Published in final edited form as:

Prog Biophys Mol Biol. 2015 January ; 117(1): 4–18. doi:10.1016/j.pbiomolbio.2015.01.008.

Uncertainty quantification of fast sodium current steady-state inactivation for multi-scale models of cardiac electrophysiology

Pras Pathmanathan^a, Matthew S. Shotwell^b, David J. Gavaghan^c, Jonathan M. Cordeiro^d, and Richard A. Gray^a

Pras Pathmanathan: pras.pathmanathan@fda.hhs.gov; Matthew S. Shotwell: matt.shotwell@vanderbilt.edu; David J. Gavaghan: gavaghan@cs.ox.ac.uk; Jonathan M. Cordeiro: jcordeiro@mmrl.edu; Richard A. Gray: richard.gray@fda.hhs.gov

^aU.S. Food and Drug Administration, 10903 New Hampshire Avenue (WO 62), Silver Spring, MD 20993, USA

^bDepartment of Biostatistics, Vanderbilt University Medical Center, 2525 West End, Ste. 11000 Nashville, TN 37203, USA

^cDepartment of Computer Science, University of Oxford, Parks Road, Oxford, OX1 3QD, UK

^dMasonic Medical Research Laboratory, 2150 Bleecker St, Utica, NY 1350, USA

Abstract

Perhaps the most mature area of multi-scale systems biology is the modelling of the heart. Current models are grounded in over fifty years of research in the development of biophysically detailed models of the electrophysiology (EP) of cardiac cells, but one aspect which is inadequately addressed is the incorporation of uncertainty and physiological variability. Uncertainty quantification (UQ) is the identification and characterisation of the uncertainty in model parameters derived from experimental data, and the computation of the resultant uncertainty in model outputs. It is a necessary tool for establishing the credibility of computational models, and will likely be expected of EP models for future safety-critical clinical applications. The focus of this paper is formal UQ of one major sub-component of cardiac EP models, the steady-state inactivation of the fast sodium current, I_{Na} . To better capture average behaviour and quantify variability across cells, we have applied for the first time an ‘individual-based’ statistical methodology to assess voltage clamp data. Advantages of this approach over a more traditional ‘population-averaged’ approach are highlighted. The method was used to characterise variability amongst cells isolated from canine epi and endocardium, and this variability was then ‘propagated forward’ through a canine model to determine the resultant uncertainty in model predictions at different scales, such as of upstroke velocity and spiral wave dynamics. Statistically significant differences between epi and endocardial cells (greater half-inactivation and less steep slope of steady state inactivation curve for endo) was observed, and the forward propagation revealed a lack of robustness of the model to underlying variability, but also surprising robustness to

Correspondence to: Pras Pathmanathan, pras.pathmanathan@fda.hhs.gov.

Disclosure The mention of commercial products, their sources, or their use in connection with material reported herein is not to be construed as either an actual or implied endorsement of such products by the Department of Health and Human Services.

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

variability at the tissue scale. Overall, the methodology can be used to: (i) better analyse voltage clamp data; (ii) characterise underlying population variability; (iii) investigate consequences of variability; and (iv) improve the ability to validate a model. To our knowledge this article is the first to quantify population variability in membrane dynamics in this manner, and the first to perform formal UQ for a component of a cardiac model. The approach is likely to find much wider applicability across systems biology as current application domains reach greater levels of maturity.

Keywords

Cardiac modelling; uncertainty quantification; population variability; nonlinear mixed effects modelling

1. Introduction

Mathematical and computational modelling of cardiac electrophysiological activity is one of the most advanced fields within systems biology, building on over fifty years of development of increasingly detailed and biophysically realistic models [29, 19]. Current computational models of whole-heart electrophysiology are inherently multi-scale and highly complex, integrating temporal models of the electrophysiology (EP) of a cardiac myocyte—referred to as cardiac *cell models*, themselves made up of sub-models representing various transmembrane and intracellular phenomena—with spatio-temporal models of electrical wave propagation through tissue. Together with experimentation, mathematical modelling has been an essential component in the cardiac EP field to unravel mechanisms and develop hypotheses. Currently, there is significant research on translating whole-heart modelling to clinical settings, such as using model predictions to optimise ablation therapy, defibrillation, and cardiac resynchronisation therapy [3, 40, 28], and as a possible means of assessing the cardiac safety of potential new drug compounds [24]. Such proposed safety-critical applications require the reliability, credibility and robustness of model predictions to be rigorously assessed, across all of the component parts of the multi-scale models. Rigorous approaches for assessing computational models already exist in various areas of engineering, and provide the potential to evaluate components of multi-scale models in a rigorous and systematic manner, but to our knowledge have not previously been applied for cardiac EP models. One such aspect is *uncertainty quantification* (UQ), which has the potential to support the consideration of the effects of intra- and inter-individual variability across the systems biology field.

UQ is the identification and quantification of the sources of uncertainty in a model, and the determination of how this uncertainty affects simulation results [26, 37]. A major source of uncertainty¹ is the *uncertainty in parameter values*. This can arise in various forms, two of the most important being uncertainty due to experimental error (values of model parameters will not usually be known exactly, due to measurement error), and uncertainty due to

¹Uncertainty is often categorised into two types: *aleatory*, or irreducible, uncertainty; and *epistemic*, or reducible, uncertainty. Aleatory uncertainty is uncertainty arising due to inherent randomness and natural variability, whereas epistemic uncertainty is uncertainty due to lack of knowledge (hence potentially reducible). For more details see [37, 22].

population variability or inherent biological variability. The aim of UQ is to characterise and quantify such uncertainties in model inputs (*uncertainty characterisation*), and then ‘propagate them forward’ through the model to determine how they affect model predictions (*uncertainty propagation*). Typically, this is accomplished using probability distributions: uncertainty in the parameters is characterised and quantified in the form of distributions which are then propagated forward to obtain resultant distributions for ‘quantities of interest’ (QOIs) (model outputs). This can be contrasted with the non-UQ approach, where a single value of each parameter is used, resulting in a single value of the QOI. UQ is related to sensitivity analysis, but in sensitivity analysis inputs are often (arbitrarily) varied about a particular parameter choice, rather than characterised through probability distributions using experimental data. For a more detailed introduction to UQ see for example [37].

UQ is part of a larger framework for assessing the credibility of computational models known as ‘verification, validation and uncertainty quantification’ (VVUQ), also known as ‘verification and validation’ (V&V), which has been utilised to great effect in modelling efforts in the physical sciences and engineering [26]. Verification is the process of confirming that a computational model (software) correctly implements an underlying mathematical model; validation compares a model with reality. In previous work, we have advocated and discussed the use of VVUQ in cardiac EP modelling [34], as well as presented novel methods for cardiac EP model verification [35]. Similarly to [35, 34], the fundamental motivation for the present article is that some understanding of how uncertainties affect predictions may be *expected* of computational heart models for future clinical applications.

UQ applied to complex physiological models is formidable task, and whilst sensitivity analyses of cardiac models have been performed (see e.g. [38, 14]), as far as we are aware no formal data-based UQ has ever been performed for parameters within cardiac models, despite the maturity of the field. UQ has been performed in a drug-based cardiac application, for parameters governing channel block of various compounds [9]. Also, in [15] a sophisticated framework for uncertainty propagation (the ‘stochastic Galerkin method’), was applied to determine the impact of uncertain rate coefficients in ion channel models. However, a simplified approach was used for the uncertainty characterisation stage, that was not based on experimental data. The difficulty in performing UQ of cardiac models is the huge number of parameters in such models; it is not uncommon for cardiac cell models to involve dozens of state variables and more than a hundred parameters. Additional complexity arises when these models are coupled to equations governing spatial processes and solved on realistic heart geometries. Figure 1 illustrates the components that make up a computational model of whole-heart EP. Processes at multiple scales are typically modelled, from channel to cell to organ. Whilst cardiac modelling is a successful application of a multi-scale systems biology modelling approach to quantitatively study physiological processes across a wide range of spatial and temporal scales, uncertainty is rarely included. In particular, this means that a concept as fundamental as natural biological variability is in general inadequately addressed.

In this article we perform rigorous uncertainty quantification for a sub-component of cardiac models, the steady-state inactivation of the fast sodium channel, I_{Na} . As far as we are aware,

this article is the first to perform formal, experimental data-based, uncertainty quantification for any part of a cardiac cell model. Accurate quantitative characterisation of I_{Na} inactivation is of significant importance because its function has been implicated in a wide range of physiological phenomena, including disease. For example, incomplete I_{Na} inactivation is the cause of the late ‘persistent’ sodium current which has been linked to a variety of disease states, such as heart failure and long QT syndrome (see [30, 31] for reviews). In addition, the gating parameters of I_{Na} inactivation appear to be different in atrial and ventricular myocytes [6], and accordingly much research is currently underway to develop pharmaceutical drug therapy including chamber-specific approaches, such as using ranolazine to treat atrial fibrillation with minimal effects on ventricular electrophysiology [2]. Therefore, it is important to accurately model I_{Na} inactivation to evaluate normal and abnormal function and the effects of drugs, including side effects [27]. Assessing the reliability of I_{Na} inactivation data, as well as characterising population variability within and across cell types, is a necessary part of these efforts.

Steady-state activation or inactivation of membrane gating variables are typically modelled by fitting sigmoidal functions to voltage clamp data, as illustrated in Figure 1. The traditional approach is to compute a single ‘best-fit’ parameter set for the chosen function given experimental data, e.g. using least-squares optimisation, possibly with some manual alterations (e.g. [33, 39]). In this paper we consider several statistical methods for analysing voltage clamp data, highlighting the difference between so-called ‘population-averaged’ and ‘individual-based’ statistical approaches. The former is closely related to the traditional approach in cardiac model development, whereas the latter has never, to the best of our knowledge, been used to analyse voltage clamp data before. We will show that the individual-based approach has significant advantages, and better characterises steady-state I_{Na} inactivation for an ‘average’ cell, as well as enabling quantification of population variability. We apply these methods to several experimental datasets on steady-state I_{Na} inactivation for both epi and endocardial canine cardiomyocytes, quantifying for the first time population variability in I_{Na} inactivation and identifying statistically significant differences between epi and endocardial cells. Then, by propagating the uncertainty forward, we will study the robustness of a cell model to the uncertainty in steady-state I_{Na} inactivation, and quantify the resultant uncertainty in model predictions. For the uncertainty propagation, we use a simple ‘brute force’ numerical integration approach, which will be feasible due to the small number of random variables involved (due in part to some model simplification that we will perform). The framework will be shown to improve the ability to validate a model.

We begin in Section 1.1 by writing the general equations of a cell model and discussing the fast sodium current, before describing in Section 1.2 the experimental canine data to be used. In Section 2.1 we describe the population-averaged and individual-based statistical methods, applying them to the data to quantify uncertainty in Section 2.2. Then, in Section 3, we propagate the uncertainty forward through the various scales using a full canine cell model, analysing the effect of uncertainty on upstroke velocity, action potential duration and spiral wave dynamics. Finally, we discuss the implications of our results in Section 4.

1.1. Single-cell models of cardiac electrophysiology

Physiological cardiac cell models are ordinary differential equations of the form

$$\mathcal{C}_m \frac{dV}{dt} + I_{\text{ion}}(V, \mathbf{u}) = I_{\text{stim}}, \quad (1)$$

$$\frac{d\mathbf{u}}{dt} = \mathbf{f}(V, \mathbf{u}), \quad (2)$$

where V is the transmembrane voltage, $\mathcal{C}_m$ the membrane capacitance per unit area, and I_{stim} is an applied stimulus current. $\mathbf{u}$ is a model-dependent set of state variables (usually probabilities of channel gates being open (see below), and ionic concentrations), $\mathbf{f}$ is a model-dependent function governing the dynamics of $\mathbf{u}$, and I_{ion} is a model-dependent function representing the ionic current across the cell membrane. The ionic current is the cumulative current arising from membrane currents corresponding to various channels, pumps and exchangers, and for example may be represented as

$$I_{\text{ion}} = I_{\text{Na}} + I_{\text{CaL}} + I_{\text{NaCa}} + I_{\text{Kr}} + I_{\text{Ks}} + I_{\text{to}} + \dots$$

with each term on the right-hand side corresponding to a different membrane current. The system (1) and (2) is solved to obtain the action potential (AP). A schematic of this process is illustrated in the central part (cell model panel) of Figure 1. See [20] for a more complete introduction.

The current we are concerned with is I_{Na} , the current resulting from transport of sodium ions through the fast sodium channel. This channel opens in response to a stimulus current, prompting a rapid influx of sodium ions and bringing about ‘all-or-none’ depolarisation. Under the most common formulation (the Hodgkin-Huxley [17] formalism) I_{Na} is represented as:

$$I_{\text{Na}} = g_{\text{Na}} m^3 h j (V - E_{\text{Na}}), \quad (3)$$

where g_{Na} is the maximum conductance of the channel, E_{Na} is the Nernst equilibrium potential of sodium ions, and $m(t)$, $h(t)$, and $j(t)$ are *gating variables*, time-dependent probabilities of channel ‘gates’ being open. The gating variables are state variables, that is, (abusing notation) $\{m, h, j\} \in \mathbf{u}$. m is an activation variable (increases with increasing voltage), while h and j are inactivation variables (the opposite). The dynamics of gating variables are typically modelled as

$$\frac{dw}{dt} = \frac{w_{\infty}(V) - w}{\tau_w(V)},$$

where w is m , h or j , and $w_{\infty}(V)$ represents the voltage-dependent steady-state activation/inactivation and $\tau_w(V)$ the voltage-dependent time-constant. Equivalently,

$$\frac{dw}{dt} = \alpha_w(V)(1 - w) - \beta_w(V)w,$$

where $\alpha_w(V)$ and $\beta_w(V)$ are voltage-dependent transition rates. Accordingly, the relationships $w_\infty = \alpha_w/(\alpha_w + \beta_w)$, $\tau_w = 1/(\alpha_w + \beta_w)$, $\alpha_w = w_\infty/\tau_w$, $\beta_w = (1 - w_\infty)/\tau_w$ hold. For each gate, the functions $w_\infty(V)$ and $\tau_w(V)$ (or equivalently $\alpha_w(V)$ and $\beta_w(V)$) are typically prescribed by fitting chosen parameterised functions to experimental measurements, as illustrated in Figure 1 (left). This is performed for each gate of each current, and is part of the reason why cell models can contain so many variables and parameters.

The paper is concerned with steady-state inactivation of the sodium channel, i.e. $h_\infty(V)$ and $j_\infty(V)$. The data described in the following section corresponds to the total steady-state inactivation, i.e. $h_\infty \times j_\infty$.

1.2. Data

We have two separate I_{Na} datasets resulting from an inactivation voltage clamp protocol, from which steady-state inactivation has been estimated. The first dataset, published in Cordeiro *et al.* 2008 [8] includes data from 16 epicardial and 16 endocardial canine cardiomyocytes. The second dataset, published in Murphy *et al.* 2011 [25], comprises ten epicardial and ten endocardial cells. The datasets will be referred to as *Cord-Epi*, *Cord-Endo*, *Murphy-Epi* and *Murphy-Endo*, and are presented in Figure 2 in separate panels.

Below, we briefly outline the experimental procedure used. For full details see [8, 25].

1.2.1. Isolation of myocytes—Myocytes were prepared from the left ventricle (LV) of canine hearts. A wedge consisting of the LV wall supplied by a descending branch of the circumflex artery (left marginal artery) was excised, cannulated and perfused with nominally Ca^{2+} -free Tyrodes solution containing 0.1% BSA for about 5 minutes. The wedge preparations were then subjected to enzyme digestion with the nominally Ca^{2+} -free solution supplemented with 0.5 mg/ml collagenase (Type II, Worthington) and 1 mg/ml BSA for 8–12 minutes. After perfusion, thin slices of tissue from the epicardium (< 2 mm from the epicardial surface) and endocardium (< 2 mm from the endocardial surface) were shaved from the wedge using a dermatome. The tissue slices were then placed in separate beakers, minced in fresh buffer containing 0.5 mg/ml collagenase, 1 mg/ml BSA and agitated. The supernatant was filtered, centrifuged and the pellet containing the myocytes was stored at room temperature.

1.2.2. Voltage clamp recordings of peak I_{Na} —Peak sodium current (I_{Na}) was recorded in low extracellular sodium and at room temperature to ensure adequate voltage control. Command voltages were delivered using a MultiClamp 700A amplifier and data acquired via a DigiData 1322 computer interface using pClamp 9 software (Molecular Devices, Foster City, CA). Patch pipettes were pulled from borosilicate glass (1.5 mm o.d. and 1.1 mm i.d.) on a vertical puller (Model PP-830, Narashige Instruments, Japan). The electrode resistance was 0.9–2.0 M Ω when filled with the internal solution². The membrane was ruptured by applying negative pressure and series resistance compensated by 75–80%.

Whole cell current data was acquired at 20–50 kHz and filtered at 5 kHz. Currents were normalised to cell capacitance.

Steady-state inactivation of I_{Na} was measured using a standard prepulse-test pulse voltage-clamp protocol (illustrated in Figure 1). The steady-state inactivation voltage clamp protocol consisted of 500 ms prepulses between -120 mV and -15 mV, followed by a 50ms test pulse to -20 mV. The holding potential was -120 mV to recruit all available sodium channels. Peak current was then normalised to the maximum current and plotted against prepulse voltage.

2. Uncertainty quantification

2.1. Statistical approaches

Let us define the following sigmoid functions

$$F_n(V; V_0, K) = \frac{1}{(1 + \exp((V - V_0)/K))^n}. \quad (4)$$

Here n is a small integer, and V_0 and K are two parameters. Such sigmoid functions are commonly used to characterise steady-state inactivation/activation for membrane gating variables.

Let $F = F_2$ be the function we will fit to experimental data. With the choice of $n = 2$, V_0 is

the voltage at which inactivation $= \frac{1}{4}$ and K is inversely proportional to the gradient of F at V_0

$$F(V_0; V_0, K) = \frac{1}{4}, F'(V_0; V_0, K) = -\frac{1}{4K}.$$

Our treatment of the parameter n as fixed, rather than unknown, is motivated by statistical and biological reasoning and evidence. The shapes of the sigmoid functions F_n grow increasingly alike as n increases. Preliminary statistical investigation suggests that, for increasing values of n , additional experimental replicates are required to estimate n with equivalent precision. By fixing the value of n , this additional uncertainty is avoided. The value $n = 2$ was selected to match the model for I_{Na} (3), where inactivation is determined by the product of two gating variables, h and j , and the notion that both h_∞ and j_∞ have an identical F_1 shape (the simplest choice). Thus, the observed experimental data should exhibit a shape that resembles $h_\infty \times j_\infty$, i.e. $F_1^2 \equiv F_2$ [39]. We will show that $n = 2$ provides a good fit to the data, and in the supplement we show that $n = 1$ does not.

Let $\mathbf{p} = (V_0, K)$. Our aim is to estimate $\mathbf{p}$ using experimental data, such that the estimate characterises steady-state I_{Na} inactivation for a ‘typical’ epi or endocardial cell. In addition,

²External solution contained (in mM): Choline Cl 120, NaCl 10, Na⁺ acetate 2.8, CaCl₂ 0.5, KCl 4, MgCl₂ 1.5, CoCl₂ 1, glucose 10, HEPES 10, NaOH 5, BaCl₂ 0.1, pH adjusted to 7.4 with NaOH/HCl. The pipette solution contained (mM): NaCl 15, CsF 120, MgCl₂ 1, KCl 5, HEPES 10, Na₂ATP 4, EGTA 10, pH adjusted to 7.2 with CsOH.

we seek to quantify the variability in $\mathbf{p}$ among cells within a population, and the uncertainty due to measurement error. We consider two methods for this purpose, which are broadly classified as: ‘population-averaged’ (Section 2.1.1) or ‘individual-based’ (Section 2.1.2). In the former category, estimated quantities are averages over a population, e.g., the average I_{Na} inactivation at transmembrane voltage V . In contrast, individual-based methods are used to estimate properties of the average individual, e.g., the I_{Na} inactivation at transmembrane voltage V for the ‘average cell’.

2.1.1. Population-averaged approach: weighted nonlinear least squares—A simple approach for estimating a population-averaged parameter, denoted $\mathbf{p}^*$, is to perform a simple weighted least squares optimisation using the average inactivation over cells at each voltage level. This is the approach most often used in cardiac model development. However, Figure 2 illustrates a potential problem with this approach. In the *Cord-Endo* sub-figure (top left), the gradients for each individual cell are similar, but the gradient of the average curve is smaller in magnitude. Indeed, the average curve does not resemble any of the individual curves. This raises the suspicion that $\mathbf{p}^*$ computed in this way may poorly approximate $\mathbf{p}$ for each individual cell, and is one motivation for the individual-based approach to be described in Section 2.1.2.

Let us briefly outline the statistical basis of weighted nonlinear least squares (NLS) regression³. Let y represent the measured steady-state I_{Na} inactivation, dependent on voltage V . Let $I = 1, \dots, M$ denote cell index and $i = 1, \dots, N$ denote voltage index, and let y_i^I denote the observed inactivation value for cell I at voltage V_i . Least squares regression is based on the assumption that observations take the form [4]

$$y_i^I = F(V_i; \mathbf{p}^*) + \varepsilon_i^I, \quad (5)$$

for unknown parameter $\mathbf{p}^*$, and where the error terms ε_i^I are assumed to be independent random variables with mean zero. If we further assume that errors are normally distributed with standard deviation σ/ω_i , for voltage-specific weights $\omega_i > 0$, then the weighted least squares estimator for $\mathbf{p}^*$ is the argument that maximises the associated likelihood function

$$\mathcal{L}(\mathbf{p}^* | \{y_i^I\}; \omega_1, \dots, \omega_N, \sigma) = \prod_{I,i} \frac{\omega_i}{\sigma \sqrt{2\pi}} \exp \left(-\frac{\omega_i^2}{2\sigma^2} (y_i^I - F(V_i; \mathbf{p}^*))^2 \right) = C \exp \left(-\frac{1}{2\sigma^2} \sum_{I,i} \omega_i^2 (y_i^I - F(V_i; \mathbf{p}^*))^2 \right) \quad (6)$$

where C is a constant with respect to $\mathbf{p}^*$. Thus, the likelihood is maximised when the weighted sum of squares $\sum_{I,i} \omega_i^2 (y_i^I - F(V_i; \mathbf{p}^*))^2$ is minimised, i.e. a weighted least-squares optimisation. Expression (6) may also be written in terms of the average inactivation at each voltage level (circles in Figure 2), denoted $\bar{y}_i$,

³Here, ‘nonlinear’ means that the model for average inactivation is not a linear function of the parameters $\mathbf{p}^*$

$$\mathcal{L}(\mathbf{p}^* | \bar{y}_1, \dots, \bar{y}_N; \omega_1, \dots, \omega_N, \sigma) = C' \exp \left(-\frac{M}{2\sigma^2} \sum_i \omega_i^2 (\bar{y}_i - F(V_i; \mathbf{p}^*))^2 \right),$$

where C' is distinct from C , but also constant in $\mathbf{p}^*$. Hence, the averages $\bar{y}_i$ are sufficient for weighted least-squares optimisation.

However, for our application, the population-averaged model (5) is a gross simplification in which all sources of variability (in particular population variability and measurement error) are represented by the errors ε_i^I . Figure 2 makes it clear that sequential cell-specific errors ε_i^I and ε_i^I are not independently distributed about the population average—if that were true, then sequential cell-specific inactivation measurements would take values both above and below the population average inactivation values. These deficiencies partly motivate the individual-based approach to be described in Section 2.1.2. To compare the two approaches, we perform weighted NLS using the `nls` function from the statistical software environment R. This implementation estimates both $\mathbf{p}^*$ and σ , for prespecified weights ω_i . The weights were computed by inverting the sample standard deviation at each voltage V_i . A Wald-type confidence region for $\mathbf{p}^*$ was computed using derivatives of F at the estimated $\mathbf{p}^*$ [4]. We denote the weighted NLS estimate of $\mathbf{p}^*$ by $\mathbf{p}^{\text{NLS}}$.

The reasons for evaluating the population-averaged approach are as follows. First, as explained above, this approach is most commonly used in the development of transmembrane current models for excitatory cells. Second, unlike the individual-based method described below, weighted NLS can be used when *only the average* inactivation values (i.e., $\bar{y}_i$) are available (which is often how such data are published). We will construct a confidence region for $\mathbf{p}^*$ using weighted NLS, and then compare this to the corresponding confidence region for $\mathbf{p}$ using the individual-based method. The weighted NLS method does not, however, provide an estimate of population variability about $\mathbf{p}$.

We note that, in addition to weighted NLS, there is a further generalised least squares method that accounts for correlation among measurements on individual cells, but cannot be applied when only the averaged data are available. In this paper, we use the conventional weighted least squares method, but briefly discuss the generalised least squares approach in the supplemental material.

2.1.2. Individual-based approach: nonlinear mixed effects modelling—

Nonlinear mixed effects (NLME) modelling is a formal statistical methodology for the analysis of population heterogeneity with respect to a nonlinear relationship. The archetypal example application for NLME modelling comes from pharmacokinetics, and involves the question of how the plasma concentration of an administered drug, as a function of time, varies across individuals. Given a model $f(t; \mathbf{p}_I)$ of concentration versus time, involving unknown individual-based parameters $\mathbf{p}_I$, the analysis provides an estimate $\mathbf{p}_I$ for each individual, together with population average parameters. See [12, 36] for an introduction to NLME modelling. For our purposes, each individual is a cell, and the independent variable is voltage rather than time.

Recalling that $\mathbf{p} = (V_0, K)$, we will make use of the transformation $\mathbf{q} = (V_0, \log(K))$. We use $\mathbf{q}$ rather than $\mathbf{p}$ in the following statistical model to force K to be positive. However, the results later will be presented in terms of $\mathbf{p}$.

The NLME model is comprised of: (i) an individual-level model:

$$y_i^I = F(V_i; \mathbf{p}_I) + \delta_i^I, \quad (7)$$

where $\mathbf{p}_I$ is cell-specific, δ_i^I represents measurement error (and also ‘realisation error’ [12], assumed to be negligible in comparison to measurement error); and (ii) a population model:

$$\mathbf{q}_I = \mathbf{q}^* + \mathbf{b}_I, \quad (8)$$

where $\mathbf{q}_I$ is related to $\mathbf{p}_I$ via the above transformation, $\mathbf{q}^*$ is a so-called ‘fixed effect’ representing the parameters for an ‘average cell’, and $\mathbf{b}_I$ is a cell-specific ‘random effect’ corresponding to the deviation between $\mathbf{q}_I$ and $\mathbf{q}^*$. We assume, as is typical in NLME modelling, that the random effects are normally-distributed with mean zero and covariance matrix Ψ . Thus, inactivation for each cell is assumed to be characterised by a unique parameter $\mathbf{q}_I$, drawn from a (bivariate) normal distribution with mean $\mathbf{q}^*$ and covariance matrix Ψ . We further assume that the measurement errors δ_i^I are independent and identically-distributed with mean zero and standard deviation σ_δ . (The validity of these assumptions will be tested as discussed shortly). The NLME problem is then to estimate $\{\mathbf{q}^*, \Psi, \sigma_\delta\}$. This is accomplished, for example, by maximising the expected likelihood function

$$\mathcal{L}(\mathbf{q}^*, \Psi, \sigma_\delta | y_1^I, \dots, y_N^I) = \prod_{I,i} \int p(y_i^I | \mathbf{b}_I, \mathbf{q}^*, \sigma_\delta) p(\mathbf{b}_I | \Psi) d\mathbf{b}_I.$$

To solve this optimisation problem, we use the statistical analysis package `nlme`, a contributed package of the software environment `R`. Let $\mathbf{p}^{\text{NLME}}$ be the estimate of $(q_1^*, \exp(q_2^*))$, i.e. the estimated fixed effect in (V_0, K) space. $\mathbf{p}^{\text{NLME}}$ represents an estimate of the parameters corresponding to an ‘average cell’, in contrast to $\mathbf{p}^{\text{NLS}}$ which estimates the parameters corresponding to the averaged inactivation curve. As with weighted NLS, the NMLE method provides a mechanism to construct a confidence region around $\mathbf{p}^{\text{NLME}}$. Estimates for each $\mathbf{p}_I$ are also available, and using the assumed normal property of $\mathbf{q}_I$, the NLME method also provides a *prediction region* for $\mathbf{p}_I$ (a (near-)ellipse centred at $\mathbf{p}^{\text{NLME}}$ with size and orientation dependent on Ψ), which represents the *population variability* in $\mathbf{p}_I$.

2.2. Application to steady-state inactivation data

First, we note that the *Cord-Epi* dataset contains one visually anomalous curve (see Figure 2, top-right). If the NLME estimates are computed with this curve included, the NLME estimate for the corresponding $\mathbf{p}_I$ falls well outside the overall 95% prediction region (see supplemental material for details). Therefore, the anomalous curve was considered an outlier and excluded from the following analysis.

Figure 3 plots the estimated inactivation curve obtained using the weighted NLS and NLME methods using, separately, the *Cord-Endo* and *Cord-Epi* datasets. The weighted NLS fit (the function $F(V; \mathbf{p}^{\text{NLS}})$) and the NLME fixed effect fit (the function $F(V; \mathbf{p}^{\text{NLME}})$) are overlaid against the experimental data and the average inactivation at each voltage. We see that the weighted NLS fits are very close to the averaged data as expected. For the *Cord-Epi* dataset the NLS and NLME results are very similar, for the *Cord-Endo* there are some slight differences, in particular in the slope of the fit. To emphasise that the NLME approach results in estimates of $\mathbf{p}_I$ for each individual cell, Figure 3 also plots the fitted curves for each individual cell. Visually, the selected sigmoidal function (4) with $n = 2$ fits the data well. A more rigorous assessment, using statistical diagnostics, is provided in the supplemental material, where evidence supporting the NLME modelling assumptions is presented, and it is also shown that $n = 1$ does not provide a good fit.

Figure 4 plots the results of the two methods in 2D parameter space, applied to each of the four datasets. For each dataset, the NLS estimate $\mathbf{p}^{\text{NLS}}$ is plotted together with the associated 95% confidence region; and similarly the NLME estimate $\mathbf{p}^{\text{NLME}}$ and 95% confidence region are plotted. Also visualised is the NLME 95% prediction region, which can be used to gauge population variability, together with NLME estimates of $\mathbf{p}_I$ for each individual cell of the dataset. This figure is highly informative, and we can make a number of observations:

NLS vs NLME—For the reasons described above we consider the NLME results as more likely to capture the behaviour of an average cell. Comparing the 95% confidence regions associated with $\mathbf{p}^{\text{NLS}}$ and $\mathbf{p}^{\text{NLME}}$ we observe that the least squares approach generally results in over-estimation of the parameter K (under-estimation of magnitude of gradient), in particular with the *Cord-Endo* and *Murphy-Epi* datasets. For three of the four datasets, the NLS and NLME confidence regions do not intersect, suggesting that the NLS approach provides a poor approximation of the average behaviour. In fact, for two datasets, $\mathbf{p}^{\text{NLS}}$ lies outside the 95% NLME prediction region—confirmation that the NLS estimate *can fail to capture the behaviour of any cell*. We conclude that the NLS approach may not be suited for parameterising functions using normalised steady-state inactivation data, especially when there is significant variability among cells.

Population variability—One major advantage of the NLME approach is that it enables quantification of population variability. Considering the 95% NLME predictions regions, we observe that except for the *Cord-Epi* dataset results, there is significant variability in $\mathbf{p}_I$, especially in V_0 , across the cell population. This could be seen in Figure 2, but Figure 4 quantifies the variability. As far as we are aware this is the first such quantification of population uncertainty, based on the underlying experimentally observed variability, for a component of a cardiac cell model. All the information on this characterised uncertainty is provided in Table 1.

Epicardial vs endocardial cardio-myocytes—Next, and again considering only the NLME results, we can compare the epicardial results with the endocardial results. The following is true for both the Cordeiro or Murphy datasets: the epicardial estimates of both V_0 and K are smaller than endocardial values, and the respective 95% confidence regions for

p do not intersect, suggesting a statistically significant difference in steady-state I_{Na} inactivation between epicardial and endocardial cells. This conclusion is consistent with the finding of [8], but involves an alternative (and given the above, arguably better) method for comparing the experimental results ([8] directly compared averaged experimental results). Whilst absolute values differ between the Cordeiro and Murphy datasets (see below), it is interesting to note that the difference in p^{NLME} between epi and endo cells is similar for the two studies (see Table 1). These last observations are only possible because of the novel application of the statistical method (NLME) that we have employed. An NLS-based analysis would conclude that V_0 (related to half-inactivation) is greater for endocardial cells, consistently in both studies, but would not lead to strong conclusions on K (related to slope). Our NLME-based analysis suggests that both half-inactivation is greater *and* the slope is smaller for endocardial cells, consistently in both studies.

Cordeiro *et al.* 2008 [8] datasets vs Murphy *et al.* 2011 [25] datasets—Finally, comparison of the Cordeiro results with the Murphy results leads to troubling conclusions on experimental reproducibility. Noticeable differences are observed in both the average behaviour and amount of variability for the Cordeiro data compared to the Murphy data. In particular V_0 for both epi and endocardial cells are not similar between the studies. It is worth noting though that estimated value of K is very similar between the studies, for both epi and endocardial cells.

3. Uncertainty in cell model predictions

We have characterised the uncertainty in the parameters of the function chosen to represent the steady-state inactivation of I_{Na} , albeit separately for the different datasets. We now investigate the extent to which this uncertainty affects predictions in a cell model. We will use the Fox *et al.* 2002 canine cell model [13], which is a midmyocardial model. In Section 3.1 we discuss some model simplification, before analysing predictions of upstroke velocity (Section 3.2), action potential duration (Section 3.3) and spiral wave dynamics (Section 3.4).

3.1. Model simplification

The fast sodium current in the Fox 2002 model is given by (3), with two inactivation gates h and j . The dynamics of h and j are prescribed through the specification of four functions α_h , β_h , α_j , β_j , which for this model take the form

$$\alpha_h(V) = 0.135e^{-0.14706(V+80)},$$

$$\beta_h(V) = \frac{7.5}{1 + e^{-0.1(V+11)}},$$

$$\alpha_j(V) = \frac{0.175e^{-0.04348(V+100)}}{1+e^{0.15(V+79)}}, \quad (9)$$

$$\beta_j(V) = \frac{0.3}{1+e^{-0.1(V+32)}}.$$

The Fox 2002 cell model therefore uses 5 independent⁴ parameters to represent h -dynamics and 7 for j -dynamics. (This is relatively little for a cardiac cell model—for example the Hund-Rudy 2004 model [18] uses 12 and 15 parameters for h - and j -dynamics respectively). For Fox 2002, the resultant steady-state functions $h_\infty = \alpha_h/(\alpha_h + \beta_h)$ and $j_\infty = \alpha_j/(\alpha_j + \beta_j)$ are actually very similar (as shown shortly), but since they contain 14 parameters in total, are not readily amenable to UQ. In fact, these parameters are highly unlikely to be *identifiable*, that is, uniquely determined from experimental data, and therefore some model simplification is likely required for uncertainty quantification.

We create a ‘simplified’⁵ Fox 2002 model by replacing the $h_\infty(V)$ and $j_\infty(V)$ formulations with the two-parameter model. Specifically, we set

$$h_\infty(V) = j_\infty(V) = F_1(V; V_0, K). \quad (10)$$

The timescales $\tau_h(V)$ and $\tau_j(V)$ are not altered. This ‘simplified’ Fox model has two parameters, (V_0, K) , to represent steady-state inactivation. The total steady-state inactivation is then $h_\infty(V) \times j_\infty(V) = F_1^2(V) = F_2(V)$, consistent with the fits performed in Section 2. We may therefore use the results of that section in this simplified Fox 2002 model to propagate forward the uncertainty in (V_0, K) and determine the uncertainty in cell model outputs.

We note that there is a choice of V_0 and K such that the simplified Fox 2002 model (10) gives virtually the same results as the original full model (9). Setting $V_0^F = -68.3862\text{mV}$ and $K^F = 4.0448\text{mV}$, if we choose $(V_0, K) = (V_0^F, K^F)$, the steady-state inactivation curves of the simplified model are almost identical to those in the full model (Figure 5 (left)). Figure 5 (right) plots an action potential obtained using the simplified and original models (computed by applying a 2ms stimulus of magnitude 1.1 times the threshold stimulus, with initial conditions taken from the CellML description of the model, using the software package Chaste [23]). The change to steady-state I_{Na} inactivation has negligible effect on the action potential—the maximum difference in transmembrane voltage over the entire AP is only 0.05mV. We will consider the choice $(V_0, K) = (V_0^F, K^F)$ as being equivalent to using the original model, and refer to V_0^F and K^F as the values ‘used in the cell model’, even though the cell model actually uses a more complex, but essentially equivalent, function.

⁴Taking into account that $ae^{b(V+c)} = de^{bV}$ with $d = Ae^{bc}$.

⁵It should be noted that whilst the inactivation formulation has been simplified, the overall model has not, since the original 14 parameters are still present in the definitions of the time constants τ_h and τ_j . It may be possible to re-formulate these using functions with less parameters, while not altering results. However, this is not relevant to this work.

In the following sections we investigate how the quantified variability of steady-state I_{Na} inactivation affect model results. In these sections we will only use the NLME prediction regions (and underlying probability distributions), which as discussed before represent *population variability*. A comment regarding NLS fits is provided at the end of Section 3.2.

3.2. Upstroke velocity (single cell)

We can use the probability distributions computed in Section 2 (Table 1) to compute the resultant probability distributions for the upstroke velocity, V_{max} . We use a simple brute force method for this calculation, using numerical integration on a 100 by 100 grid in parameter space (see supplemental material for details). Figure 6 (top-left) plots upstroke velocity as a function of the two parameters V_0 and K , and overlays the 95% NLME prediction regions for each of the four datasets, as well as the point (V_0^F, K^F) . Figure 6 (top-right) plots the corresponding probability distributions for upstroke velocity, together with the upstroke velocity provided by the original model (corresponding to (V_0^F, K^F)), as well as experimental measurements of upstroke velocity in intact tissue taken from Cordeiro *et al.* 2008 [8] (188.6 ± 15.5 mV/ms for epicardial tissue, 228.1 ± 18.8 mV/ms for endocardial tissue).

The computed probability distributions are quite wide, between them covering the range ~ 240 mV/ms to ~ 350 mV/ms, although they do overlap to some extent. The figure illustrates the difference between a traditional parameter-fit with no characterisation of uncertainty and the present approach, and shows how uncertainty quantification improves the ability to compare simulation results with experiment—i.e. to perform validation. Whilst the experimental measurements display the same epi/endo trends as the model, we see that both experimental results fall in the tails of the model probability distributions. The experimental measurements of upstroke velocity were made using intact tissue, and therefore expected to be smaller than in the isolated cell (see Section 4.1 for more discussion). The important point is that the uncertainty quantification enables a much more informed comparison of simulation and experiment.

For completeness, this analysis is repeated in the supplemental material using the maximum action potential amplitude, V_{max} , as the quantity of interest, and we note that time of maximum upstroke velocity, and resting potential, were both found to be essentially independent of V_0 and K .

We also note that, for any given dataset, there is little difference between upstroke velocity when computed using $\mathbf{p}^{NLS}$ compared to $\mathbf{p}^{NLME}$ (results not presented). This is because V_{max} is largely independent of K . This type of uncertainty propagation is required to determine if the differences in fitted values is significant for specific quantities of interest. Hence, for upstroke velocity as the quantities of interest an NLS fit to the underlying data would provide very similar results to an NLME fit. However, the fact that NLME enables quantification of population variability illustrates its general superiority over NLS.

3.3. Action potential duration (single cell)

We now repeat this process using the action potential duration (APD) as the quantity of interest. We measure APD₉₀, the time to 90% repolarisation. It might be expected that steady-state I_{Na} inactivation has little effect on APD. However, it turns out that for a large region in parameter space, no APD can be computed, as the model fails to fully repolarise, which is clearly unphysiological. This occurs when V_0 is at greater potentials, which increases the I_{Na} ‘window’ current and can lead to additional steady state solutions of the single cell ODEs [1]. Figure 6 (bottom panel, top left) plots APD as a function of V_0 and K , with the region of ‘model failure’ coloured in black. In the region where APD can be computed, it does not vary greatly with (V_0, K) (as expected)—it is only on the border of the failure region where it changes significantly. Crucially, the 95% prediction regions overlap with this failure region for three of the four experimental datasets. Thus, the uncertainty quantification *exposes the fact that with this model APD is highly sensitive to steady-state I_{Na} inactivation, and not robust to the true underlying variability*. We should note though that the Fox 2002 cell model was not developed to take into account transmural differences, and this lack of robustness could be due to mixing an individual-based approach for steady-state I_{Na} inactivation with a population-averaged approach for the rest of the model.

Since APD is not defined everywhere in the parameter-space, we cannot compute probability distributions for APD as was carried out for upstroke velocity. Figure 6 (bottom panel, top right) instead plots 95% prediction *bands* for the action potential, obtained by sampling over points within the 95% prediction regions and computing the action potential. The variability in the *Cord-Epi* dataset leads to little uncertainty in action potential. However, for the other datasets, there is large uncertainty in phases 3 and 4 of the AP, and with the *Murphy-Endo* dataset there is also some phase 2 uncertainty.

3.4. Spiral wave dynamics (tissue)

Finally, we consider the effect of uncertainty in steady-state I_{Na} inactivation on spiral wave dynamics in 2D. We solve the monodomain equations (see e.g. [20]) governing spatio-temporal propagation, using our ‘simplified’ Fox 2002 cell model, on a 2D planar domain of size 10 cm by 10 cm (regular grid with spatial stepsize $h = 0.02$ cm), using the software package Chaste [23]. These simulations are computationally expensive, and we only consider a small number of parameter choices, rather than propagating forward full distributions.

First, using the model parameters (V_0^F, K^F) (and the same initial conditions as in the single-cell simulations), a spiral wave was induced by stimulating the right-half of the domain while imposing a gradient in initial h in the y -direction, and the system then allowed to evolve for 1900ms for a spiral wave to develop. At $t = 1900$ ms the parameters were altered from (V_0^F, K^F) to the parameter choice being tested, and the system allowed to evolve until $t = 4000$ ms. Six parameter choices were compared: the control (V_0^F, K^F) (i.e. no change); (V_0^{NLME}, K^{NLME}) for each of the four datasets (values in Table 1); and $(-50, 9)$, as a value for which APD is infinite in the single-cell simulation (see Figure 6), which we refer to as a single-cell failure (SCF) parameter choice.

The trajectories of spiral wave tips were computed as done previously [16]. Briefly, a cardiac phase variable was computed at each site from reconstructed state space as

$$\theta(x, y, t) = \arctan \left(\frac{V(x, y, t + \tau) - V_{\text{ref}}}{V(x, y, t) - V_{\text{ref}}} \right)$$

(using an arctan function that takes quadrant into account, i.e. ‘atan2’), where $V_{\text{ref}} = -30$ mV and $\tau = 10$ ms, and the line integral

$$\oint \nabla \theta \cdot d\mathbf{l}$$

was then computed around each pixel using the surrounding 8 pixels. The integral is only non-zero (taking values $\pm 2\pi$) when a phase singularity is present.

Figure 7 (top) illustrates all six parameter choices on top of single-cell APD90 as a function of (V_0, K) . The model and NLME parameter choices all lie in the finite-APD region, although the *Murphy-Endo* choice lies on the border of the model-failure region (APD90 equal to 221ms for *Murphy-Endo*, compared to 184ms for the others).

2D spiral wave results are presented in Figure 7 (bottom). Somewhat surprisingly, the 2D spiral wave results for each of the five simulations excluding the SCF simulation were *virtually identical*. Figure 7 (bottom left) plots the spiral wave tip trajectory for $t = 1900$ ms using the *Cord-Endo* parameter choice. The trajectories using the other parameter choices excluding SCF are indistinguishable if also plotted. All the NLME trajectories were within 0.06 cm of the control trajectory at each timestep, and the proportion of timesteps for which the algorithm returned an identical tip location to the control was: 99% for *Cord-Endo*, 97% for *Cord-Epi*, 86% *Murphy-Endo*, and 99% for *Murphy-Epi*. Similarly, there were almost no differences in voltage across all spatial locations and timesteps⁶. For the SCF choice, no prolonged action potentials occurred, and a spiral wave did occur, following a similar but not identical tip trajectory to the control, as shown in Figure 7 (bottom right) (trajectory within 1.01 cm of control, identical to control for first 9%).

It is remarkable that these simulation results at the tissue-scale are so similar, given that upstroke velocity and APD at the single-cell scale were sensitive to the parameter values. This observation, which is a promising result in the context of quantifying uncertainty for clinical applications, may be an illustration of a fundamental property in systems biology: inherent robustness to natural variability in function across multiple scales, or *emergent robustness*. However, more detailed analysis is certainly needed.

⁶Root-mean-squared difference in voltage (average taken over all spatial locations and timesteps) compared to control: 0.075 mV for *Cord-Endo*, 0.16 mV for *Cord-Epi*, 0.68 mV for *Murphy-Endo*, 0.065 mV for *Murphy-Epi*. Max difference in voltage (over all spatial locations and timesteps) compared to control: 9.6 mV for *Cord-Endo*, 26.0 mV for *Cord-Epi*, 52.9 mV for *Murphy-Endo*, 14.1 mV for *Murphy-Epi*. The fact that the maximum difference values are less than the difference between activated and resting potential shows that wavefront location is almost identical at all times; overall the differences are miniscule.

4. Discussion

In this paper we have presented a methodology to perform uncertainty quantification (UQ) of steady-state I_{Na} inactivation in computational models of cardiac electrophysiology. Multi-scale mechanistic models in physiology, such as models of the heart, tend to be extremely complex, and integrating uncertainty into such models is a huge challenge. In fact, it is unclear the extent to which UQ is possible for whole heart electrophysiological simulations as typically used in simulation studies. To make the problem tractable, we have restricted ourselves to uncertainty in steady-state I_{Na} inactivation only. As far as are aware, this is the first formal UQ analysis in cardiac modelling and therefore much insight is provided by this work, as well as numerous open questions, as discussed in this section.

4.1. Insight from application of framework to steady state I_{Na} inactivation

We have applied an ‘individual-based’ statistical methodology, nonlinear mixed effects (NLME) modelling, to analyse voltage clamp data taken from a population of cells, and contrasted this approach with a traditional approach of presenting data averaged over cells and fitting curves using a least squares minimisation. NLME is commonly used in pharmacokinetic applications, but as far we are aware has never been used to analyse voltage clamp data. The individual-based approach better characterises the properties of an ‘average cell’, as well as providing a framework with which to quantify population variability. The traditional approach was seen to underestimate slope of steady-state inactivation when variability was not small. We believe NLME will prove to be a powerful technique for analysing physiological data and in future work we aim to apply it to data on I_{Na} activation.

The statistical methods were applied to four datasets, including both epicardial and endocardial cardio-myocytes. The results presented in Figure 4 are highly informative, illustrating population variability (which is quantified in the form of probability distributions; see Table 1) and differences across experiments. Our results support the findings of [8] in suggesting there exists a difference in steady-state I_{Na} inactivation between epi and endocardial cells, and specifically we have shown that endocardial cells exhibit decreased slope as well as greater half-inactivation. The observation that the slope of steady state I_{Na} inactivation may be dependent on transmural location was only possible because of our use of the individual-based method.

The fact that the results using the Cordeiro *et al.* [8] datasets differed to those using the Murphy *et al.* [25] datasets is somewhat surprising, since all the data was collected in the same laboratory under similar conditions, although the effects of gender (unknown) and seasonal/yearly variations may play a role. Whilst reproducibility of experimental data is well beyond the scope of this study, it is clearly of paramount importance in regards to VVUQ. The best approach to deal with conflicting data within the UQ framework is not clear. Once again, an individual-based approach may prove beneficial in addressing such experimental variability.

The parameter uncertainty for each of the four datasets were propagated through the Fox 2002 cell model [13], and we computed probability distributions for model outputs V_{max} ,

$V_{\max}$ and V_{rest} , as well as analysing the effect on APD and spiral wave dynamics. As expected, we found that the choice of the inactivation parameters (V_0 and K) affects $V_{\max}$ (Figure 6) and $V_{\max}$ (supplemental material), and we found that model predictions for these quantities of interest are consistent with relative differences between epi and endocardial canine cells found experimentally (although absolute values do not match experiment, as discussed further below). As anticipated, (V_0 , K) did not affect V_{rest} . In general it is important to explore how model parameters affect important model outputs even when an effect is not expected, and this is illustrated in Figure 6, in which APD was shown to become unphysiologically large/infinite for a large region in parameter space. Although the parameters values used in the model, (V_0^F, K^F) , lie outside this region, the 95% prediction regions corresponding to population variability intersect this region for three of the four datasets. However, despite this lack of robustness to ‘real’ biological variability in simulations of the isolated cell, the underlying uncertainty had surprisingly little effect on spiral wave dynamics. Overall, the UQ framework has been shown to improve the ability to perform validation, and identify robustness or lack thereof of model predictions to natural variability.

The absolute values of $V_{\max}$ and $V_{\max}$ differing between experiment and simulation may be due to comparison of action potentials recorded in tissue with single-cell simulations. Another factor may be various ‘conditions’ assumed by the model that do not match the experimental environment exactly. For cardiac models, some of the important conditions to consider are temperature and concentrations of intra- and extracellular ions. It should be noted that the conditions used to record peak I_{Na} (the voltage clamp data, i.e. our ‘calibration data’) and action potentials (our ‘validation data’) in the experimental studies were quite different [8, 25]. While action potentials were recorded from canine wedge in normal external Na^+ at 36 °C, peak I_{Na} had to be recorded at room temperature in greatly reduced external Na^+ in isolated myocytes to ensure adequate voltage control. These condition deviate from the model conditions and could be partly responsible for the differences between model predictions and experimental results of $V_{\max}$ and $V_{\max}$.

4.2. Uncertainty quantification in whole-heart modelling

Overall, this article represents a first step towards the ultimate aim of whole-heart simulations with all appropriate uncertainties (relevant to the specific quantities of interest to be used in decision-making) rigorously accounted for. Here we studied the steady-state inactivation of one current (I_{Na}), assuming that $h_{\infty}(V) = j_{\infty}(V)$ and using four datasets (acquired at low temperature and external sodium concentration). This approach can be directly applied to other voltage-dependent gating currents.

Quantifying uncertainty for all aspects of each membrane current in a complete physiological cell model, and propagating forward this uncertainty in whole-heart simulations, may be difficult or impossible with current technology, due to the large number of parameters in such models and the high computational cost of whole-heart simulations. However, the following approach may be possible: (i) perform initial model simplification (such as in Section 3.1), to reduce the number of parameters in an overall model; (ii) use both physiological insight and sensitivity analyses, to estimate/determine which parameters

in the simplified model heavily impact a particular quantity of interest; and (iii) perform UQ analyses using the framework developed in this paper, on simplified situations such as performed in this paper, together with limited uncertainty propagation using full whole-heart simulations, to gauge the expected impact of relevant uncertainties in whole-heart simulations. For this process (which can be regarded as an intermediate approach between no UQ and fully rigorous UQ), and for the long-term aim in which all appropriate uncertainties are rigorously quantified and propagated, numerous challenges need to be overcome, which we now discuss. We believe a hierarchical and iterative approach to tackle these problems will be needed.

First, our methodology should be applied to the activation and slow inactivation of the sodium channel, as well as the corresponding time constants. The methodology is directly applicable to all voltage-gated variables; however, simultaneously quantifying uncertainty of all gating variables for a given current may require extension of our methodology. Typically, time constants are fitted using more complex functions than the sigmoidal functions used for steady-state activation/inactivation, and experimental measurement of time constants can be difficult, which will likely result in a larger number of parameters and greater experimental variability. Still, since such analyses can be performed in parallel, quantifying uncertainty for various current models is feasible. Quantifying uncertainty in the rapid delayed rectifying current, I_{Kr} (the hERG channel), may in particular provide important insight.

Propagating forward quantified uncertainty remains a major challenge, as there may be a large number of uncertain parameters to include in the analysis. Model simplification (discussed further below) should be considered to reduce the number of model parameters in advance of the analysis, and as mentioned above physiological insight together with sensitivity analyses may be required to decide which parameters to include in the uncertainty propagation. Even so, advanced methods for performing uncertainty propagation in high dimensions may be required, such as the Monte-Carlo or Markov Chain Monte-Carlo method, Gaussian Process emulator based methods [32, 21], or generalised polynomial chaos based methods, including Galerkin and collocation polynomial chaos approaches [41].

Model simplification, such as that performed in Section 3.1, is clearly a vital first stage of the overall process. We envision that hybrid physiological-phenomenological models (i.e. physiological models for some currents of particular importance, simplified phenomenological models for the sum of all other currents) may provide a means for limiting the number of model parameters. Such models may be less realistic than wholly physiological models, but a simplified model with UQ may be far more useful than a complex model with no UQ [26]. Model simplification is related to the question of parameter identifiability and over-parameterisation (see e.g. [7, 11]). This is an important issue with larger models, which we touched upon in Section 3.1 where we showed that the two-parameter model (10) of steady state inactivation gives essentially the same results as the 12 parameter version in the original Fox *et al.* model. This suggests that the 12 parameter version is over-parameterised, or equivalently, that the 12 parameters are not (locally) identifiable with respect to the steady state inactivation function. Identifiability analysis will be an important tool in the process of model simplification and uncertainty

quantification, since parameters are estimated, and parameters that are not identifiable are not estimable.

Finally, optimal experimental design techniques [10] may be used to identify the experimental frameworks and statistical estimation procedures that render the most precise and efficient estimates of model parameters. This paper is focused on natural cell-to-cell variability, which cannot be reduced through design. However, the quality of estimates of this variability, and the complexity of the parameterisations that can be assessed in this framework, are affected by experimental design. For example, optimal design methods may be used to optimise the sequence of test voltages in voltage clamp experiments (i.e. sample more densely among the voltages that are most informative about the model parameters, as opposed to voltage steps at fixed increments as in Figure 2). As a further example, the optimal design framework may aid in evaluating and maximising the information that arises under different types of experiments, especially when these data are used simultaneously to estimate model parameters.

4.3. Conclusion

Cardiac EP modelling is one of the success stories in the application of the multi-scale systems biology approach, but even here, as we have shown, there has to date been comparatively little work in developing our understanding of how underlying uncertainty/variability affects predictions. With such complex multi-scale models, advanced statistical methodologies (together with new ways of assessing experimental results and developing models) provide the potential for assessing and integrating such uncertainty, but huge challenges still remain. The approach described in this article has potential application across multi-scale systems biology, but the challenges we have exposed remain common to all such applications.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

The authors would like to thank Leonardo Angelone and Tina M. Morrison (Center for Devices and Radiological Health, FDA) for their feedback. PP would like to acknowledge support from the UK EPSRC-funded '2020 Science' programme (grant number: EP/I017909/1). MS would like to acknowledge grant NIH/NHLBI 1R01HL118392-01. Tissues simulations were carried out using FDA supercomputing resources, and PP would like to thank Mike Mikailov (FDA Supercomputing Laboratory) for technical assistance.

References

1. Angelino E, Brenner M. Excitability constraints on voltage-gated sodium channels. *PLoS Computational Biology*. 2007; 3(9):e177.
2. Antzelevitch C, Burashnikov A, Sicouri S, Belardinelli L. Electrophysiologic basis for the antiarrhythmic actions of ranolazine. *Heart Rhythm*. 2011; 8(8):1281–1290. [PubMed: 21421082]
3. Ashikaga H, Arevalo H, Vadakkumpadan F, Blake R, Bayer J, Nazarian S, Muz Zviman M, Tandri H, Berger R, Calkins H, Herzka D, Trayanova N, Halperin H. Feasibility of image-based simulation to estimate ablation target in human ventricular arrhythmia. *Heart Rhythm*. 2013; 10(8):1109–1116. [PubMed: 23608593]

4. Bates, D.; Watts, D. *Nonlinear Regression Analysis and Its Applications* (Second Edition). Wiley-Interscience; 1988.
5. Bishop M, Plank G, Burton R, Schneider J, Gavaghan D, Grau V, Kohl P. Development of an anatomically detailed MRI-derived rabbit ventricular model and assessment of its impact on simulations of electrophysiological function. *American Journal of Physiology—Heart and Circulatory Physiology*. 2010; 298(2):H699. [PubMed: 19933417]
6. Burashnikov A, Di Diego J, Zygmunt A, Belardinelli L, Antzelevitch C. Atrium-selective sodium channel block as a strategy for suppression of atrial fibrillation differences in sodium channel inactivation between atria and ventricles and the role of ranolazine. *Circulation*. 2007; 116(13): 1449–1457. [PubMed: 17785620]
7. Cserecsik D, Hangos K, Szederkényi G. Identifiability analysis and parameter estimation of a single Hodgkin–Huxley type voltage dependent ion channel under voltage step measurement conditions. *Neurocomputing*. 2012; 77:178–188.
8. Cordeiro J, Mazza M, Goodrow R, Ulahannan N, Antzelevitch C, Di Diego J. Functionally distinct sodium channels in ventricular epicardial and endocardial cells contribute to a greater sensitivity of the epicardium to electrical depression. *American Journal of Physiology-Heart and Circulatory Physiology*. 2008; 295(1):H154–H162. [PubMed: 18456729]
9. Elkins R, Davies M, Brough S, Gavaghan D, Cui Y, Abi-Gerges N, Mirams G. Variability in high-throughput ion-channel screening data and consequences for cardiac safety assessment. *Journal of pharmacological and toxicological methods*. 2013; 66:112–122. [PubMed: 23651875]
10. Fedorov, V.; Leonov, S. *Optimal design for nonlinear response models*. CRC Press; 2013.
11. Fink M, Noble D. Markov models for ion channels: versatility versus Identifiability and speed. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 2009; 367(1896):2161–2179.
12. Fitzmaurice, G.; Davidian, M.; Verbeke, G.; Molenberghs, G. *Longitudinal data analysis*. CRC Press; 2008.
13. Fox J, McHarg J, Gilmour R Jr. Ionic mechanism of electrical alternans. *American Journal of Physiology—Heart and Circulatory Physiology*. 2002; 282:H516–H530. [PubMed: 11788399]
14. Gemmell P, Burrage K, Rodriguez B, Quinn T. Quantifying importance of repolarization currents underlying physiological variability in rabbit cellular electrophysiology. *Biophysical Journal*. 2012; 102 Supplement 1(3) 540a –.
15. Fox J, McHarg J, Gilmour R Jr, Geneser S, Kirby R, Xiu D, Sachse F. Stochastic Markovian modeling of electrophysiology of ion channels: Reconstruction of standard deviations in macroscopic currents. *Journal of theoretical biology*. 2007; 245:627–637. [PubMed: 17204291]
16. Gray R, Pertsov A, Jalife J. Spatial and temporal organization during cardiac fibrillation. *Nature*. 1998; 392(6671):75–78. [PubMed: 9510249]
17. Hodgkin A, Huxley A. A quantitative description of membrane current and its application to conduction and excitation in the nerve. *J Physiol*. 1952; 117:500–544. [PubMed: 12991237]
18. Hund T, Rudy Y. Rate dependence and regulation of action potential and calcium transient in a canine cardiac ventricular cell model. *Circulation*. 2004; 110(20):3168–3174. [PubMed: 15505083]
19. Hunter P, Crampin E, Nielsen P. Bioinformatics, multi-scale modeling and the IUPS physiome project. *Briefings in bioinformatics*. 2008; 9(4):333–343. [PubMed: 18477639]
20. Keener, J.; Sneyd, J. *Mathematical Physiology*, volume 8 of *Interdisciplinary Applied Mathematics*. Springer; 1998.
21. Kennedy M, O’Hagan A. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2001; 63(3):425–464.
22. Kiureghian A, Ditlevsen O. Aleatory or epistemic? Does it matter? *Structural Safety*. 2009; 31(2): 105–112.
23. Mirams G, Arthurs C, Bernabeu M, Bordas R, Cooper J, Corrias A, Davit Y, Dunn S-J, Fletcher A, Harvey D, Marsh M, Osborne J, Pathmanathan P, Pitt-Francis J, Southern J, Zemzemi N, Gavaghan D. Chaste: An open source C++ library for computational physiology and biology. *PLoS Computational Biology*. 2013; 9(3):e1002970. [PubMed: 23516352]

24. Mirams G, Davies M, Cui Y, Kohl P, Noble D. Application of cardiac electrophysiology simulations to pro-arrhythmic safety testing. *British Journal of Pharmacology*. 2012; 167(5):932–945. [PubMed: 22568589]
25. Murphy L, Renodin D, Antzelevitch C, Di Diego J, Cordeiro J. Extracellular proton depression of peak and late Na⁺ current in the canine left ventricle. *American Journal of Physiology—Heart and Circulatory Physiology*. 2011; 301(3):H936–H944. [PubMed: 21685271]
26. National Research Council. Assessing the Reliability of Complex Models: Mathematical and Statistical Foundations of Verification, Validation, and Uncertainty Quantification. The National Academies Press; 2012.
27. Nesterenko V, Zygmunt A, Rajamani S, Belardinelli L, Antzelevitch C. Mechanisms of atrial-selective block of Na⁺ channels by ranolazine: II. Insights from a mathematical model. *American Journal of Physiology—Heart and Circulatory Physiology*. 2011; 301(4):H1615–H1624. [PubMed: 21821780]
28. Niederer S, Fink M, Noble D, Smith N. A meta-analysis of cardiac electrophysiology computational models. *Experimental physiology*. 2009; 94(5):486–495. [PubMed: 19139063]
29. Noble D. A modification of the Hodgkin-Huxley equations applicable to Purkinje fibre action and pacemaker potentials. *J. Physiol.* 1962; 160:317–352. [PubMed: 14480151]
30. Noble D, Noble P. Late sodium current in the pathophysiology of cardiovascular disease: consequences of sodium–calcium overload. *Heart*. 2006; 92:iv1–iv5. [PubMed: 16775091]
31. Moreno J, Clancy C. Pathophysiology of the cardiac late Na current and its potential as a drug target. *Journal of molecular and cellular cardiology*. 2012; 52(3):608–619. [PubMed: 22198344]
32. Oakley J, O'Hagan A. A Bayesian inference for the uncertainty distribution of computer model outputs. *Biometrika*. 2002; 89(4):769–784.
33. O'Hara T, Virág L, Varró A, Rudy Y. Simulation of the undiseased human cardiac ventricular action potential: model formulation and experimental validation. *PLoS Computational Biology*. 2011; 7(5):e1002061. [PubMed: 21637795]
34. Pathmanathan P, Gray R. Ensuring reliability of safety-critical clinical applications of computational cardiac models. *Frontiers in physiology*. 2013; 4
35. Pathmanathan P, Gray R. Verification of computational models of cardiac electrophysiology. *International journal for numerical methods in biomedical engineering*. 2014; 30(5):525–544. [PubMed: 24259465]
36. Pinheiro, J.; Bates, D. Mixed effects models in S and S-PLUS. Springer; 2000.
37. Roy C, Oberkampf W. A complete framework for verification, validation, and uncertainty quantification in scientific computing. 48th AIAA Aerospace Sciences Meeting and Exhibit. 2010:2010–0124.
38. Sarkar A, Christini D, Sobie E. Exploiting mathematical models to illuminate electrophysiological variability between individuals. *The Journal of physiology*. 2012; 590(11):2555–2567. [PubMed: 22495591]
39. ten Tusscher K, Noble D, Noble P, Panfilov A. A model for human ventricular tissue. *American Journal of Physiology—Heart and Circulatory Physiology*. 2004; 286(4):1573–1589.
40. Trayanova N. Computational cardiology: The heart of the matter. *ISRN cardiology*. 2012; 2012
41. Xiu, D. Numerical methods for stochastic computations: a spectral method approach. Princeton University Press; 2010.

Figure 1. Components of a multi-scale model of the heart

Top centre: a schematic of the channels, pumps and exchangers modelled in a (relatively simple) cell model. Inactivation dynamics of the fast sodium current can be experimentally measured using the illustrated voltage clamp protocol (top left), and curves fitted to experimental results (bottom left) and embedded in the cell model. Cell models can be used to predict quantities such as action potential (bottom centre), or they may be coupled to partial differential equations governing spatial propagation [20], and, together with a computational mesh of the heart (top right: a high-resolution mesh of a rabbit heart comprised of 21 million elements [5]), used to simulate whole-heart electrophysiological activity (bottom right). [Colour in online version].

Figure 2.

Steady-state I_{Na} inactivation for endocardial (left) and epicardial (right) cells from datasets used in Cordeiro *et al.* 2008 [8] (top) and Murphy *et al.* 2011 [25] (bottom). Small gray stars correspond to individual cells, large black circles to the average inactivation at each voltage. (Note: lines connect data points, they are not fits to the data).

Figure 3.

Weighted nonlinear least squares (NLS) and nonlinear mixed effects (NLME) fits using the Cordeiro *et al.* 2008 datasets. Left: experimental observations (thin dashed lines), average of experimental observations at each voltage (circles), and NLS and NLME fits (thick dashed & solid lines). Right: experimental observations (circles), and individual NLME fits (solid lines). [Colour in online version].

Cordeiro et al. 2008 endo **Cordeiro et al. 2008 epi** **Murphy et al. 2011 endo** **Murphy et al. 2011 epi** **Figure 4.**

Confidence and prediction regions in parameter space using each of the four datasets.

Plotted are $\mathbf{p}^{\text{NLS}}$ and $\mathbf{p}^{\text{NLME}}$ (filled circles), 95% confidence regions for $\mathbf{p}^{\text{NLS}}$ and $\mathbf{p}^{\text{NLME}}$ (solid lines), the NLME 95% prediction regions (dashed lines) which represent *population variability*, and estimates of $\mathbf{p}_I$ for each individual cell of the dataset (stars). [Colour in online version].

Figure 5.

Left: Original $h_\infty(V)$ and $j_\infty(V)$ in Fox 2002 (dashed lines; $h_\infty(V)$ indistinguishable from solid curve, $j_\infty(V)$ almost indistinguishable from solid curve), and new $h_\infty(V) = j_\infty(V)$ curve (solid line). Right: Action potential with original (dashed line) and new (solid line) $h_\infty(V)$ and $j_\infty(V)$ (APs indistinguishable—max difference is 0.05mV)

Figure 6.

Propagation of uncertainty through a cardiac cell model. Top panel: upstroke velocity (Section 3.2); bottom panel: action potential duration (APD) (Section 3.3). Upstroke velocity panel, left: upstroke velocity as a function of the two parameters V_0 and K , together with NLME 95% prediction regions (for legend see top-right figure), and (V_0^F, K^F) , the parameter choice used in the model (filled circle). Upstroke velocity panel, right: corresponding probability distributions for upstroke velocity, as well as value computed using (V_0^F, K^F) and experimental measurements (filled circles). APD panel, top left: APD90

as a function of the two parameters V_0 and K , together with 95% prediction regions and (V_0^F, K^F) (filled circle). The black region indicates parameter values such that the model does not repolarise. APD panel, bottom: sample action potentials for different regions of parameter space. APD panel, top right: 95% action potential bands, corresponding to the 95% prediction regions. [*Colour in online version*].

Figure 7.

Top: Several points in parameter space, against a backdrop of (single cell) APD. Plotted are parameter values (V_0^F, K^F) used by the Fox 2002 model (labelled 'Model'), together with four points (V_0^{NLME}, K^{NLME}) computed using each of the four datasets, and the point $(-50, 9)$ arbitrarily chosen in the 'failure region' (labelled SCF for 'single-cell failure'). Bottom left: spiral wave simulation on a 10 cm by 10 cm 2D domain as described in text, using the *Cord-Endo* parameter choice (background colour: trans-membrane voltage at $t = 1900$ ms; solid line: trajectory of spiral wave tip for following 2000 ms). Identical tip trajectories are

obtained with all parameter choices except SCF. Bottom right: comparison of SCF trajectory vs trajectory with other parameter choices. [*Colour in online version*].

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 1

Quantified population variability for steady-state I_{Na} inactivation, using each of the four datasets. Each cell is assumed to have parameters (V_0, K) such that $(V_0, \log(K))$ is drawn from a bivariate normal distribution $N((V_0^{NLME}, \log(K^{NLME})), \Psi)$. The covariance matrices are symmetric 2×2 matrices; provided below are the three independent components in the form $[\Psi_{11}, \Psi_{12}, \Psi_{22}]$. Note that K is used in the second column but $\log(K)$ is used in the covariance matrices.

	V_0^{NLME}	K^{NLME}	$\Psi = \text{cov}(V_0, \log(K))$
<i>Cord-Endo</i>	-69.3	6.80	[54.5, 0.248, 0.00671]
<i>Cord-Epi</i>	-78.7	5.93	[5.04, 0.0561, 0.00416]
<i>Murphy-Endo</i>	-59.7	7.30	[23.9, 0.243, 0.0173]
<i>Murphy-Epi</i>	-73.1	5.91	[49.5, 0.804, 0.0161]