mathcal{N}(\boldsymbol{\mu}, \Sigma), \tag{19}$$

where $p(\boldsymbol{g}|\boldsymbol{\phi})$ is the standard Gaussian prior with covariance K and $Z_{EP} = p(\boldsymbol{t}|\boldsymbol{\phi}) = \int p(\boldsymbol{t}|\boldsymbol{g})p(\boldsymbol{g}|\boldsymbol{\phi})$ is the marginal likelihood. The posterior mean and variance is given by [20]

$$
\begin{aligned}
\boldsymbol{\mu} &= \Sigma\tilde{\Sigma}^{-1}\tilde{\boldsymbol{\mu}}, \\
\Sigma &= (K^{-1} + \tilde{\Sigma}^{-1})^{-1}.
\end{aligned}
\tag{20}
$$

Note that we have made the parameter dependency explicit in the prior $p(\boldsymbol{g}|\boldsymbol{\phi})$ while defining the marginal likelihood as in the regression case.

The task of EP is then to find each of the site parameters iteratively so that the marginal posterior is as accurate as possible. To this end, we first combine

the prior and the site functions into an approximate marginal distribution, also known in machine learning parlance as a *cavity* distribution [20]:

$$q_{-i}(g_i) \propto \int p(\boldsymbol{g}) \prod_{j \neq i} \xi_j(g_j; \tilde{\mu}_j, \tilde{\sigma}_j, \tilde{Z}_j) dg_j, \tag{21}$$

where the subscript '$-i$' means "all but the i-th". The cavity distribution is an approximation to the marginal distribution of the latent function at the i-th site obtained by combining the prior $p(\boldsymbol{g})$ with $n - 1$ (all but the i-th) approximate likelihood terms ξ_i. The simplest way to to obtain the cavity distribution is by first finding the i-th approximate posterior from the joint in equation (19) and then dividing it by the i-th site function ξ_i. Thus we have the cavity distribution at the i-th site as [20]

$$q_{-i}(g_i) = \mathcal{N}(g_i | \mu_{-i}, \sigma^2_{-i}), \tag{22}$$

where $\mu_{-i} = \tilde{\sigma}^2_{-i}(\sigma_i^{-2}\mu_i - \tilde{\sigma}_i^{-2}\tilde{\mu}_i)$ and $\sigma^2_{-i} = (\sigma_i^{-2} - \tilde{\sigma}_i^{-2})^{-1}$.

We then find, for each site, a new un-normalised Gaussian marginal with parameters $\{\hat{\mu}_i, \hat{\sigma}_i, \hat{Z}_i\}$ which best approximates the product of the cavity distribution and the exact likelihood at each site.

$$\hat{q}(g_i) := \xi_i(g_i; \hat{\mu}_i, \hat{\sigma}_i, \hat{Z}_i) \approx q_{-i}(g_i)p(t_i|g_i). \tag{23}$$

The parameters of the Gaussian $\hat{q}(g_i)$ are found by moment matching with the right hand side of equation (23). Finally, the site parameters $\{\tilde{\mu}_i, \tilde{\sigma}_i, \tilde{Z}_i\}$ of the likelihood approximation ξ_i are obtained in turn from the updated moments of $\hat{q}(g_i)$. We refer the reader to [20] for a detailed derivation of the desired moments.

This procedure is carried out iteratively where in each sweep all the individual site functions are fitted and the sweeps are carried out until the convergence of the site parameters of ξ_i for all the sites. In practice a fixed number of sweeps, say 20, generally suffices for convergence. The converged site parameters are used to obtain the posterior latent function, at test locations [20],

$$p(\boldsymbol{g}_*|\boldsymbol{t}) \approx q(\boldsymbol{g}_*|\boldsymbol{t}) = \mathcal{N}(\mathbf{m_{g_*}}, \boldsymbol{\Sigma_{g_*}}), \tag{24}$$

where,

$$\begin{aligned} \mathbf{m_{g_*}} &= \mathbf{K_{*,g}}(\mathbf{K_{g,g}} + \tilde{\boldsymbol{\Sigma}})^{-1}\tilde{\boldsymbol{\mu}}, \\ \boldsymbol{\Sigma_{g_*}} &= \mathbf{K_{*,*}} - \mathbf{K_{*,g}}(\mathbf{K_{g,g}} + \tilde{\boldsymbol{\Sigma}})^{-1}\mathbf{K_{g,*}}. \end{aligned} \tag{25}$$

Substituting this value of $p(\boldsymbol{g}_*|\boldsymbol{t}) \approx q(\boldsymbol{g}_*|\boldsymbol{t})$ into equation (16), and approximating the sigmoidal function $sig(t_i|g_i) \approx \Phi(t_i|g_i)$ by a probit function[1] we can evaluate the integral in equation (16) to obtain the posterior class membership probability

---

[1]A probit function $\Phi(x)$ is the standard CDF of a normal distribution given by $\Phi(x) = \int_{-\infty}^{x} \mathcal{N}(y)dy$. EP gives us a very good Gaussian approximation as $\mathcal{N}(\mathbf{m_{g_*}}, \boldsymbol{\Sigma_{g_*}})$.

as

$$\hat{\pi} := p(\boldsymbol{t}_* = +1|\boldsymbol{t}) = \int \Phi(\boldsymbol{t}_*|\boldsymbol{g}_*)\mathcal{N}(\mathbf{m}_{\mathbf{g}_*}, \boldsymbol{\Sigma}_{\mathbf{g}_*})\mathbf{d}\boldsymbol{g}_*,$$

$$= \Phi\left(\frac{\mathbf{m}_{\mathbf{g}_*}}{\sqrt{1 + \boldsymbol{\Sigma}_{\mathbf{g}_*}}}\right). \tag{26}$$

We also use the site parameters to maximise $\ln(Z_{EP})$ to obtain the maximum likelihood estimates of the hyperparameters. From the likelihood approximations we directly obtain the marginal likelihood as the function of the site parameters given by

$$\ln(Z_{EP}) \approx \ln \int \prod_{i=1}^{n} \xi_i(g_i; \tilde{\mu}_i, \tilde{\sigma}_i, \tilde{Z}_i) p(\boldsymbol{g}|\boldsymbol{\phi}) d\boldsymbol{g}$$

$$= \sum_{i=1}^{n} \ln \frac{\tilde{Z}_i}{\sqrt{2\pi}} - \frac{1}{2}\boldsymbol{V}^T\left(\mathbf{K_{g,g}}^{-1} + \mathbf{K_{g,g}}^{-1}\tilde{\Sigma}^{-1}\mathbf{K_{g,g}}^{-1}\right)\boldsymbol{V} - \frac{1}{2}\ln\left|\mathbf{K_{g,g}} + \tilde{\Sigma}^{-1}\right|, \tag{27}$$

where,

$$\boldsymbol{V} = \left[\mathbb{I} - \mathbf{K_{g,g}}\left(\mathbf{K_{g,g}} + \tilde{\Sigma}^{-1}\right)\right]\mathbf{K_{g,g}}\tilde{\Sigma}\tilde{\boldsymbol{\mu}}, \tag{28}$$

and $|A|$ denotes the determinant of matrix A. Note that each iteration for maximizing $\ln Z_{EP}$ with respect to $\boldsymbol{\phi}$ in equation (27) requires the estimate of site parameters and thus a number of sweeps of EP. Thus the computational cost of maximising the marginal likelihood is much higher in this case compared to regression.

1.4   Sparse approximations

GP models suffer from high computational load for inference computations. For n training points exact inference as used in GP regression requires $O(n^3)$ effort while for EP approximation a sequence of $O(n^3)$ operations are required.

There is an active line of research whose aim is to alleviate this computational bottleneck by using a sparse approximation of the true covariance function. Some of these methods are reviewed in [24]. The common approach advocated by these methods is to use a set of m inducing (or imaginary) inputs $\boldsymbol{X}_u$ with associated latent function $\boldsymbol{u}$ to reduce the computational load to $O(nm^2)$. We denote the $n \times n$ covariance matrix between the training inputs as $\mathbf{K}$, the $m \times n$ covariance matrix between the inducing and training inputs as $\mathbf{K_u}$ and the $m \times m$ covariance matrix between the inducing inputs as $\mathbf{K_{u,u}}$. The most widely used approximation scheme [24], the FITC approximation $\hat{\mathbf{K}}$ to the full covariance $\mathbf{K}$ is given by

$$\mathbf{K} \approx \hat{\mathbf{K}} = \mathbf{Q} + \text{diag}(\mathbf{K} - \mathbf{Q}), \tag{29}$$

where $\text{diag}(A)$ is a diagonal matrix whose elements match the diagonal of A and

the matrix $\mathbf{Q}$ is given by

$$\mathbf{Q} = \mathbf{K_u}^T \mathbf{Q_{u,u}} \mathbf{K_u}, \tag{30}$$

$$\mathbf{Q_{u,u}} = \mathbf{K_{u,u}} + \sigma_{n_u}^2 \mathbb{I}, \tag{31}$$

where $\sigma_{n_u}^2$ is the noise from inducing inputs. $\hat{\mathbf{K}}$ has the same diagonal elements as $\mathbf{K}$ and the off-diagonal elements are the same as for $\mathbf{Q}$. This sparse approximation was first introduced in [25] to scale the GP regression and later it was introduced in the context of a GP classifier in [26] within the EP approximation.

2 Methods

Having introduced the GP machinery we will now proceed towards applying the GP classification and regression to build a two-step emulator.

2.1 A GP classifier for segmenting the APD response surface: boundary detector

As mentioned in the previous section, our approach of using a classifier is primarily motivated by the idea of boundary detection and as a consequence segmentation of the input domain. However, unlike the previous domain segmentation attempts in the realm of computer experiments where a tree-based or non-stationary GP is constructed, in our application we can define a priori a set of possible labels to the APD_{90} values (and the corresponding region of input space) based on the depolarisation/repolarisation pattern of the membrane voltage. The goal is to use a small number of simulated APD_{90} values obtained by varying the parameters to train a classifier to predict the labels for a much larger set of test inputs with a quantifiable measure of uncertainty. We denote the inputs (here, scaling factors between zero and one applied to each maximal conductance of an ion current) as $\boldsymbol{R} = (R_1, \ldots, R_D)$, where D is the number of ion currents under consideration (and dimension of the input space) and each element of which can take any value between $[0, 1]$.

For the n-th input vector our $\boldsymbol{R}_n$ our simulator $\mathcal{S}$ returns a set $\{y_n, k_n\} = \mathcal{S}(\boldsymbol{R}_n)$, where y_n is the APD_{90} value and $k_n \in \{1, 2, 3\}$ is a label associating the n-th input with any one of the three observed categories of action potential as shown in Figure 1(a). We chose the following convention for labelling the action potential (see Fig. 1): $k = 1$ for no-depolarisation, $k = 2$ for a valid action potential, and $k = 3$ for no-repolarisation. Note that the simulator only returns an APD value in y_n when the input is within the valid AP region $k = 2$, and an error code denoting which of the other regions it is in otherwise.

As our problem is inherently a multi-class classification we adopt a One-versus-Rest (OVR) method of classification. Using OVR we build one binary GP classifier, as introduced in section 1.3.2, for each class $k \in \{1, 2, 3\}$ with associated labels $\boldsymbol{t}^k \in \{+1, -1\}$ (+1 for the k-th class and -1 for the rest of the classes), to predict the probability

$$\pi_*^k := p(t_*^k = +1 | R_*, \boldsymbol{R}, \boldsymbol{t}^k), \tag{32}$$

that a test input $\boldsymbol{R}_*$ given the training inputs $\boldsymbol{R} = (\boldsymbol{R}_1, \ldots, \boldsymbol{R}_n)$ and OVR labels $\boldsymbol{t}^k$ belongs to the k-th class. The predicted class label k_* for the test input $\boldsymbol{R}_*$ is then simply the most likely class:

$$k_* = \operatorname*{argmax}_{k} (\pi_*^k). \tag{33}$$

2.2 GP regression response surface prediction

Again we consider a simulated dataset where $\{\boldsymbol{R}_n, y_n\} = \mathcal{S}(\boldsymbol{R}_n)$ is the n-th input-output pair which gives rise to a valid action potential and the associated APD_{90} value returned by the simulator. We wish to learn a latent function $\boldsymbol{f}$ that is an emulator of the simulator $\mathcal{S}$. Thus we have the output of the emulator y_n (the APD_{90} values) given by

$$y_n = \boldsymbol{f}(\boldsymbol{R}_n). \tag{34}$$

Now we place a zero mean GP prior on $\boldsymbol{f}$ as

$$\boldsymbol{f} \sim \mathcal{N}(0, K(\boldsymbol{R}, \boldsymbol{R}'; \boldsymbol{\phi})) \tag{35}$$

where $K(\boldsymbol{R}, \boldsymbol{R}'; \boldsymbol{\phi})$ is a covariance function parametrized by hyperparameters $\boldsymbol{\phi}$. Note that this covariance is separate from the classifier covariance although we may use the same kernel function.

Given the training data $\{\boldsymbol{R}_*, \boldsymbol{y}\}$, the posterior mean at a new test point $f_* := f(\boldsymbol{R}_*)$ is given by (using equation (11))

$$\mu(f_*) = K(\boldsymbol{R}, \boldsymbol{R}_*) - K(\boldsymbol{R}, \boldsymbol{R})^{-1}\boldsymbol{y}, \tag{36}$$

and the posterior variance as (again using equation (11))

$$Var(f_*) = K(\boldsymbol{R}_*, \boldsymbol{R}_*) - K(\boldsymbol{R}, \boldsymbol{R}_*)^T K(\boldsymbol{R}, \boldsymbol{R})^{-1} K(\boldsymbol{R}_*, \boldsymbol{R}). \tag{37}$$

The hyperparameters are obtained as maximum likelihood estimates by maximizing the log marginal likelihood of the GP given by (following equation (12))

$$\log P(\mathbf{y}|\boldsymbol{\phi}) = -\frac{1}{2}\boldsymbol{y}^T K(\boldsymbol{R}, \boldsymbol{R})^{-1}\boldsymbol{y} - \frac{1}{2}\log|K(\boldsymbol{R}, \boldsymbol{R})| - \frac{1}{2}\log(2\pi). \tag{38}$$

2.3 Two-step emulator

We combine the boundary detector (using GP classification) and surface emulator (using GP regression) in a sequential manner to design a two-step emulator. In the training phase we use the simulator to create a training dataset of n points: $\mathcal{D}_{train} = \{(\boldsymbol{R}_i, y_i, k_i) \text{ for } i = 1, \ldots, n\}$. We draw the values of $\boldsymbol{R}_i$ from $\mathcal{U}(0, 1)$. We then learn the GP hyperparameters associated with the boundary detector (with the OVR method using binary GP classifiers) and the surface emulator using the training dataset $\mathcal{D}_{train}$. Note that we use a subset of $\{\boldsymbol{R}_i, y_i\}$ for training the surface emulator. This subset contains only those inputs that generate a valid action potential — that is, all training points in this subset have the same associated class label $k = 2$.

In the test/prediction phase for test input vector $\boldsymbol{R}_*$ we first use the boundary detector and obtain the class labels k_* which we subsequently use to segment those test inputs into three domains: $\boldsymbol{R}_*^{NoDep}$ for which the membrane potential does not depolarise (that is no AP is generated); $\boldsymbol{R}_*^{AP}$ where we observe an

AP; and $\boldsymbol{R}_*^{NoRep}$ where the membrane potential does not repolarise after the occurrence of an AP as shown in Fig. 1. Since we are interested in the AP region we pass $\boldsymbol{R}_*^{AP}$ to the surface emulator to obtain the posterior predictive given by (using equation (36))

$$y^{AP} \approx \hat{f}_* \sim \mathcal{N}\big(\mu(f_*), Var(f_*)\big). \tag{39}$$

2.4 Choice of GP covariances

In order to use both the classifier and surface GP one has to choose a suitable covariance function a-priori which embeds our prior assumption about the function that we are trying to model with a GP. For the classifier GP we use a squared exponential covariance given by equation (6) and for the surface GP we use the rational quadratic covariance function given by

$$k_{RQ}(\boldsymbol{R}_i, \boldsymbol{R}_j) = \nu^2 \left(1 + \frac{\|\boldsymbol{R}_i - \boldsymbol{R}_j\|^2}{2\alpha l^2}\right)^{-\alpha}, \tag{40}$$

where $\{\nu, \alpha, l\}$ are the covariance hyperparameters. This covariance is equivalent to adding many squared exponential covariance with different lengthscales. As a result we expect to see functions varying across different lengthscales. Our prior intuition about the APD response surface is that it is smoother away from the boundary and varies much more rapidly near the boundary. Thus we have used this covariance function to accommodate different lenghtscales (degree of variation). Here the hyperparameter α determines the relative weighting of large-scale and small-scale variations.

We have used the GPML package [27] called from MATLAB to implement the surface and boundary GPs. That code is available from `http://www.gaussianprocess.org/gpml/code/matlab` and our simulator and emulator codes are available as described in Section 6.

2.5 Active learning

Since we are using a probabilistic framework for both classification and regression, we can exploit the uncertainty associated with the predictions to choose the training inputs using some form of adaptive scheme, as opposed to picking training points at random. This is known as *active learning* [28]. The main idea behind active learning is to sequentially add inputs to the training set to reduce uncertainty in predictions away from the training locations. Choosing the inputs actively has the potential to significantly reduce the required training budget, i.e. the number of simulations needed to generate $\mathcal{D}_{train}$. This is what we turn to next.

2.5.1 Active learning for boundary detection using GP classification

To carry out active learning of a GP classifier our goal is to augment a set of n_1 initial training inputs $\boldsymbol{R}_\emptyset$ with either a single new training point $\boldsymbol{R}_*$ or a set of such points $\left\{\boldsymbol{R}_{new_j}\right\}_{j=1,\ldots,n_s}$ of size n_s that convey more information about the

14

boundaries in comparison to an equal number of randomly chosen points. We would like to point out that previously we have denoted a test point as $\boldsymbol{R}_*$ where the subscript $*$ denotes a point where we draw predictions from a GP. In this context, the same notation applies to those points where we draw predictions during the active learning process. We do this by finding those inputs about which the GP classifier is most uncertain, by using a suitable criterion to quantify this uncertainty. Information theoretic criteria such as the conditional entropy have been used to quantify uncertainty and carry out active learning for a GP classifier previously [29]. To quantify uncertainty through conditional entropy, we need to evaluate posterior expectations that do not have a closed form (due the the sigmoidal likelihood of the GP classifier). Gaussian approximation is known to work well for a binary GP classifier [29] and is considered a state-of-the-art technique to carry out active learning for a classifier GP. Although it is possible to extend the active learning approach proposed in Houlsby *et al.* [29] to the multi-class OVR method, we propose an alternative quantification of uncertainty using the posterior prediction probability as:

$$c(\boldsymbol{R}_*) = \max(\pi_*^k) - \max(\pi_*^{k_-}), \tag{41}$$

where π_*^k is the classification probability for the most-likely class, given by equation (32), and the second term is the probability of classification into the second-most-likely class. We term c the *certainty* of classifier predictions as $c = 1$ means the classifier is absolutely certain and $c = 0$ represents equal probability of being in either of the two most likely classes. We are essentially characterising the regions in the input space that lead to an overlap of possible class distributions, and these are quantified as $c \to 0$. We propose to use a particle based optimisation algorithm to find regions of minimum certainty, and (after this has converged) to add n_s new points returned by the particle based optimiser to our training set to obtain the actively learnt input set $\boldsymbol{R} = \boldsymbol{R}_\emptyset \cup \boldsymbol{R}_{new}$. Moreover, we can carry out this procedure sequentially while using the optimiser at each step to find $\boldsymbol{R}_{new}$ and then setting $\boldsymbol{R}_\emptyset = \boldsymbol{R}_\emptyset \cup \boldsymbol{R}_{new}$. Repeating this for r rounds we obtain the active set of inputs $\boldsymbol{R}$ consisting $n_3 = r \times n_s$ new active training points, and thus a total training set at which the full simulator must be run of size $N = n_1 + n_3$. A flow diagram of active classifier learning is presented in Fig. 3. For carrying out the optimisation we use a particle swarm optimisation (PSO) algorithm [30, 31] and to be synonymous with the terminology of PSO we will denote the set of points $\boldsymbol{R}_{new}$ considered at each round as a *swarm* of possible training points.

Note that to generate distinct and useful training points for the classifier we purposefully refrain from running the PSO up to convergence, to maintain some distance (in input space) between each of the particles in the swarm. Thus in practice we stop the PSO iterations when the average certainty of the swarm goes below a threshold θ. We found satisfactory performance by setting $\theta = 0.5$. To visualise the learning mechanism we evaluate this active learning scheme on a 2-input simulator (and emulator) of the O'Hara model with inputs as:

i) The sodium channel conductance scalings – R_{Na}; and

15

ii) The hERG channel conductance scalings – R_{Kr};

with APD$_{90}$ as the output.

We started with $n_1 = 10$ initial training points drawn randomly over the input space. We ran $r = 10$ rounds of active learning with a swarm size of $n_s = 10$ and show the consecutive uncertainty contours for rounds $r = \{3, 4, 5\}$ in Fig. 4. We also carried out a random design by adding 10 points at random to the initial 10 training points, and compare the certainty with that of active learning. The contour plots are shown in Fig. 4.

It is evident from the contour plots for active learning that after round $r = 4$ the classifier is able to get an estimate of the boundaries and the active points start to gather in the regions of low certainty. At the final ($r = 10$) round the active learning method is able to increase the certainty significantly more than the random design.

2.5.2   Active learning for surface emulation

Active learning for GP regression is a well studied topic and different schemes have been proposed. The schemes differ mainly in the criteria with which uncertainty is quantified. The most widely used of these is the entropy criterion [32–34]. Alternative criteria such as mutual information [19] and integrated variance [35] have been proposed recently. Note that actively picking training points based on the entropy as well as mutual information is a NP-hard problem [36] and thus greedy algorithms are used while using any of these information theoretic criteria. Note that optimal algorithms have also been proposed [28, 37] recently that integrate active learning with covariance hyperparameter learning, so that the selection of new training locations are optimal for carrying out Bayesian inference of the hyperparameters.

Most active learning methods have been used in spatial modelling where i) the training locations are 2D grids and ii) the number of training and test data points are much smaller than we can afford in our application where full simulator evaluations are relatively cheap (but not so cheap that an emulator is not desirable). In our application, the dimensionality of the inputs will be over two or three, and thus using a full grid is impractical. For these reasons we approach the active learning problem using a greedy algorithm, for computational tractability, that utilises the entropy criteria.

In order to set up the active learning consider a pool of *candidate points* $\left\{ \boldsymbol{R}_{o_j} \right\}_{j=1,\ldots,N_c}$ which are locations spanning the input domain, drawn randomly, from which we choose training points based on a chosen uncertainty criterion. Our aim is then to choose a new training point $\boldsymbol{R}_* \in \boldsymbol{R}_o$ such that it gives us more information about the domain than choosing $\boldsymbol{R}_*$ randomly, resulting in greater prediction accuracy at lower simulation cost. This can be achieved by quantifying the uncertainty associated with the point $\boldsymbol{R}_*$, having observed a small initial dataset $\boldsymbol{R}_\emptyset$ consisting of n_1 points, where $n_1 \ll N_c$ for which the simulator returns a valid AP. Intuitively we want to choose $\boldsymbol{R}_*$ as points of maximum uncertainty. We can quantify this uncertainty using posterior

conditional entropy between the latent function $f(\boldsymbol{R}_*)$ and the observations $\boldsymbol{y}_\emptyset = \mathcal{S}(\boldsymbol{R}_\emptyset)$ returned by the simulator. This conditional entropy is given by [19]

$$H(f(\boldsymbol{R}_*)|\boldsymbol{y}_\emptyset) = -\int p(f(\boldsymbol{R}_*), \boldsymbol{y}_\emptyset) \log p(f(\boldsymbol{R}_*)|\boldsymbol{y}_\emptyset) df(\boldsymbol{R}_*) d\boldsymbol{y}_\emptyset, \qquad (42)$$

which can be evaluated in closed form, since $p(f(\boldsymbol{R}_*)|\boldsymbol{y}_\emptyset)$ is the posterior Gaussian density of the surface GP, as

$$H(f(\boldsymbol{R}_*)|\boldsymbol{y}_\emptyset) = \frac{1}{2} \log \left(\mathrm{Var}(f(\boldsymbol{R}_*)) \right) + \frac{1}{2} \log \left((2\pi) + 1 \right), \qquad (43)$$

where $Var(f(\boldsymbol{R}_*))$ is the posterior variance given by equation (37). We can now choose the point $\boldsymbol{R}_*$ from the N_c candidate points $\boldsymbol{R}_o$ that has the largest conditional entropy:

$$\boldsymbol{R}_* = \underset{k}{\mathrm{argmax}} \left[H\left(f(\boldsymbol{R}_o|\boldsymbol{y}_\emptyset) \right) \right] \qquad (44)$$

In order to implement such a scheme we have to be aware of the discontinuities. We start with a small random training set evaluated using the simulator at $\boldsymbol{R}_\emptyset$ for the surface (discarding non-AP points). We use a small n_1 to minimise the simulation cost. We also obtain an initial estimate of the hyperparameters using n_1. To remove the non-AP points from the candidates in $\boldsymbol{R}_o$ we can use the classifier. Inclusion of the classifier alleviates full simulations of all the candidate points. Carrying out the above procedure we get the new training set as $\boldsymbol{R} = \boldsymbol{R}_\emptyset \cup \boldsymbol{R}_*$. We can also do this sequentially by setting $\boldsymbol{R}_\emptyset = \boldsymbol{R}_\emptyset \cup \boldsymbol{R}_*$ and then again finding one new point $\boldsymbol{R}_*$ using equation (44).

We can carry on this iterative procedure for n_2 rounds to collect n_2 active training input points. Note that due to misclassification some of the candidate points in $\boldsymbol{R}_\emptyset$ may be in non-AP regions. When we evaluate the simulator on the n_2 active points we remove these non-AP points before forming the surface GP training set. Thus the resulting set of active points is of size $\hat{n}_2 \le n_2$ and our new training set after carrying out active learning consists of $N = n_1 + \hat{n}_2$ training points. Also note that the classifier is used only once on the entire candidate set before adding an active point. The various steps involved in the surface active learning are summarised and presented in Fig. 5 as a flow diagram.

We again set up the 2 input problem exactly as illustrated previously for classification. We started with $n_1 = 11$ initial training points drawn randomly over the input space and carried out active learning for $n_2 = 90$ rounds using candidate inputs $\left\{ \boldsymbol{R}_{o_j} \right\}_{j=1,\dots,N_c}$ of size $N_c = 10{,}000$. We also carried out a random design by adding points sequentially to the training set starting with the initial n_1 points for a total of 90 rounds. Fig 6 shows a comparison of consecutive entropies during active learning rounds $\{10, 11, 12\}$. We also show in the same figure (bottom row) the comparison of entropies between active learning and random design at the final round. We observe during rounds $\{10, 11, 12\}$ that the active points (blue circles with cross-hair in Fig 6) are placed at places of high entropy. Active learning is able to reduce the entropies better than random design.

17

In our active learning scheme we restrict the candidate points to be within the valid AP region by i) before simulation filtering out non-AP candidate points using the classifier; and ii) after simulation, discarding any non-AP points that were misclassified. During the active learning process the entropies are higher in the regions near the boundaries as only a few candidate points (due to misclassification) exist in the non-AP regions. A well-studied pathology of entropy based active learning [28] is that a lot of training points are selected near the edges of the surface, that is the regions of highest entropies. As the active learning progresses the boundary regions emerge as regions of high uncertainty for the reasons described above. This is something that we notice in the contour plots. However, this behaviour actually works in our favour for this application as the surface changes most rapidly near the boundary region and thus having training points in those regions leads to a better surface prediction.
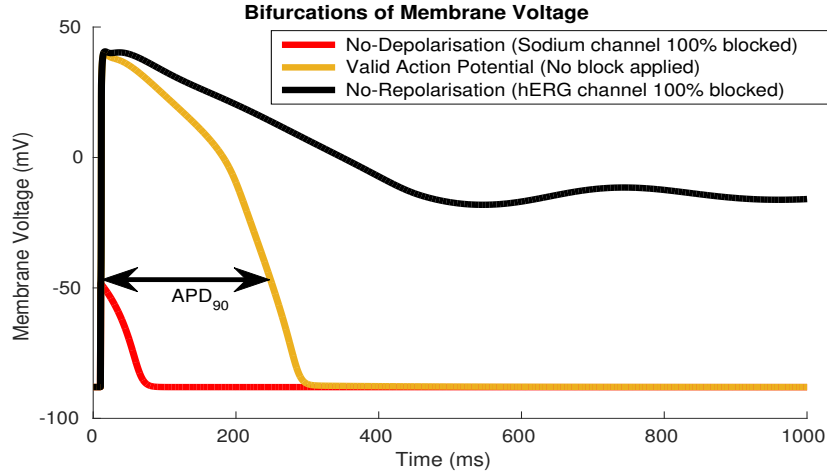
(a)
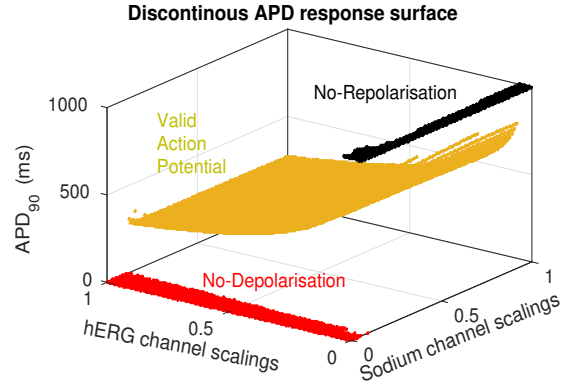
(b)

Figure 1: **Bifurcation induced discontinuities in the APD biomarker response surface for the O'Hara model**. a) Bifurcations in the O'Hara model dynamics [12] lead to three distinct types of behaviour, with no smooth transition in APD_{90}. The black line indicates that the membrane voltage fails to return to the resting membrane potential. The orange line shows a valid AP. The red line shows the failure of depolarisation. b) Resulting discontinuous APD response surface evaluated on a dense grid of 100×100 sample points. The parameters are conductance scalings of the hERG and sodium channels. The three regions of the parameter space are shaded and colour coded according to the depolarisation/repolarisation patterns as seen above.
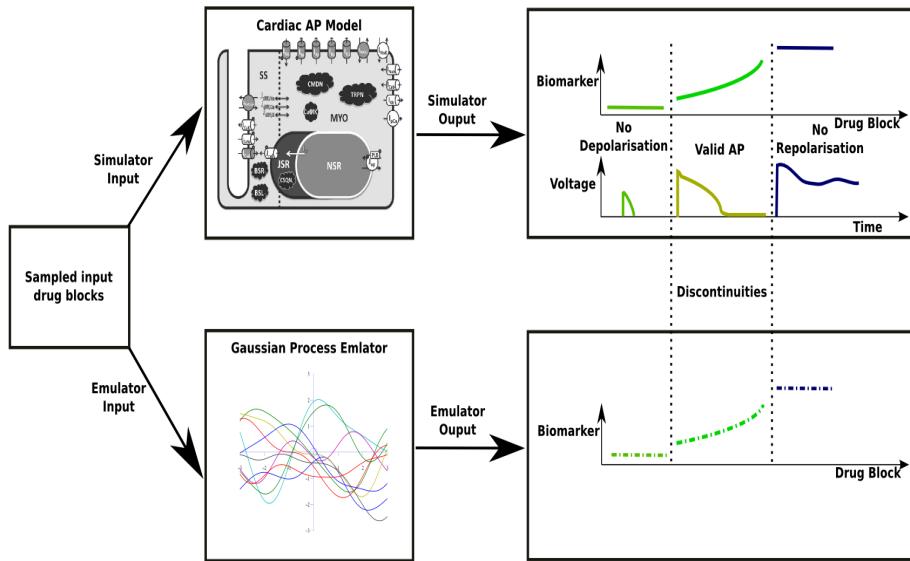
Figure 2: Comparison of the steps involved in simulation vs. emulation of discontinuous biomarker response surface.
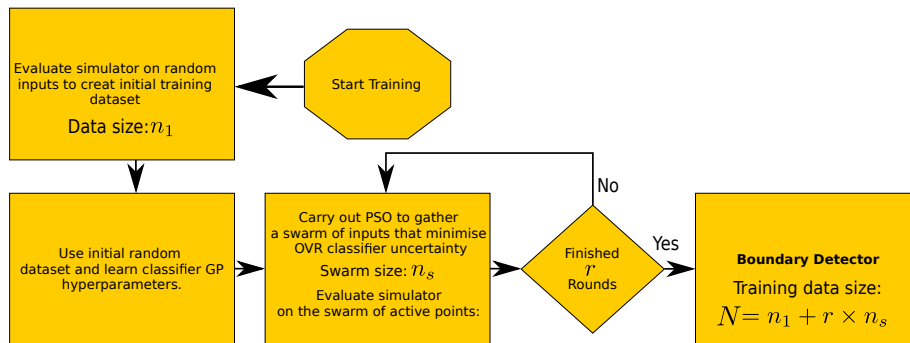
Figure 3: **Steps involved in active learning for the boundary detector**. The algorithmic settings, supplied by the user, for this active learning process are the initial data size n_1, number of rounds r and the swarm size n_s.
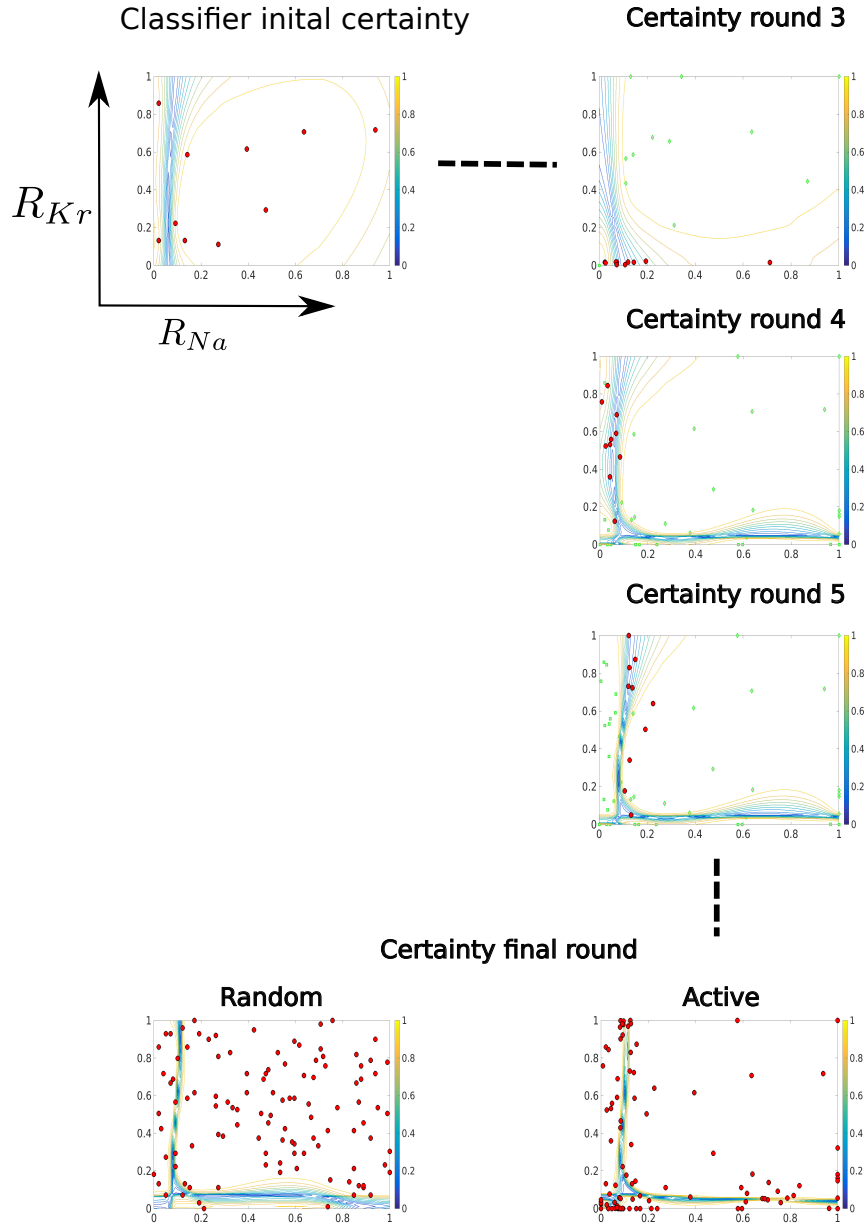
Figure 4: **Comparison of the measure of certainty c in classification between actively and randomly adding training points**. The certainty as defined by equation (41) is shown as contour plots. The training points (accumulated thus far) are shown as red circles. The darker shade of the contours represents the areas of least certainty spread along the boundaries (discontinuities) as estimated by the classifier. We show three consecutive rounds of active learning $r = \{3, 4, 5\}$ in the right column. In the top left we show the initial certainty. In bottom row we compare the final certainty between active and random design. After round $r = 24$ the active method starts discovering the boundaries and puts swarm points in this region. After the final round the active method chooses more points around the class boundaries, with most of the points placed in the region where the three classes meet. The region where three class boundaries intersect is the region of least certainty.
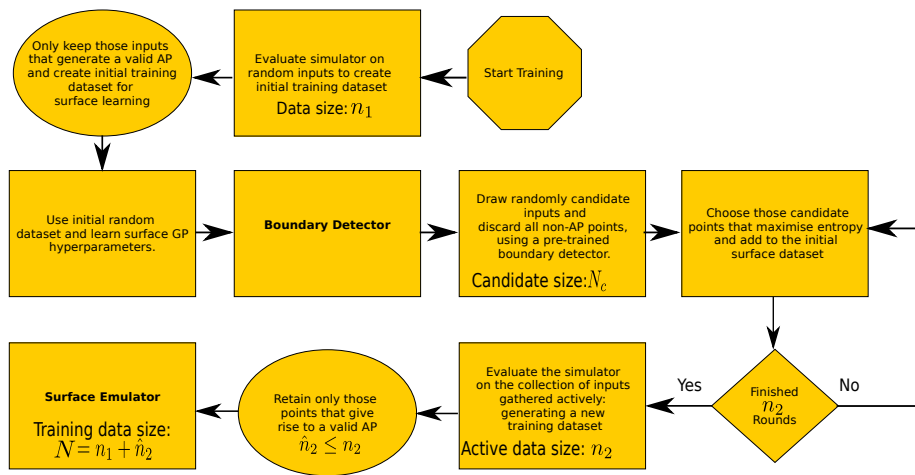
Figure 5: **Steps involved in active learning for the surface predictor**. We have used the oval shapes to indicate steps where we are removing non-AP points from the training set. The rest of the shapes used here follow standard flowchart notations. The algorithmic settings that needs to be supplied by the user are the initial data size n_1, the candidate set size N_c and the number of sequential rounds n_2.
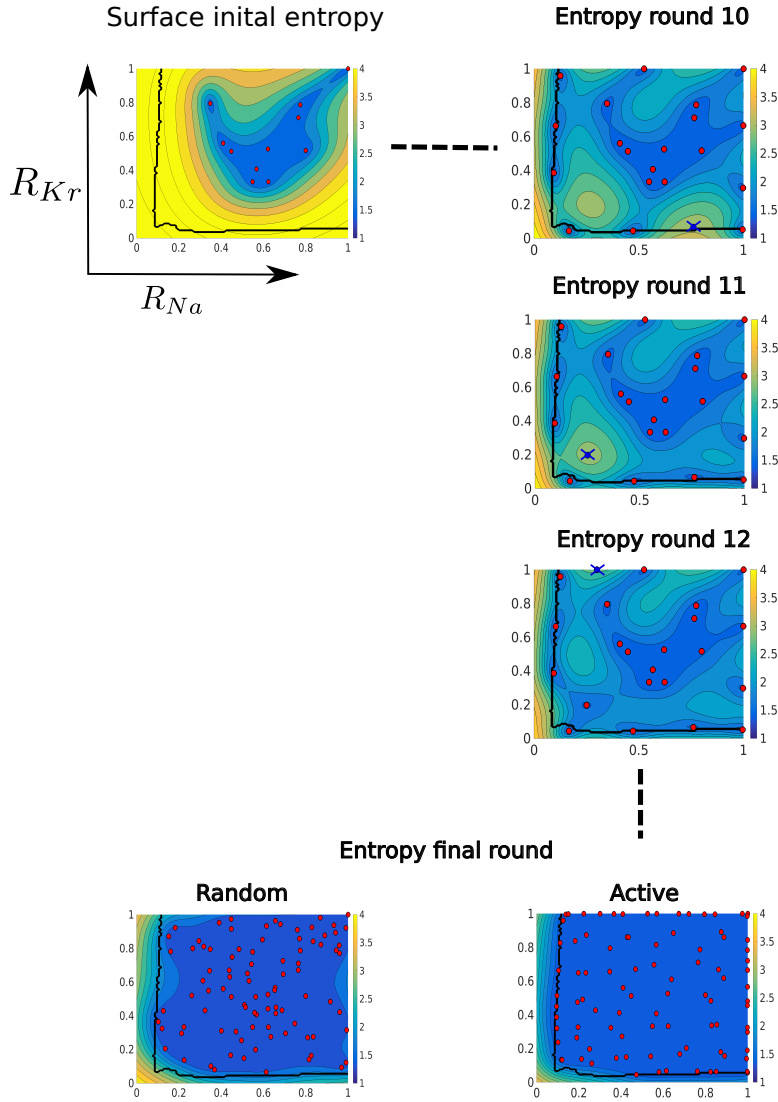
Figure 6: **Comparisons of entropies from active vs. random training of the surface GP**. The entropy defined by equation (44) is shown as a contour plot. A darker shade represents less entropy and thus less uncertainty in surface predictions. The black line demarcates the real boundary (obtained using a 100×100 grid for visualisation, and shown in Fig. 1(b)) between the non-AP and valid AP regions. The red circles show all the training points accumulated at specific stages of training in both active and random learning schemes. The active points picked during intermediate rounds are shown as blue circles with black cross-hair. We show three consecutive rounds of active learning $r = \{10, 11, 12\}$ in the right column. The next training point is always placed in the most uncertain (high entropy) region on the input space. In the top left we show the initial uncertainty on $n_1 = 11$ training points. In the bottom row we compare the final uncertainty between the active and random schemes after $n_2 = 90$ rounds. The active method is able to reduce the uncertainty noticeably over the input domain after the final round. Also note that in the intermediate rounds such as $r = \{10, 11, 12\}$ some active points are picked in the non-AP region. This happens due to misclassification of the candidate points. However, after we finish the active learning these points are discarded based on the actual simulator evaluation.

3 Results

We are going to test different aspects of the two-step emulator applied to predicting the APD response surface as obtained by simulating the O'Hara [12] cardiac electrophysiology model. We evaluate the performance of the emulator by observing the rate of decrease in prediction error, a *learning curve*, of the surface and boundary GP respectively as we grow the size of the training dataset. We have designed two sets of tests. In the first set, we evaluate the learning curve of the surface and classifier GP without active learning, which we denote as random learning. In the second, we evaluate similar learning curves using the active learning schemes described in sections 2.5.2 & 2.5.1.

Note that the error in the surface GP cannot be evaluated at points that belong to either of the non-repolarising or non-depolarising regions. To circumvent this, we only use test points that are associated with a valid AP to obtain learning curves (in both random and active learning experiments) for the surface GP, but we also track the misclassification rate of all points.

We also compared the learning curves in the first set of experiments, the random case, with the learning curves of a look-up table based interpolator used within [9] and subsequently in a web-based APD prediction application [38]. This interpolator uses a look up table and a space partitioning algorithm to interpolate a range of biomarkers including the APD. Following the emulator tests we only test the interpolator on inputs generating a valid AP while comparing the surface predictions. To test the classifier GP's performance we use inputs on the entire domain.

3.1 Simulator setup

As mentioned previously, our simulator is the ventricular action potential model from [12]. For the experiments below, the 4D input to the simulator is the conductance scalings of four ion currents:

i) Fast sodium channel conductance scaling — R_{Na};

ii) Rapid delayed rectifying potassium channel conductance scaling — R_{Kr};

iii) Slow delayed rectifying potassium channel conductance scaling — R_{Ks};

iv) L-type calcium channel conductance scaling — R_{CaL}.

The output of the simulator is the APD_{90} value under these channel scalings. For each evaluation of the simulator we pace the cardiac model with 100 1 Hz paces (using the stimulus defined in the CellML file/original [12] paper) in order to allow the state variables to settle towards their 1 Hz limit-cycle (a larger number may be required in practice). Thus for each input combination the simulator calls the underlying ODE solver 100 times. We have used the Chaste cardiac modelling package [39] to implement the simulator using a CellML representation [40] of the model. Code is openly available as described in Section 6.

24

3.2 Surface and classifier GP learning curves: random method

For this experiment we used the simulator (setup as described previously) to generate a training, $\mathcal{D}_{train}$, and a test, $\mathcal{D}_{test}$, dataset each containing a different $N = 100{,}000$ points. For evaluating the surface GP we keep only those points, in both the training and test datasets, that are associated with a valid AP. To obtain the learning curve we generate a random subsample $\mathcal{D}^*train = \{500, \dots, 50\,000\}$ of the training data to fit the hyperparameters by MLE for both the surface and boundary GPs and draw prediction on the entire 100,000 test points. Denoting the test dataset as $\mathcal{D}_{test} = \{(\boldsymbol{R}_i, y_i, k_i), i = 1, \dots, N\}$ we define the surface prediction error $\mathrm{E}_{surface}$ as the mean absolute error in the APD_{90} given by

$$\mathrm{E}_{surface} = \frac{1}{N_{\mathrm{AP}}} \sum_{i=1}^{N_{\mathrm{AP}}} |y_i - \mu(f_i)|, \tag{45}$$

where N_{AP} is the number of test points associated with a valid AP, $\mu(f_i)$ defines the posterior mean prediction of the surface emulator GP, and $|.|$ defines the absolute value. The error for the boundary detector $\mathrm{E}_{boundary}$ is defined as the percentage misclassification rate given by

$$\mathrm{E}_{boundary} = \frac{100}{N} \sum_{i=1}^{N} \mathbb{I}(k_i \neq \widehat{k}_i), \tag{46}$$

where $\widehat{k}_i$ is the most likely class prediction from the OVR classifier and $\mathbb{I}(\cdot)$ denotes the indicator function.

Also note that we start using sparse covariances, using the FITC approximation, when number of training points is greater than 10,000 and 5000 for the surface and classifier GPs respectively. We used the FITC approximation with 1000 and 300 training points (see section 1.4) for the surface and classifier GPs respectively. This number of training points is sufficient to produce an approximation comparable (and superior) to that of the Look-Up-Table interpolator.

We plot the learning curves for the surface and classifier GPs in Figs. 7 & 8. In both of these figures the black line distinguishes between the part of learning curves generated using the true and the FITC covariance.

Unlike the classifier GP learning curve (Fig. 8(b)) the surface GP error stops decreasing as the FITC approximation is introduced and this error remains relatively constant despite the increase in training data size (see Fig. 7(b)). The error generated by the surface GP with a training data size of 10,000 is the minimum error achievable, on the test dataset that we have used, and thus introducing more training points is futile in decreasing the error. In fact with the FITC covariance the error goes up due the introduction of the covariance approximation. The surface GP clearly outperforms the linear interpolator; this is expected as the interpolator employs a very simple function estimation compared to GP regression.
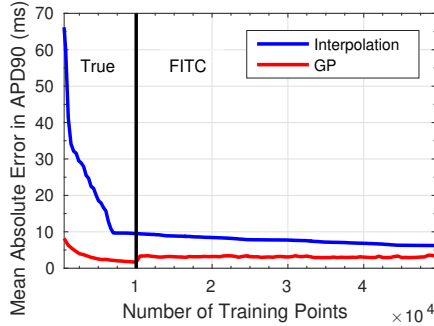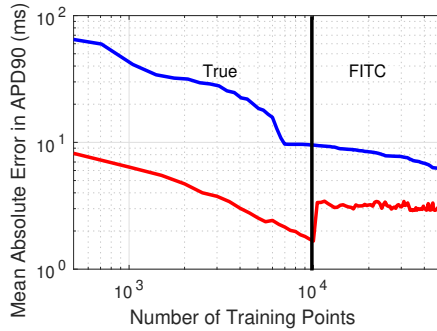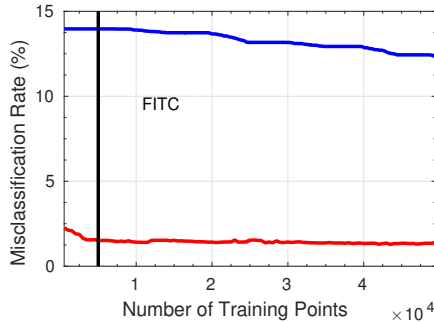
Figure 7: **Learning curves for the surface GP**. a) Original axes, b) Logarithmic axes. The black line distinguishes between the part of learning curves generated using the true covariances and the FITC approximations.

From Fig. 8(a) it is evident that the classifier error decreases steadily at a faster rate up to 4000 training points. Although the error keeps on decreasing beyond this training data size, the rate of decrease is reduced noticeably. Also notice in Fig. 8(b) that using a sparse covariance approximation does not appear to hinder the accuracy of the GP classifier.

3.3 Performance with active learning: surface emulation

In section 2.5.2 we gave an overview of the surface active learning scheme. Although entropy based active learning is a well studied method, doing so on a discontinuous surface brings about a new set of challenges so as to confine the learning scheme within the boundaries of discontinuities. This is achieved using the classifier. We start with an initial dataset $\mathcal{D}_\emptyset \in \mathcal{D}_{train}$ of size $n_1 = 500$, a random subsample of $\mathcal{D}_{train}$. We also fix the candidate set $\left\{ \boldsymbol{R}_{o_j} \right\}_{j=1,\dots,N_c}$ size as $N_c = 10,000$. For evaluation purposes we train (learn the hyperparameters)

26

(a)
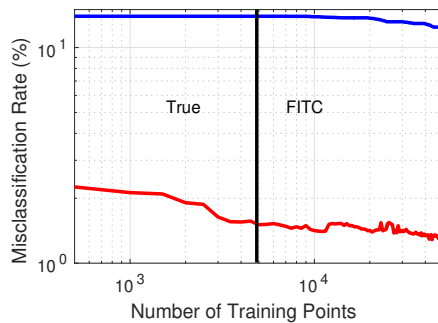
(b)

Figure 8: **Learning curves for the classifier GP**. a) Original axes, b) Logarithmic axes. The black line distinguishes between the part of learning curves generated using the true covariances and the FITC approximations.

of the surface and classifier GPs using $\mathcal{D}_\emptyset$. Note that for actual applications we will have an (actively) pre-trained classifier GP. We test this approach below in section 3.6. We also fix the number of active points to $n_2 = 2500$. With no misclassification the final training set size would be $N = n_1 + n_2 = 3000$. However, we end up with a training size $N_{AP} < N$ after discarding the inputs with invalid APs.

To generate a learning curve we calculate the surface error using equation (45) by sequentially generating predictions with the active dataset after the inclusion of 25 new training points. Note alternatively we can evaluate the learning curve after inclusion of every other active point. However, to keep parity with the classifier learning curves (presented in section 3.4), evaluated on each new swarm of active points, we adopt the aforementioned frequency of evaluation. The prediction is made on the test dataset $\mathcal{D}_{test}$, containing N_{AP} test points. To compare this with randomly adding training points to the initial set $\mathcal{D}_\emptyset$ we calculate the same sequential prediction errors by drawing 25 new

points randomly from $\mathcal{D}_{train}$, having only valid AP inputs, as opposed to actively learning them. We keep on adding these 25 random points for $n_2/25 = 100$ rounds. To highlight the variability in randomly growing the training dataset we repeat the random learning exercise 10 times. Furthermore, we draw the predictions in both schemes with the set of hyperparameters learnt using the same initial dataset $\mathcal{D}_\emptyset$. We plot the resulting learning curves in Fig. 9(a). From Fig. 9(a) it is evident that the active learning scheme outperforms the random update method. In Fig. 9(b) we plot the learning curves for the 2-input problem that we used for visualisations of the surface entropy in section 2.5.2. Only in this case we change the initial points size to $n_1 = 100$ to improve the hyperparameter estimate, and to compare with random design we keep adding 25 new randomly drawn points sequentially to the initial training set as done above for the 4-input problem. We generated a test dataset containing 10,000 points placed on a 2-dimensional grid for this 2-input case. We extended the rounds to 500 in order to have a smoother learning curve. The swarm size for evaluating error is maintained as 25. The increased accuracy in the 2-input case, for both active and random learning, results from the reduced dimensionality of the discontinuity surface. Also note that the error is reduced much faster during the earlier rounds of active learning than the later ones.

The entropy criteria will eventually start picking all those points from the candidate set that are located near the surface boundaries (regions of higher entropy including the class boundaries), once a certain number of points are picked in the valid AP region to reduce the entropy significantly. Although it is difficult to quantify the number of points (in the valid AP region) required to reduce entropy optimally we start to see this effect for the simpler 2-input problem in Fig. 9(b). Once the number of active training points is greater than 400 the learning curve improvement slows. At this point the active learning scheme is adding points mostly around the surface boundaries, whereas for the random addition of training points more points are accumulated in the valid AP region and thus the random error keeps on decreasing. It is worth noting that in applications such as ours some of the boundary regions (no ion channel block in certain dimensions) are of particular interest, and extra accuracy there is beneficial. Due to the nature of the entropy criteria we expect the same thing to happen in the 4-input case but for a higher number of training points.

3.4 Evaluating classifier active learning

We adopt the same procedure for testing the classifier active learning as for the surface. That is, we try to compare the misclassification error produced by actively growing the training dataset to that of a random design.

We start with an initial dataset $\mathcal{D}_\emptyset \in \mathcal{D}_{train}$ of size $n_1 = 500$ and sequentially grow the training dataset using PSO as described in section 2.5.1 to collect a swarm of inputs of size $n_s = 50$ in each round. We retain the same test dataset and we repeat this process for a 2-input problem set up exactly as in the surface active learning experiment. The random learning is carried out as previously, using 10 randomly drawn swarms in each round from $\mathcal{D}_{train}$. For the 4-input

problem we perform $r = 20$ rounds of PSO, collecting $n_3 = 20 \times 50 = 1000$ active input points, and thus $N = n_1 + n_3 = 1500$ training points in total. For the 2-input case we start with $n_1 = 100$ initial points and restrict the swarm size to 25 points, thus collecting $n_3 = 25 \times 20 = 500$ active input points and $N = n_1 + n_3 = 600$ training points in total. We retain the same test dataset which was used for evaluating the surface active learning in the 2-input case. We use a cut-off value $\theta = 0.5$ specified by the average certainty of the swarm particles to stop the PSO iterations.

We plot the learning curves for the boundary classifier in Fig. 10(a). The learning curve for the 2-input problem is shown in Fig. 10(b). We have created an animation (see Supplementary Material) visualising the first 8 rounds of PSO and subsequent changes to contour plots of the certainty. In both these plots we see that the active error decreases much more rapidly than in the surface case. Furthermore, we notice that the variability (among the 10 repetitions) of random errors for the classifier GP is higher than the surface error plots.

However, for the 2-input case we see the same flattening of the active learning curve as in the surface active learning. This is because after going through 20 rounds the active learning scheme manages to put the necessary amount of input points covering all the uncertain regions on the input space. Adding further inputs does not affect the overall uncertainty noticeably.

3.5 Learning times and swarm sizes

While designing the emulator we have to consider the fact that training and prediction of GPs are limited due to the computationally expensive covariance inversion step. For the classifier GP both training and prediction involve a number of such covariance inversion steps due to the expectation propagation algorithm. In Figs. 11 & 12 we show the training and prediction times for random learning in the case of the surface and classifier GP respectively. The corresponding learning curves are presented in Figs. 7 & 8. We have also plotted the simulation time (shown as a solid blue line) required to evaluate the corresponding number of training points for each of the GPs. The simulation time for evaluating all the points in the test set is shown as a horizontal broken line in all the plots. The total time for training the GP and simulating the training set is also plotted in Figs. 11 & 12.

For the surface GP the total training plus prediction time is negligible in comparison to the simulation time for the entire test set, which is an expected result. Furthermore, using a FITC covariance we see further speed-up albeit at the cost of reduced accuracy (see Fig. 7(b) for the corresponding learning curve). However, for the classifier GP interestingly we see that while using a true covariance the training time exceeds the simulation time for evaluating the entire test set (mainly due to the PSO rounds). Even for the FITC covariance the training and prediction times are significantly higher than that of the surface GP. Thus in order to make our two-step approach work in a practical manner we need to use a small training dataset for the classifier GP. This is possible using the active learning scheme as we can use fewer points than the corresponding

random learning.

With a small number of training points we can reduce the training time of a classifier GP. But since each PSO round incurs many classifier predictions involving the EP algorithm it is important to find out the time spent in carrying out the classifier active learning, especially within the PSO iterations. In Fig. 13 we plot the reduction in misclassification for the classifier active learning, averaged over 10 repetitions, with increasing computation time. The computation time is represented as the cumulative training time which is the sum total of the training time upto the r-th round. We carried out the training for 30 rounds with a swarm size of $n_s = 50$. We also plot the same graph for the FITC covariance (red line). For each of the 10 repetitions we used a separate initial set of $n_1 = 500$ random training points to learn the true and FITC covariance hyperparameters.

It is evident from Fig. 13 that initially both using a FITC as well as the true covariance the average error is reduced almost at the same rate. However, we get marginally higher reduction using the FITC covariance.

Another important algorithmic setting that has a potential impact on the run-time of the classifier active learning is the swarm size. A smaller swarm of points has to go through many more rounds of PSO in comparison to a larger swarm to collect the same amount of actively-learnt training points. However, more PSO rounds will enable the active learning scheme to hone in on different uncertain regions of the input space.

In order to test this we ran the FITC active learning, with three different swarm sizes $\{1000, 500, 100\}$. We chose the initial dataset size n_1 to be 1000 points. We carried out active learning for each of the swarm sizes to collect 5000 active points. Hence, the GP with a swarm size of 100 finished $l = 50$ rounds while the GP with swarm size of 1000 finished only 5 rounds. In Fig. 14 we plot the misclassification rate for the three GPs concerned. The smallest swarm has the largest run-time but is able to reduce the error much more than the larger swarm variants. In fact the smallest swarm variant reduces the classifier error significantly more within 1 hour (this is demarcated with the black line) than the larger variants. The reason for this is that the smallest swarm can complete more rounds of PSO within the same time compared to others and thus can put points across more distinct uncertain regions than the larger swarms do. Covering more regions of uncertainty (with fewer points) is more important than covering fewer regions with more points.

3.6   Performance of the two-step method

Finally to test the performance of the complete two-step emulator we carry out active learning for both the classifier and surface GPs. For the classifier GP we used a swarm of $n_s = 50$ points with $r = 30$ rounds of PSO generating the active set of points $n_3 = r \times n_s = 1500$. We retain the same PSO threshold of $\theta = 0.5$. We used the FITC covariance for the classifier GP.

Subsequently for the surface GP we carry out active learning to gather $n_2 = 3000$ points sequentially. For the surface active learning we have used a candidate set of 10,000 input points and used the actively learnt classifier to

filter out all the non-AP points. We chose an initial dataset $\mathcal{D}_\emptyset$ of size $n_1 = 500$, drawn randomly, to learn the classifier and surface GP hyperparameters. Thus, the total training data size for the two-step emulator is $n_1 + \hat{n}_2 + n_3 = 5000$. As we filter out the non-AP points from the active surface training inputs our actual training size is slightly less than 5000.

To carry out a fair comparison between active and random learning the hyperparameters were learnt using $\mathcal{D}_\emptyset$ and were not updated after finishing all the rounds. However, while evaluating the two-step emulator we have updated the hyperparameters after finishing the active learning. This update step affects the run-time of the two-step method and one can optionally avoid this step.

In Table 1 we present the performance of the two-step GP which we denote as **Two-step** including the parameter re-learning/update step. We have also presented the performance of our previously described linear interpolator [9] denoted as the **Interpolator**, here trained using 5000 points. We also show the time required for both the GP based and interpolation method to complete the emulation task, which we denote as prediction time, and simulation (Chaste ODE evaluations) for the entire test dataset. In Table 2 we breakdown the run-time of the emulator into training and prediction times for both the surface and boundary GP. In Table 2 the training time of classifier consists of the sum total of hyperparameter learning on $n_1 = 500$ initial points, active learning to collect $n_3 = 1500$ points and re-learning hyperparametrs on 2000 points. For the surface GP the total training time consists of hyperparameter learning on valid AP points in the initial set of $n_1 = 500$, active learning to collect $n_2 = 3000$ points and re-learning hyperparametrs on the AP ones out of 3500 points. Note that the surface and classifier GP training times include the time for ODE simulation during the respective active learning processes. The prediction is drawn on 100,000 test points.

Table 1: **Performance of two-step emulator training for a 4-input problem**. We evaluated emulator performance by comparing against a test dataset of 100,000 points, picked from the 4D input space at random. Simulating the full ODE solutions for these points took 44.2 hours. The training times for both the methods are listed below and include walltime for ODE simulation at the training points. These performance tests were carried out on single processor (3GHz with 32GB RAM). Although the two-step emulator's training time is double that of the interpolator it outperforms the interpolator in terms of prediction accuracies.

Method	Training size	Training time (h)	Prediction time (s)	$E_{boundary}$ (%)	$E_{surface}$ (ms)
Two-step-GP	5000	5.5085	68.8980	1.5770	2.8742
Interpolator	5000	2.4320	1.8951	13.9670	17.9525

Although the new two-step GP emulator is slower than the Look Up Table-based interpolator, considering the time needed in simulating the entire test dataset it is a reasonable alternative to the interpolator due to its improved accuracy in the presence of discontinuities. Also note that once the emulator is trained we can use it to evaluate the response surface for a large number of

Table 2: **Breakdown of run-time of the two-step emulator**. Here the training time of the classifier and surface GPs constitute the sum total of hyperparameter learning on the initial and final training set as well as the active learning. The total number of training points for the classifier is $n_1 + (n_s \times r) = 500 + (30 \times 50) = 2000$ and for the surface is $n_1 + n_2 = 500 + 3000 = 3500$. Predictions are made for a test-set of 100,000 points.

Method	Training time (h)	Prediction time (minutes)
Classifier GP	4.2939	0.6320
Surface GP	1.2146	0.5163

inputs repeatedly. In such a scenario the fast prediction time, as seen in Table 2, is extremely valuable for an uncertainty propagation task.

3.7 Drug Block Case Study

Finally, we evaluate the performance of the two-step emulator within a study of drug action. This is relevant to the work being undertaken in the Comprehensive in-vitro Pro-arryhthmia Assay initative [41] and will make the Uncertainty Quantification undertaken there [42] faster to perform. Recently, Crumb *et al.* [43] published dose-response screening data for 30 compounds on 7 different ion channels along with point estimates of $[IC50]$ and the Hill coefficient n. Johnstone *et al.* [44] then implemented a method to derive a probability distribution for the drug block parameters, as given by equation (1), on various ion-channels. To propagate the uncertainty, as captured through the marginal posterior distributions of these parameters, APD_{90} values were simulated using a Monte Carlo method for the corresponding samples of $[IC50]$ and n. We will test our emulator in the same setting, to establish whether it can provide the same insights in a more computationally tractable fashion.

At a high concentration ($[C]$ in equation (1)) of quinidine, the reduction in hERG ion channel conductance pushes the model into the non-repolarising region. Thus, to evaluate our proposed emulator near to the discontinuous regime we repeat the uncertainty quantification task while blocking the hERG channel based on quinidine as the chosen drug. We refer the reader to Johnstone *et al.* [44] (section 4 in particular) for further details of the characterisation of input uncertainties, and have generated samples of the $[IC50]$ and Hill coefficient n using the technique and code they provided. We used the hierarchical model variant from that paper and generated 2000 samples inferring the underlying drug effect (rather than including a prediction of future experiment-level variability in our samples) using the concentration effect curve given by equation (1) for i) a moderate dose — $0.3\,\mu$M of quinidine producing $\approx 50\%$ block and ii) a higher dose — $3\,\mu$M of quinidine producing $\approx 85\%$ block, obtaining distributions of conductance scalings $\boldsymbol{R}_{Kr}$ shown in Fig. 15. We picked these concentrations to test the emulator for uncertainty quantification on a distribution of APD_{90} that

is i) entirely on the surface and ii) straddling a discontinuity. Our distributions shown in Fig. 15 are analogous to those shown in the original publication [44] (Fig. 12E in that paper).

Since in many drug action studies we would be testing a single ion channel, the conductance for that channel would be blocked, whereas for other channels the conductances would be set to the maximal conductance with no blocking. To account for this scenario we augment the initial random dataset $\mathcal{D}_\emptyset$ with 2^4 training points which have the scalings $\boldsymbol{R}$ set to 0/1, and use a full 4D emulator. With this new augmented $\mathcal{D}_\emptyset$ we repeat the active learning for both the classifier and surface GP as described in the previous section. The total training set size for the two-step emulator in this case is $n_1 + 2^4 + n_2 + n_3 = 5016$. Addition of these corner points does not change the prediction accuracy noticeably for a general case where we draw test points which have some amount of scaling applied to each channel such as our test dataset $\mathcal{D}_{test}$ used so far for testing the two-step emulator in the previous sections. Drawing predictions on $\mathcal{D}_{test}$ we found the surface error $E_{surface} = 3.0597$ ms and the classifier error $E_{boundary} = 1.5580\%$. We notice little difference between these error values and the ones reported in Table 1.

We perform a prediction for all the 2000 samples of $\boldsymbol{R}_{Kr}$ at each of the two concentrations using the emulator to obtain estimates of the APD$_{90}$ surface using equation (39). To facilitate the visualisation of the classifier and surface GP predictions we plot one-dimensional slices, in Figure 16, of the pre-trained classifier posterior probabilities and the surface GP mean and variance for an artificial test dataset with 1000 samples of $\boldsymbol{R}_{Kr}$ spread evenly between 0 and 1. The posterior probabilities consist of the probabilities π^k of the binary GP classifiers for all the three regions. Note that for visualisation we are passing all the artificial test points to the surface GP, unlike the two step method where we pass only those which are classified as valid AP points to the surface GP. We set the scaling of other channels as: $(\boldsymbol{R}_{Na}, \boldsymbol{R}_{Ks}, \boldsymbol{R}_{CaL}) = 1$, to represent no block at those channels. Whilst this means a 1D emulator could be used, we wish to test our more general 4D emulator in what follows.

In previous sections we considered the posterior mean of the surface GP at test points to define APD$_{90}$ surface. However, in a real application such as this drug action study we also want to include uncertainty due to the emulator, as estimated by its posterior variance; and we wish to propagate this uncertainty into the corresponding APD$_{90}$ distribution too.

Thus to calculate the full uncertainty we simply sum the continuous Gaussian distributions for APD given by the emulator at each discrete sample of block, and re-normalise to produce a full probability distribution for APD. We denote the number of test points classified as being in the 'full AP region' as N^{AP}, so

$$p_{APD_{90}} = \frac{1}{N^{AP}} \sum_{i=1}^{N^{AP}} \mathcal{N}\big(\mu(f_i), Var(f_i)\big), \tag{47}$$

where $\mu(f_i)$ (equation (36)) and $Var(f_i)$ (equation (37)) are the posterior mean and variance of the surface GP at the i-th (out of N^{AP}) test point. Computa-

33

tionally, we discretise the above continuous equation into 1000 values of APD_{90} between 0 and $1000\,\mathrm{ms}$ (the bounds for a $1\,\mathrm{Hz}$ simulation).

Fig. 17 shows the distribution of APD_{90} from taking samples from Fig. 15 and mapping them through the classifier and Surface GP shown in Fig. 16 and finally evaluating equation (47). We also evaluate $p_{APD_{90}}$ using the ODE simulator directly on all samples of block for comparison.

For both moderate and high dose cases we find that the distributions of APD_{90} obtained from the emulator and simulator match well. For the moderate dose there is no misclassification. For the high dose 548 out of the 2000 $\boldsymbol{R}_{Kr}$ samples belong to the non-repolarising region when running the full ODE system. The emulator assigns 477 points to this region and thus misclassifies 71 points. The slight increase in misclassification rate 3.55% from what is reported in Table 1 for predictions on $\mathcal{D}_{test}$ happens due to the presence of most of the test samples in this UQ task being near the class boundary.

We notice from Fig. 16 that the binary classifier probabilities for the AP and the non-repolarising regions are almost the same near the class boundary; as a result the certainty is very low in this region. If using the emulator for safety critical applications (such as the high dose quinidine action study considered here) if some test input points are located right along the discontinuities then one switch to performing full simulation simulation for points where the certainty c is less than a threshold, say 0.8.

4   Discussion

Uncertainty and variability is intrinsic to a plethora of biological processes that we want to understand, model and predict. In cardiac modelling, sources of uncertainty stem from the experimental 'error' in the measurements from our protocols, lack of knowledge about the underlying mechanisms leading to 'structural error' in our models, variability due to differences in cell and ion channel states due to cells being in different settings and gene expression patterns, and variability due to the inherent stochasticity of some of these processes exhibited at multiple time and spatial scales [18]. To accommodate mathematical/phenomenological models in safety-critical clinical practice and drug development, it is therefore of utmost importance to quantify and propagate these uncertainties to model predictions. As a consequence we need to examine our model predictions over large parameter domains. This is where emulation becomes necessary to reduce the computational burden associated with uncertainty quantification and propagation. However, many mathematical models, especially in cardiac electrophysiology, have bifurcations in behaviour as we move through parameter domains, rendering traditional Gaussian Process emulation infeasible. In this work we have addressed this specific issue of emulating cardiac action potential models having bifurcations in dynamics, and as a result, discontinuous output/response surfaces. We proposed a two-step emulator combining GP classification and regression to emulate the discontinuous action potential duration biomarker response surface and applied our method to a study of drug action.

34

Looking at the computational complexity of the GP classifier it is natural to ask whether a simpler classifier could be used for boundary detection. To achieve a good degree of separation we have to use many more training points with a simple classifier such as a linear softmax classifier (this classifier can generate probabilistic predictions). Furthermore, using the probabilistic framework of the GP classifier we can quantify the uncertainty of the boundary locations, somewhat accounting for the fact that numerical errors become important close to the bifurcation point and so the notion of a somewhat random and probabilistic outcome here is helpful even though the ODE system is completely deterministic in this case. We use this probabilistic property to build an active learning scheme which reduces the necessary training dataset size significantly. Using a complex boundary detector we become less sensitive to simulation errors and are able to use the simulator more sparsely.

We tried an alternative GP classifier using the Laplace approximation [27] which reduced the training and prediction times dramatically (results not shown). However, there was a significant drop in accuracy, and so we retained use of the EP algorithm.

If there is sufficient computing power available, the classifier certainty metric could be used to confine use of the emulator to locations that we are confident in the class. Thus, we can associate a threshold, say 0.9, and if for a particular test point the certainty is below the threshold then we can use simulation for finding the true output value. We did not explore an adaptive train-use-refine scheme here, but it would then be intuitive to add the simulation points to the GP training sets.

We have so far not discussed the timing implications of the surface active learning. This is because the surface active learning took less than 30 minutes (for $n_2 = 3000$) to finish. This speed can be further reduced by using the FITC approximation, but we recommend the usage of true covariance as the run-time is insignificant in comparison to classifier active learning.

In this paper we have confined our emulation to one biomarker: the APD. However, the effect of bifurcation is observed in many other summary statistics of the action potential too, and the techniques are transferable. Using multiple GPs one can build a two-step emulator that covers the output space consisting of all the relevant biomarkers of the action potential trace. Or, using a multi-output GP [45, 46], correlations among the generated biomarkers can also be captured to enhance prediction accuracies.

5   Conclusion

In this paper we presented a two-step emulator of the discontinuous APD_{90} biomarker response surface generated by the O'Hara cardiac AP model under varying drug block. The proposed emulator produces good prediction accuracies on an artificial test dataset containing 100,000 test points. Our two-step emulator was trained at a fraction ($\approx 10\%$) of the computational expense of simulation. The proposed emulator requires ≈ 1 minute for drawing predictions on the entire

test dataset. We achieve this by using sparse GP approximation (FITC) in conjunction with a novel active learning scheme. A significant amount of the emulation time is consumed by the classifier GP due to its inherent computational limitations stemming from repeated covariance inversion within the EP algorithm.

We have applied our two-step method for uncertainty quantification in a drug action study. In this application we found the two-step method was able to perform uncertainty quantification of APD_{90} with high prediction accuracy. Our proposed method can be easily extended to accommodate other biophysical models (that go through bifurcations) and biomarkers.

6 Materials and Methods

All the codes we used to generate the results in this study are available to download from `https://github.com/sanmitraghosh/ApPredict_GP`.

7 Acknowledgements

References

1. P. Pathmanathan, M. S. Shotwell, D. J. Gavaghan, J. M. Cordeiro, and R. A. Gray, "Uncertainty quantification of fast sodium current steady-state inactivation for multi-scale models of cardiac electrophysiology," *Progress in Biophysics and Molecular Biology*, vol. 117, no. 1, pp. 4 – 18, 2015. Multi-scale Systems Biology.

2. J. Oakley and A. O'Hagan, "Bayesian inference for the uncertainty distribution of computer model outputs," *Biometrika*, vol. 89, no. 4, pp. 769–784, 2002.

3. J. Oakley, *Bayesian uncertainty analysis for complex computer codes.* PhD thesis, University of Sheffield, 1999.

4. M. R. Davies, H. B. Mistry, L. Hussein, C. E. Pollard, J.-P. Valentin, J. Swinton, and N. Abi-Gerges, "An in silico canine cardiac midmyocardial action potential duration model as a tool for early drug safety assessment," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 302, no. 7, pp. H1466–H1480, 2012.

5. G. R. Mirams, M. R. Davies, Y. Cui, P. Kohl, and D. Noble, "Application of cardiac electrophysiology simulations to pro-arrhythmic safety testing," *British Journal of Pharmacology*, vol. 167, no. 5, pp. 932–945, 2012.

6. M. R. Davies, K. Wang, G. R. Mirams, A. Caruso, D. Noble, A. Walz, T. Lavé, F. Schuler, T. Singer, and L. Polonchuk, "Recent developments in using mechanistic cardiac modelling for drug safety evaluation," *Drug Discovery Today*, vol. 21, no. 6, pp. 924–938, 2016.

7. T. Brennan, M. Fink, and B. Rodriguez, "Multiscale modelling of drug-induced effects on cardiac electrophysiological activity," *European Journal of Pharmaceutical Sciences*, vol. 36, no. 1, pp. 62–77, 2009.

8. R. C. Elkins, M. R. Davies, S. J. Brough, D. J. Gavaghan, Y. Cui, N. Abi-Gerges, and G. R. Mirams, "Variability in high-throughput ion-channel screening data and consequences for cardiac safety assessment," *Journal of Pharmacological and Toxicological Methods*, vol. 68, no. 1, pp. 112–122, 2013.

9. G. R. Mirams, M. R. Davies, S. J. Brough, M. H. Bridgland-Taylor, Y. Cui, D. J. Gavaghan, and N. Abi-Gerges, "Prediction of thorough QT study results using action potential simulations based on ion channel screens," *Journal of Pharmacological and Toxicological Methods*, vol. 70, no. 3, pp. 246–254, 2014.

10. E. T. Chang, M. Strong, and R. H. Clayton, "Bayesian sensitivity analysis of a cardiac cell model using a gaussian process emulator," *PloS one*, vol. 10, no. 6, p. e0130252, 2015.

11. C.-h. Luo and Y. Rudy, "A dynamic model of the cardiac ventricular action potential. i. simulations of ionic currents and concentration changes.," *Circulation research*, vol. 74, no. 6, pp. 1071–1096, 1994.

12. T. O'Hara, L. Virág, A. Varró, and Y. Rudy, "Simulation of the undiseased human cardiac ventricular action potential: model formulation and experimental validation," *PLoS Comput Biol*, vol. 7, no. 5, p. e1002061, 2011.

13. R. H. Johnstone, E. T. Chang, R. Bardenet, T. P. de Boer, D. J. Gavaghan, P. Pathmanathan, R. H. Clayton, and G. R. Mirams, "Uncertainty and variability in models of the cardiac action potential: Can we build trustworthy models?," *Journal of Molecular and Cellular Cardiology*, vol. 96, pp. 49–62, jul 2016.

14. B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, "Taking the human out of the loop: A review of bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2016.

15. J. Snoek, K. Swersky, R. S. Zemel, and R. P. Adams, "Input warping for bayesian optimization of non-stationary functions.," in *Proceedings of the 31st International Conference on International Conference on Machine Learning*, pp. 1674–1682, 2014.

16. R. B. Gramacy and H. K. H. Lee, "Bayesian treed gaussian process models with an application to computer modeling," *Journal of the American Statistical Association*, vol. 103, no. 483, pp. 1119–1130, 2008.

17. J.-A. M. Assael, Z. Wang, B. Shahriari, and N. de Freitas, "Heteroscedastic treed bayesian optimisation," *arXiv preprint arXiv:1410.7172*, 2014.

18. G. R. Mirams, P. Pathmanathan, R. A. Gray, P. Challenor, and R. H. Clayton, "Uncertainty and variability in computational and mathematical models of cardiac physiology," *The Journal of Physiology*, vol. 594, pp. 6833–6847, dec 2016.

19. A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies," *Journal of Machine Learning Research*, vol. 9, pp. 235–284, 2008.

20. C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning.* MIT Press, 2006.

21. D. Barber and C. K. I. Williams, "Gaussian processes for bayesian classification via hybrid monte carlo," in *Advances in Neural Information Processing Systems 9* (M. C. Mozer, M. I. Jordan, and T. Petsche, eds.), pp. 340–346, Neural Information Processing Systems Foundation, Inc., 1997.

22. T. P. Minka, "Expectation propagation for approximate bayesian inference," in *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pp. 362–369, 2001.

23. H. Nickisch and C. E. Rasmussen, "Approximations for binary gaussian process classification," *Journal of Machine Learning Research*, vol. 9, no. Oct, pp. 2035–2078, 2008.

24. J. Quiñonero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate gaussian process regression," *Journal of Machine Learning Research*, vol. 6, no. Dec, pp. 1939–1959, 2005.

25. E. Snelson and Z. Ghahramani, "Sparse gaussian processes using pseudo-inputs," in *Advances in neural information processing systems*, pp. 1257–1264, 2006.

26. A. Naish-Guzman and S. Holden, "The generalized fitc approximation," in *Advances in Neural Information Processing Systems*, pp. 1057–1064, 2008.

27. C. E. Rasmussen and H. Nickisch, "Gaussian processes for machine learning (gpml) toolbox," *The Journal of Machine Learning Research*, vol. 11, pp. 3011–3015, 2010.

28. A. Krause and C. Guestrin, "Nonmyopic active learning of gaussian processes: an exploration-exploitation approach," in *Proceedings of the 24th international conference on Machine learning*, pp. 449–456, 2007.

29. N. Houlsby, F. Huszar, Z. Ghahramani, and J. M. Hernández-Lobato, "Collaborative gaussian processes for preference learning," in *Advances in Neural Information Processing Systems*, pp. 2096–2104, 2012.

30. R. C. Eberhart and J. Kennedy, "Particle swarm optimization," in *Proceedings of the IEEE international conference on neural networks*, pp. 1942–1948, 1995.

31. Y. Shi and R. C. Eberhart, "Empirical study of particle swarm optimization," in *Proceedings of the 1999 Congress on Evolutionary Computation-CEC99*, vol. 3, pp. 1945–1950, 1999.

32. D. J. MacKay, "Information-based objective functions for active data selection," *Neural computation*, vol. 4, no. 4, pp. 590–604, 1992.

33. S. Seo, M. Wallat, T. Graepel, and K. Obermayer, "Gaussian process regression: Active data selection and test point rejection," in *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, vol. 3, pp. 241–246, 2000.

34. M. C. Shewry and H. P. Wynn, "Maximum entropy sampling," *Journal of applied statistics*, vol. 14, no. 2, pp. 165–170, 1987.

35. A. Gorodetsky and Y. Marzouk, "Mercer kernels and integrated variance experimental design: Connections between gaussian process regression and polynomial approximation," *SIAM/ASA Journal on Uncertainty Quantification*, vol. 4, no. 1, pp. 796–828, 2016.

36. C.-W. Ko, J. Lee, and M. Queyranne, "An exact algorithm for maximum entropy sampling," *Operations Research*, vol. 43, no. 4, pp. 684–691, 1995.

37. T. N. Hoang, K. H. Low, P. Jaillet, and M. Kankanhalli, "Nonmyopic ϵ bayes-optimal active learning of gaussian processes," in *Proceedings of the 31st International Conference on Machine Learning*, pp. II–739–II–747, 2014.

38. G. Williams and G. R. Mirams, "A web portal for in-silico action potential predictions," *Journal of pharmacological and toxicological methods*, vol. 75, pp. 10–16, 2015.

39. G. R. Mirams, C. J. Arthurs, M. O. Bernabeu, R. Bordas, J. Cooper, A. Corrias, Y. Davit, S.-J. Dunn, A. G. Fletcher, D. G. Harvey, *et al.*, "Chaste: an open source c++ library for computational physiology and biology," *PLoS computational biology*, vol. 9, no. 3, p. e1002970, 2013.

40. J. Cooper, R. J. Spiteri, and G. R. Mirams, "Cellular cardiac electrophysiology modeling with Chaste and CellML," *Frontiers in Physiology*, vol. 5, p. 511, jan 2015.

41. B. Fermini, J. C. Hancox, N. Abi-Gerges, M. Bridgland-Taylor, K. W. Chaudhary, T. Colatsky, K. Correll, W. Crumb, B. Damiano, G. Erdemli, *et al.*, "A new perspective in the field of cardiac safety testing through the comprehensive in vitro proarrhythmia assay paradigm," *Journal of biomolecular screening*, vol. 21, no. 1, pp. 1–11, 2016.

42. K. C. Chang, S. Dutta, G. R. Mirams, K. A. Beattie, J. Sheng, P. N. Tran, M. Wu, W. W. Wu, T. Colatsky, D. G. Strauss, *et al.*, "Uncertainty quantification reveals the importance of data variability and experimental design considerations for in silico proarrhythmia risk assessment," *Frontiers in physiology*, vol. 8, p. 917, 2017.

43. W. J. Crumb, J. Vicente, L. Johannesen, and D. G. Strauss, "An evaluation of 30 clinical drugs against the comprehensive in vitro proarrhythmia assay (cipa) proposed ion channel panel," *Journal of Pharmacological and Toxicological Methods*, vol. 81, pp. 251 – 262, 2016. Focused Issue on Safety Pharmacology.

44. R. Johnstone, R. Bardenet, D. Gavaghan, and G. Mirams, "Hierarchical bayesian inference for ion channel screening dose-response data [version 2]," *Wellcome Open Research*, vol. 1, no. 6, 2017.

45. E. V. Bonilla, K. M. Chai, and C. Williams, "Multi-task gaussian process prediction," in *Advances in neural information processing systems*, pp. 153–160, 2007.

46. R. Durichen, M. A. Pimentel, L. Clifton, A. Schweikard, and D. A. Clifton, "Multitask gaussian processes for multivariate physiological time-series analysis," *IEEE Transactions on Biomedical Engineering*, vol. 62, no. 1, pp. 314–322, 2015.
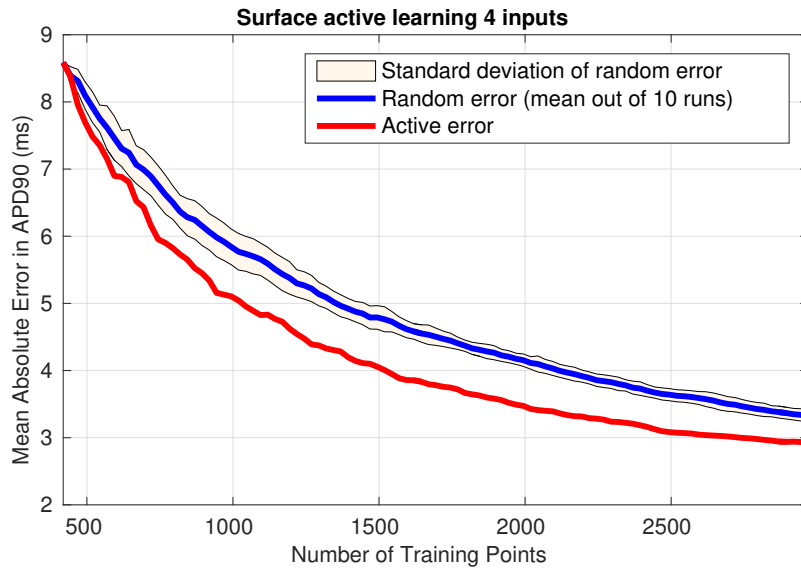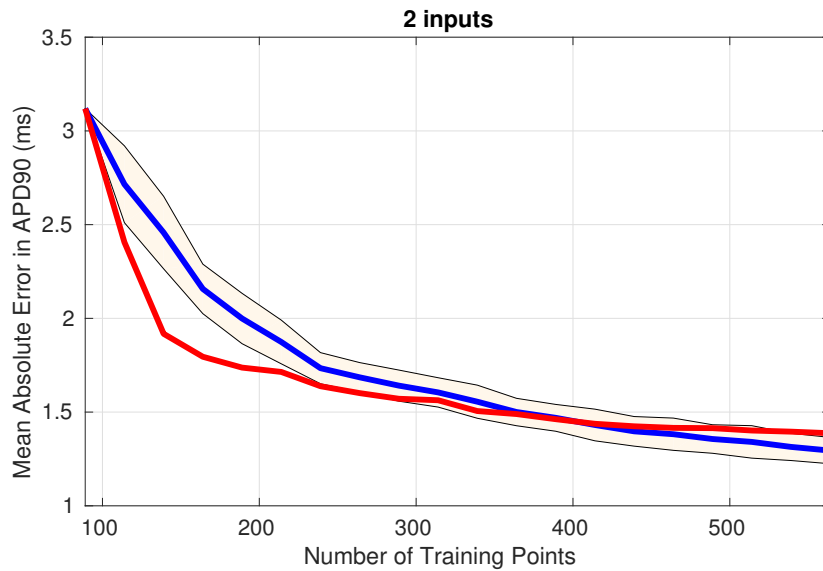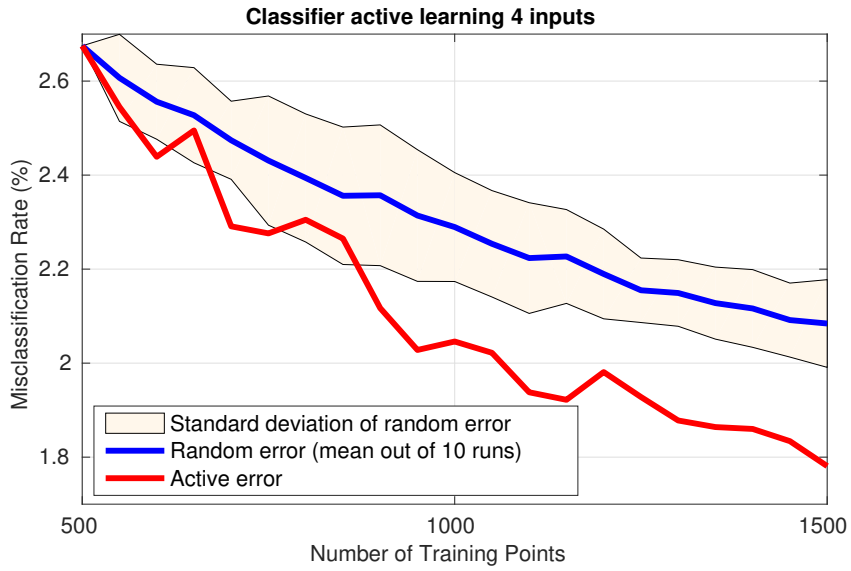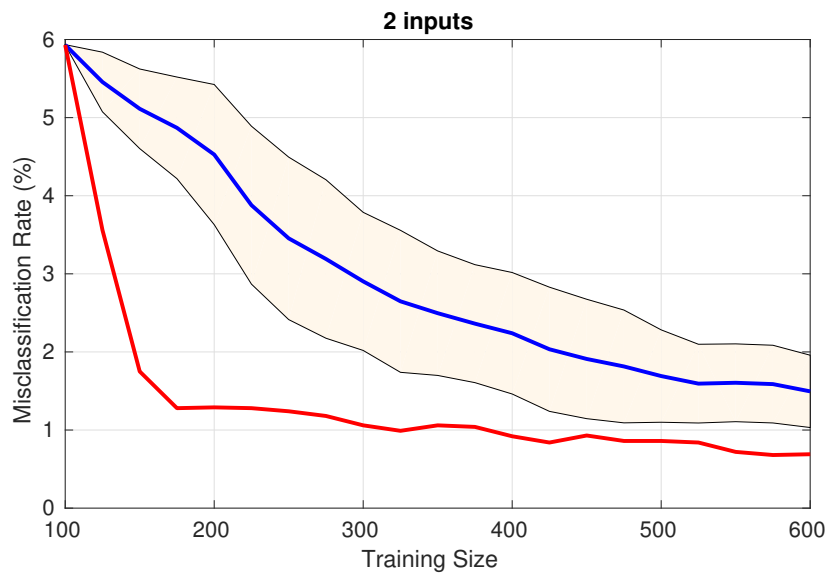
Figure 9: **Learning curves for surface active vs. random learning for the a)** 4-**input** & **b)** 2-**input problem**. The red line shows the error for training the GP using actively learnt inputs. The blue line shows the average, out of 10 repetitions, GP error for randomly drawing training inputs. The shaded area shows the standard deviation of this error, reflecting the variability in performance if a random design is carried out.
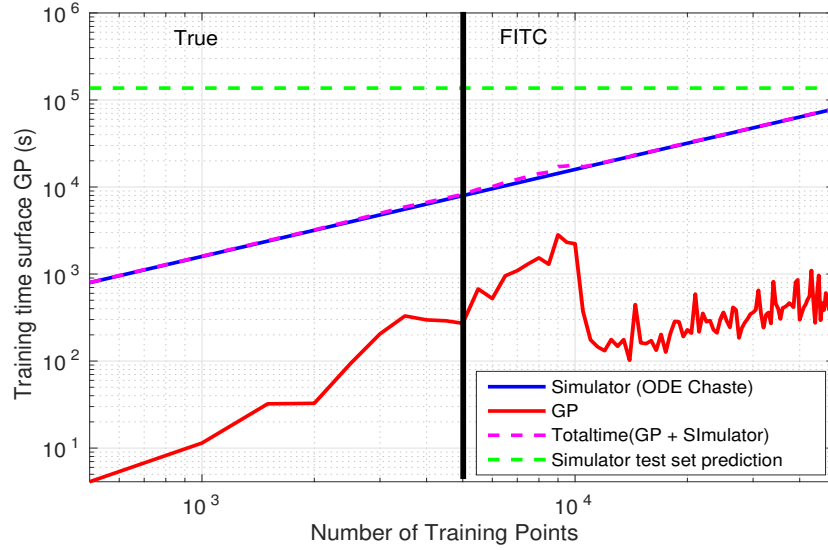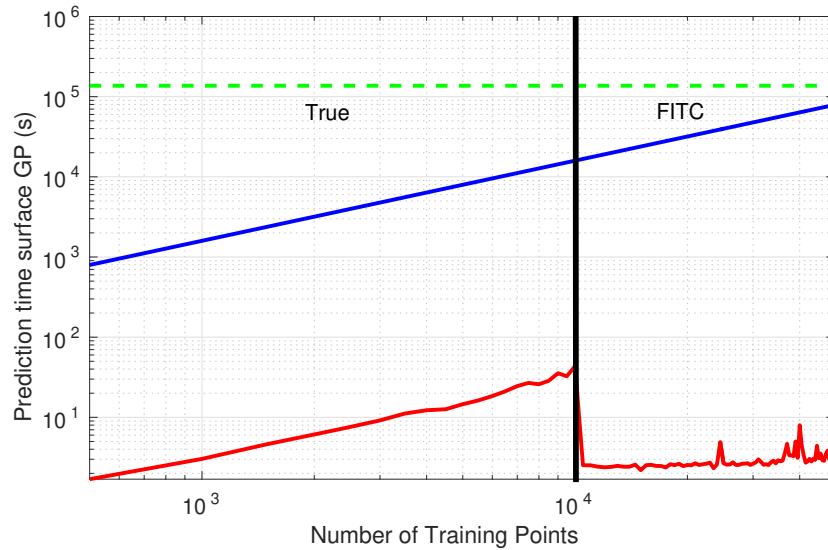
Figure 10: **Learning curves for classifier active vs. random learning for the a) 4-input & b) 2-input problem**.

(a)

(b)

Figure 11: **Surface GP timing performance for the 4-input problem, see section 3.2 for the learning curves.** a) Training, and b) Prediction times with increasing numbers of training inputs. The predictions are drawn over the test dataset $\mathcal{D}_{test}$ which contains $100,000$ test points. The green broken line shows the time required by the simulator to evaluate $\mathcal{D}_{test}$. The blue line shows the simulation time for an increasing number of inputs and the red line shows the GP training (hyperparameter learning) and prediction time. The magenta line shows the total training time which is the sum total of the simulation and GP training time. The black vertical line demarcates the training size beyond which we use the FITC covariance. Here we see the potential benefit in terms of speed when using the FITC method, despite the slightly larger error that we observed in Fig. 7.
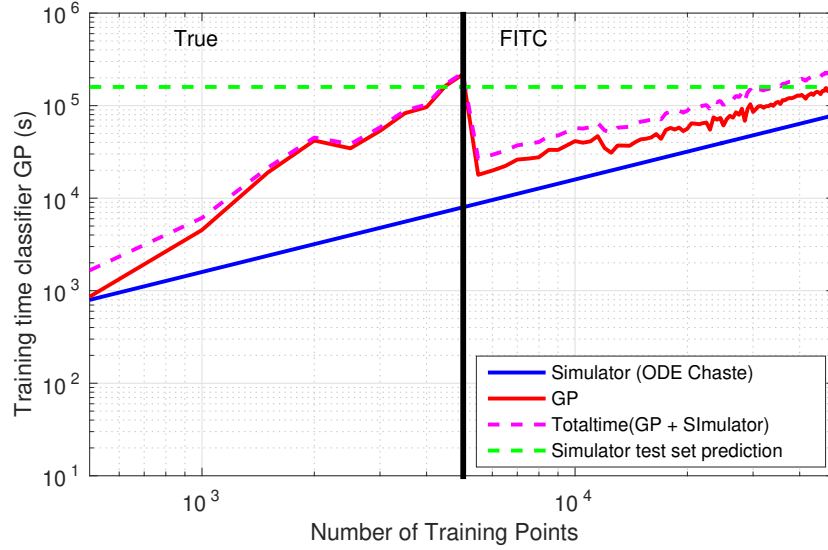
(a)
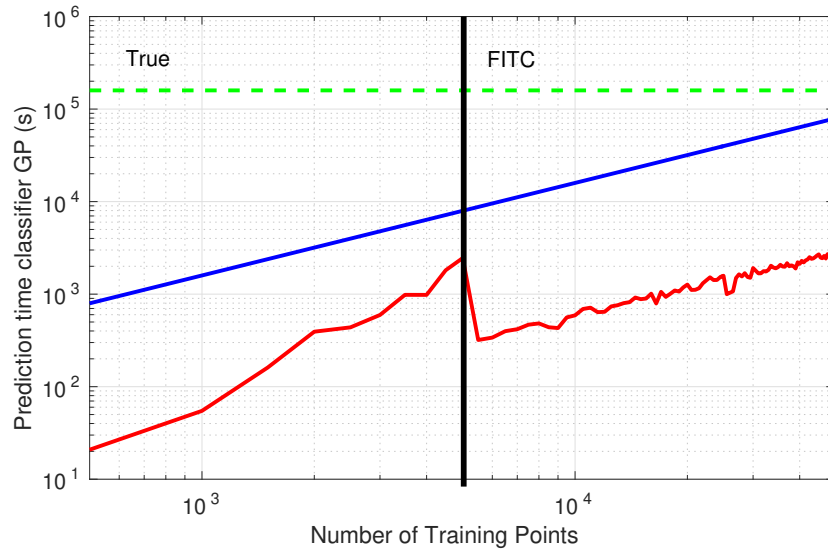
(b)

Figure 12: **Training a) and Prediction b) time for the classifier GP for the 4-input problem, see section 3.2 for the learning curves, with increasing number of training inputs**. The predictions are drawn over the test dataset $\mathcal{D}_{test}$ which contains $100,000$ test points. The green broken line shows the time required by the simulator to evaluate $\mathcal{D}_{test}$. The blue line shows the simulation time for increasing number of inputs and the red line shows the GP training (hyperparameter learning) and prediction time. The magenta line shows the total training time which is the sum total of the simulation and GP training time. The black vertical line demarcates the training size beyond which we use the FITC covariance.
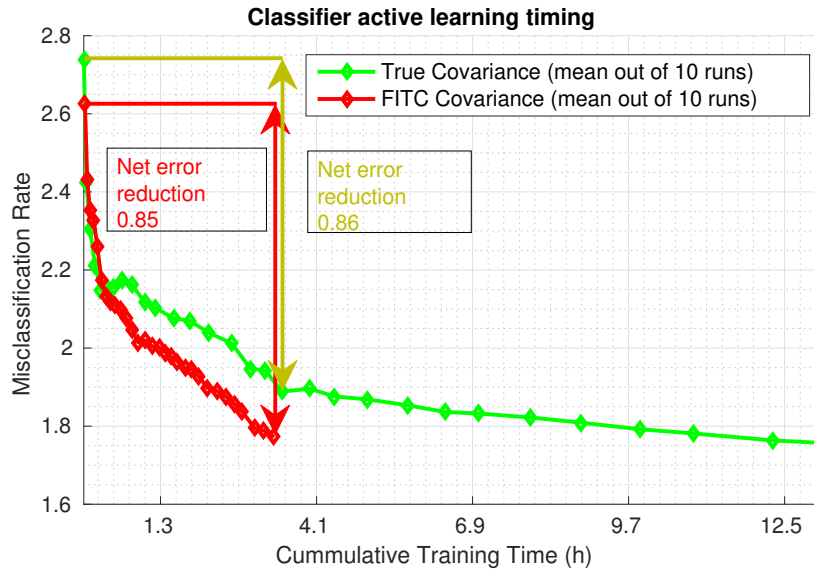
Figure 13: **Misclassification error reduces at the same rate using FITC and true covariance for carrying out active learning.** Comparison of true vs. FITC covariance for classifier active learning. Active learning is repeated 10 times using different initial datasets. We notice same average (out of 10 runs) rate of error reduction using both covariances. The horizontal and vertical lines point out the average error reduction observed after spending the time required for carrying out active learning using FITC (red lines). We used a swarm size of $n_s = 50$ and a separate initial set of $n_1 = 500$ random training points for each repetition. Within this same time budget we achieve similar average error reduction using the true covariance (green line).
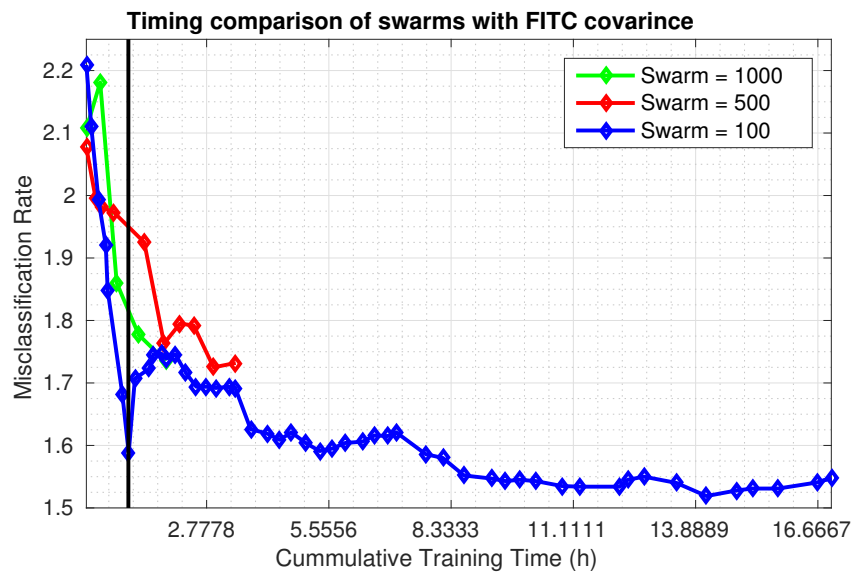
Figure 14: **More rounds of PSO with smaller swarm sizes reduces error most efficiently.** Effect of swarm sizes on active learning time with FITC covariance.
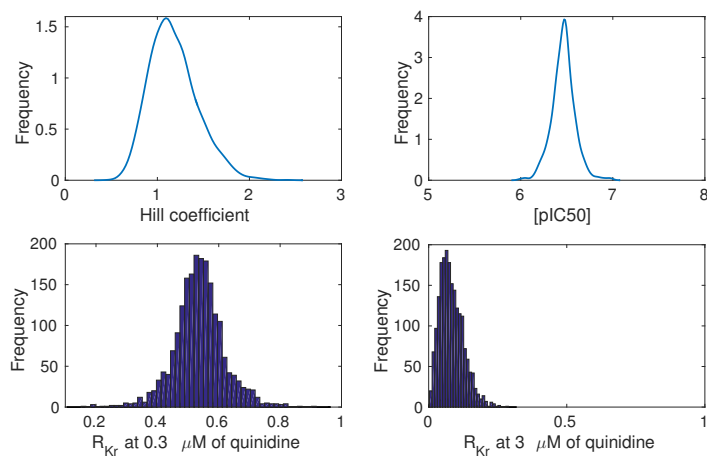
Figure 15: **Distributions of the concentration-effect-curve parameters as obtained in [44] through MCMC, and the corresponding distribution of R_{Kr} (hERG ion channel scalings) at different doses of quinidine.** The top row shows the marginal posterior distributions (as kernel density estimates) of Hill coefficient n and $[pIC50] = -\log_{10}([IC50])$ estimated from the dose-response data in [43] for quinidine compound action on hERG channel. The bottom row shows the corresponding distributions of $\boldsymbol{R}_{Kr}$, as histograms, at quinidine concentrations $0.3\,\mu\text{M}$ (left) and $3\,\mu\text{M}$ (right) calculated using equation (1). Each of the kernel density estimates and histograms are made using 2000 samples.
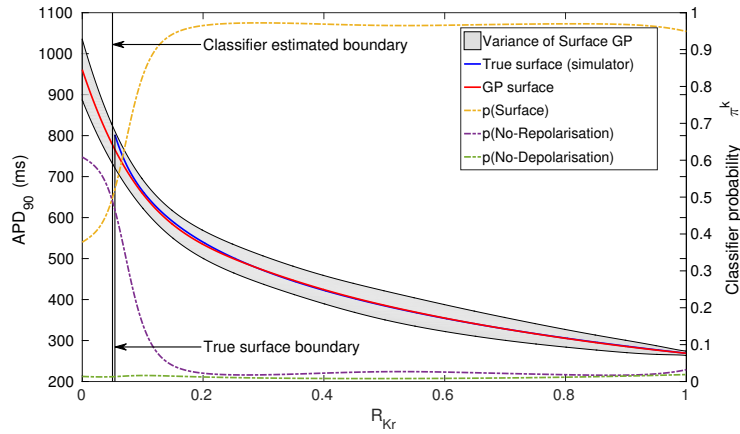
Figure 16: **1-dimensional slice of surface prediction along with binary classification posterior probabilities on a linear grid of R_{Kr}. Misclassification is associated with low classifier probabilities.** The surface GP prediction of APD_{90} values is plotted (red line) along with the emulator variance (shaded area), scale shown on left hand axis. The true solution from running the ODE solver (simulator) is plotted as a blue line. The true and the classifier's estimated class boundary are shown as vertical lines. The posterior class probability p(Surface) of the surface (valid AP vs rest of the regions) region is shown on the right hand axis (orange line) and reduces rapidly in the misclassification region between the two vertical lines. The variance of the surface GP also reduces in this region. The corresponding posterior class probabilities p(No-Repolarisation) and p(No-Depolarisation) of the non-repolarising and non-depolarising regions vs rest of the regions respectively, are also shown on the right hand axis as broken lines. Note that the true solution is contained within the emulator variance.
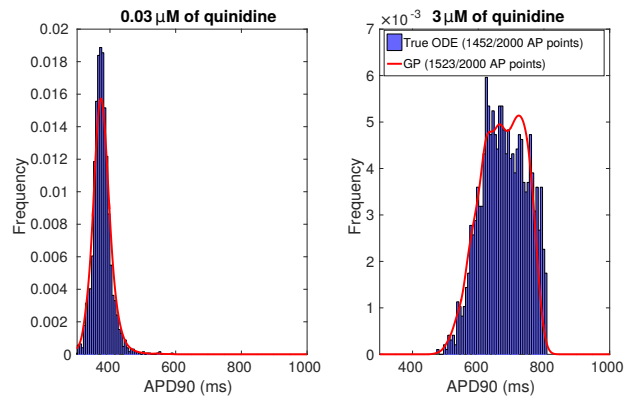
Figure 17: **Uncertainty propagation from concentration-effect-curve parameters to APD$_{90}$.** Distributions of APD$_{90}$ obtained through emulation (evaluated as given in equation (47)) are plotted as red lines for corresponding $\boldsymbol{R}_{Kr}$ values as represented in Fig. 15. Distributions of APD$_{90}$ obtained through Monte Carlo samples with a full simulation of the ODE system are shown as histograms.