

Optimal experimental designs for characterising ion channel gating by filling the phase-voltage space of model dynamics

Gary R. Mirams, Michael Clerx, Dominic G. Whittaker & Chon Lok Lei

To cite this article: Gary R. Mirams, Michael Clerx, Dominic G. Whittaker & Chon Lok Lei (2024) Optimal experimental designs for characterising ion channel gating by filling the phase-voltage space of model dynamics, *Mathematics in Medical and Life Sciences*, 1:1, 2375494, DOI: [10.1080/29937574.2024.2375494](https://doi.org/10.1080/29937574.2024.2375494)

To link to this article: <https://doi.org/10.1080/29937574.2024.2375494>



© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



[View supplementary material](#)



Published online: 23 Jul 2024.



[Submit your article to this journal](#)



Article views: 129



[View related articles](#)



[View Crossmark data](#)

Optimal experimental designs for characterising ion channel gating by filling the phase-voltage space of model dynamics

Gary R. Mirams ^a, Michael Clerx ^a, Dominic G. Whittaker ^a and Chon Lok Lei ^{b,c}

^aCentre for Mathematical Medicine & Biology, School of Mathematical Sciences, University of Nottingham, Nottingham, UK; ^bInstitute of Translational Medicine, Faculty of Health Sciences, University of Macau, Macau SAR, China; ^cDepartment of Biomedical Sciences, Faculty of Health Sciences, University of Macau, Macau SAR, China

ABSTRACT

Voltage-clamp waveforms are imposed in the patch-clamp electrophysiology technique to provoke ion currents, the particular waveform that is used is known as the “voltage-clamp protocol”. Designing protocols to probe and quantify how gating for a particular ion channel occurs has typically been done manually and results in a suite of long protocols. It is desirable to gain the same, or even more, information in a shorter time, and also to automate the process of designing these protocols. In this paper we introduce a new optimal experimental design objective for ion channel characterisation, which involves considering a 3-dimensional phase space for the channel states combined with the voltage, using room-temperature hERG/Kv11.1 currents as an example. A range of designs are proposed, the best of which visits 82% of the discretised phase-voltage space in a 9 s protocol. This new protocol design strategy results in a simulated current visiting a wide range of channel gating states, at a wide variety of voltages, and we therefore expect these designs to be very useful in characterising ion currents, parameterising models, as well as being a challenging test of assumptions made about ion channel gating.

PLAIN LANGUAGE SUMMARY

Many cells pass signals through tissues with waves of electric charge, such as neurons firing in the brain or heart cells coordinating their beating. At the molecular level, proteins called ion channels span the outer membrane of cells, and allow specific types of ions to flow across. This movement of ions carries enough charge to alter the voltage across the membrane and pass these electrical signals. Many ion channels themselves open or close in response to changes in voltage, so there is a complicated dependence between channels being both affected by and influencing the voltage across the cell membrane in a feedback loop. Over the past 70 years mathematical models of the currents flowing through these channels have been used to understand the biological processes better, and to make predictions of what will happen in new experiments. This research paper proposes a new way to do experiments that probes a wide range of channel behaviours, across the full span of biologically-plausible membrane voltages. The results are short experiments that should allow the way ion channels open and close to be captured very quickly and distilled into mathematical models that can be used in further scientific studies. A novel aspect is that the suggested experiments should highlight any limitations or shortcomings of the mathematical models that are used in the experiment designs, providing motivation and data for use in improving models and better understanding of how the channels work in future.

ARTICLE HISTORY

Received 4 May 2024
Revised 17 June 2024
Accepted 18 June 2024

KEYWORDS

Experimental design;
mathematical model; ion
channel; electrophysiology;
hERG

1. Introduction

Ion channels are trans-membrane proteins that permit the passage of ions from one side of a biological membrane to the other. A general challenge is to characterise the kinetics of an ion channel – that is its opening and closing (or *gating*), usually in response to trans-membrane voltage changes or ligand binding. We will be concerned here with voltage-gated ion channels and macroscopic whole-cell electrophysiology. That is,

recording the ionic current that flows through a large number of ion channels in a cell’s membrane, rather than single-channel measurements. Some of the most important voltage-dependent ion channels for various electrophysiological functions are selectively expressed in excitable tissues such as neurons or muscles such as the heart.

A common method for studying ion channel gating is a whole-cell voltage-clamp experiment: through an

CONTACT Gary R. Mirams  gary.mirams@nottingham.ac.uk  Centre for Mathematical Medicine & Biology, School of Mathematical Sciences, University of Nottingham, University Park, Nottingham NG7 2RD, UK

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

ingenious system including a micropipette and electrical amplifier one can clamp the cell's trans-membrane voltage and record the currents flowing through its ion channels. A decision then is what voltage should be applied through time to gain the most information about the ion current's gating behaviour. Experimenters conceptually separate processes such as activation/deactivation and inactivation/recovery to generate protocols where one process dominates. To understand these protocols and their outputs by eye, typically they include long periods at "holding potentials" to return channels to a resting state. A suite of "conventional" protocols for studying every process is assembled, and running the complete set of protocols takes a long time, leading to them often having to be recorded in different cells and averages taken.

In recent studies we have proposed short and information-rich experimental protocols, although these have been designed manually [1–3], without any strict mathematical criteria, but with a general aim of parameter identifiability [4,5] as well as model/equation-structure selection and robustness to experimental artefacts [6].

This paper describes a new *space-filling* methodology to design short, high-information protocols algorithmically. It is based on exploring as many combinations of gating behaviour at different voltages as possible by defining the possibilities as a mathematical "space" that we can try to "fill" by optimising voltage-clamp protocols. The method should be immediately applicable to any current that models channel kinetics with Hodgkin-Huxley style gating variables, and could easily be extended to more flexible Markov model representations as well.

2. Methods

Our method is motivated by the dynamics of a simple I_{Kr} model, illustrated in Clerx et al. [2]. In that paper, a phase-plane analysis was presented as an educational tool to understand the motivation behind the voltage-clamp protocols that are commonly used to interrogate the gating of hERG channels (which carry cardiac rapid-delayed rectifying potassium current, I_{Kr}).

2.1. Ion Current Model

Our I_{Kr} or hERG model has been presented before [1–3], but is described here again as the gating processes are important for the experimental design that follows. We use the standard Ohmic approach for the current dependence on the voltage and channel open probability:

$$I_{Kr} = g_{Kr} \cdot O \cdot (V(t) - E_K). \quad (1)$$


Here "." indicates multiplication and g_{Kr} is the maximum I_{Kr} conductance. The membrane voltage, $V(t)$, changes

through time according to the voltage-clamp protocol, whilst O is the open probability (itself a function of time via the protocol). E_K is the reversal potential which we will set to -88.6 mV throughout; generally it can be calculated using the Nernst equation – a function of temperature and potassium concentrations either side of the membrane – for any particular experimental conditions [7].

In the model we use here, the open probability O can be specified as the product of two independent Hodgkin-Huxley gating processes [2,8]. We will discuss how the approach could be easily extended to Markov models with arbitrary numbers of states later. The naming of I_{Kr} gating is somewhat confusing, as deactivation and inactivation are two different processes. *Deactivation* is a relatively slow process for I_{Kr} which reduces its open probability at low voltages, whilst *inactivation* is a relatively fast process that reduces open probability at high voltages. In a Hodgkin-Huxley style current model these processes are treated as independent so there is one gating variable (a) for *activated* (the term for not-deactivated channels) which specifies the proportion channels in this state. A second gating variable (r) represents the proportion of channels *recovered* from inactivation. Being independent processes, this means that the channels can be both "activated" and "inactivated" at the same time, so the sooner we get to describing things with mathematics rather than words, the better. The open probability is given by simply

$$O = a \cdot r. \quad (2)$$

The gates themselves are considered to undergo voltage-dependent chemical reactions to transition between their activated/deactivated or inactivated/recovered states:



with $k_1 \dots k_4$ being reaction rates. Ordinary differential equations (ODEs) governing a and r can be written as

$$\frac{da}{dt} = k_1(1 - a) - k_2a, \quad \frac{dr}{dt} = -k_3r + k_4(1 - r), \quad (4)$$

based on mass-action kinetics. Following [1], each of the rates k_1 to k_4 follows a first-principles rate theory reaction for movement of a voltage sensor as a function of voltage [7]:

$$k_1 = p_1 \exp(p_2 \cdot V(t)), \quad k_3 = p_5 \exp(p_6 \cdot V(t)), \quad (5)$$

$$k_2 = p_3 \exp(-p_4 \cdot V(t)), \quad k_4 = p_7 \exp(-p_8 \cdot V(t)). \quad (6)$$

Where $p_1 \dots p_8$ are positive real scalar parameter values. Equation (4) can be recast equivalently as

$$\frac{da}{dt} = \frac{a_\infty - a}{\tau_a}, \quad \frac{dr}{dt} = \frac{r_\infty - r}{\tau_r}, \quad (7)$$

Table 1. Model parameter values, taken from Beattie et al. [1] for Cell #5 of that study.

Parameter	Value	Unit
g_{Kr}	1.524×10^{-1}	μS
p_1	2.260×10^{-4}	ms^{-1}
p_2	6.990×10^{-2}	mV^{-1}
p_3	3.448×10^{-5}	ms^{-1}
p_4	5.460×10^{-2}	mV^{-1}
p_5	8.730×10^{-2}	ms^{-1}
p_6	8.910×10^{-3}	mV^{-1}
p_7	5.150×10^{-3}	ms^{-1}
p_8	3.158×10^{-2}	mV^{-1}

where a_∞ and r_∞ denote the steady-states for each gate, and τ_a and τ_r denote the time constants. These can be defined in terms of the transition rates as

$$\tau_a = 1/(k_1 + k_2), \quad \tau_r = 1/(k_3 + k_4), \quad (8)$$

$$a_\infty = k_1 \cdot \tau_a, \quad r_\infty = k_4 \cdot \tau_r. \quad (9)$$

Note that indirectly these terms are all also dependent on time since rates themselves are voltage dependent and $V = V(t)$, as above. The model parameters for all our simulations in this paper represent room-temperature I_{Kr} from Beattie et al. [1] and are shown in Table 1.

In a typical patch-clamp experiment, multiple protocols are applied to each cell, between which they are kept at a *holding potential*. For our current of interest (I_{Kr}), the holding potential is commonly chosen to be -80 mV; close to cardiac myocytes' resting potential, where I_{Kr} channels are strongly deactivated and mostly recovered. For all the simulations that follow, we start with the model at its steady state for a holding potential of -80 mV, given by simply $a(0) = a_\infty|_{V=-80} = 3.09 \times 10^{-4}$ and $r(0) = r_\infty|_{V=-80} = 0.601$, with total open probability $O(0) = a(0) \cdot r(0) = 1.86 \times 10^{-4}$, so very little current flows. Steady state is a reasonable approximation for experimental conditions because the slowest timescale of decay towards the steady state (Equation (7)) is $\tau_a|_{V=-80} = 367$ ms and the cells will typically have been at holding potential for tens of seconds prior to the protocol being applied.

Many more complex Markov models have been proposed to capture the details of hERG gating kinetics more completely (e.g. [9–13]), but even this very simple two gate Hodgkin-Huxley style ion current model is capable of reproducing the majority of the features of I_{Kr} and predicting currents during action potential waveforms extremely well [1–3,14]¹.

¹ As an aside, a two gate Hodgkin-Huxley scheme was used in many previous I_{Kr} models, beginning with Zeng et al. [15], although this (and many others) models inactivation as an instantaneous process ($r \equiv r_\infty$). Indeed, the only other two-gate HH model that uses first-principles Eyring theory for the reaction rates is the Winslow et al. [16] action potential model which used an $O = a \cdot r_\infty$ scheme with a corresponding six kinetic parameters. The rest

2.2. Phase-voltage space

One way to study dynamical systems like this model is to examine behaviour in “phase-space”. That is, looking at steady states and trajectories of model variables on graphs where the axes are state variables instead of plotting trajectories through time. Generally these “*phase portraits*” are very informative in terms of qualitatively understanding model behaviour, evolution to steady states, etc. (see reference [17] for an accessible introduction to phase portraits for biological modelling).

In [2] we used phase portraits to understand how “conventional” voltage protocols were designed to probe gating processes such as (e.g.) activation or inactivation. Here, one axis of the phase portrait is the variable a and the other is the variable r . In Figure 1 we show how a model trajectory evolves on this diagram through a series of voltage steps. Conventional electrophysiology protocols tend to use very high or low voltages (as in the example shown in Figure 1) to push one of the gates towards zero/one whilst probing the other. The intention is often to allow one gating process to be ignored in data analysis – for example, whilst constructing a conductance/voltage curve for one gate we might assume the other gate is fully open. It has been acknowledged that this may introduce inaccuracies in later interpretation of results [18], and this was studied in the context of this model in terms of the inaccuracies introduced when fitting Equations (8) and (9) directly to current/voltage and time-constant/voltage summary curves [2].

We can visualise voltage as an additional dimension, as shown in Figure 2(C), to create what we are calling a *phase-voltage cube*. Note that in Figure 2 we can see how the conventional “time constant of activation” protocol predominantly stays at the sides of the phase-voltage cube. Even restricted to exploring two perpendicular sides (or any vertical slices) of the cube shown in Figure 2(C), then in theory – assuming their independence – we should be able to measure everything we need to know about activation and inactivation. However we may have doubts about the reaction schemes and their independence, as most published Markov models for I_{Kr} do. For instance, we may suspect intermediate conformational states may be involved, or an inactive state may be inaccessible from closed states. In this case, a motivation arises to test whether the experimental system evolves as predicted when the model is exhibiting different combinations of a and r .

of the literature hERG models appear to incorporate empirical modifications of (e.g.) a_∞ or τ_a featuring additional parameters or terms unrelated to the reactions shown in Equation (3).




Figure 1. The evolution of the model's trajectory through phase-space for a series of voltage steps. *Top*: the voltage clamp protocol, with sequential shading of different steps as they are applied through time. This is the first +40 mV test voltage of a “time-constant of activation protocol” (Protocol Pr2 from Beattie et al. [1] and Clerx et al. [2]). *Bottom*: in blue we show the evolution of the model's state variables as these steps are applied, with trajectories shown from the start of these steps to the end of the highlighted sections above. The red line indicates the function $O_{\infty}(V) = a_{\infty}(V) \cdot r_{\infty}(V)$ as voltage varies from large negative voltages (top left) to large positive voltages (bottom right). This plot is adapted from the supplementary material of Clerx et al. [2] under a CC-BY licence.




Figure 2. Construction of a 3D phase portrait. *A*: Top – the voltage clamp protocol, here sweeps with different durations for step P2 are performed and shown in different colours. Bottom – the resulting current. *B*: The phase portrait, constructed as shown in Figure 1 for the first sweep, but here for all six sweeps, colour coded according to the sweep shown in panel A. *C*: The 3D portrait showing the same trajectory through the 3D phase-voltage cube. The function $O_{\infty}(V)$ is represented here by the red-dashed line (see Figure 1 caption for how this is generated). This plot is adapted from Clerx et al. ([2], Figure 1 and supplement) under a CC-BY licence.

Similarly, if we were happy to assume that the voltage-dependence is completely understood, and perfectly follows Equations (5)–(6), then we might examine only two voltages to identify the two parameters in these equations. Indeed, traditional statistical optimal experimental design can result in very high/low voltage steps being suggested, as maximum sensitivity to parameters happens at these extremes. But if we have any doubts about the voltage-dependence of rates in Equations (5)–(6), or how their voltage-dependence interacts with any other imperfections in the model, we might wish to test the model behaviour across the whole physiological range of voltages and gating processes.

These notions of (i) testing the Hodgkin-Huxley independence assumption; and (ii) testing the voltage-dependence assumptions encoded in the rates, together motivate the idea of gathering experimental data from

as complete an exploration of the phase-voltage cube as possible. We are calling the designs that aim to do this *space-filling* protocols.

2.3. Experimental Design Algorithm

In this section we describe the algorithm that designs space-filling protocols, with the aim of exploring experimental behaviour across as many combinations of the two states and voltage as possible. For simplicity, we divide the phase-voltage space into a series of discrete “boxes”. The principal design aim is then simply to force the channel gating to visit as many of these boxes as possible. We have chosen to discretise each dimension into six, for a total of $6^3 = 216$ boxes as shown in Figure 3. The rationale here is to balance an exhaustive exploration of the space (which more divisions and boxes would




Figure 3. The discretisation of the phase-voltage space into “boxes” used in the optimisation of the experimental protocol. The objective principally maximises the number of these boxes that are visited by the model at some point during the protocol, whilst keeping the protocol short.

provide) with the limited number of steps available to encode in a protocol on automated patch platforms. A continuous version of this design approach would be possible, in terms of maximising the distance in phase space between trajectories in some sense, but would come with added computational complexity and a number of choices to make about distance measures.

There is a fixed 2.4 s voltage clamp section added to the start of all our protocol designs, defined in Table 3, and shown in Figure 4. This section contains steps which we use to estimate leak current and conductance, and being identical across all protocol designs it is also useful for quality control [1,3]. This pre-clamp needs to be considered in terms of setting the initial conditions for the design phase: we run the model from an initial steady state at -80 mV at $t = 0$ through the fixed 2.4 s section to get state variables at $t = 2.4$ s, these state variable values are then used as initial conditions for proposed designs discussed below. We also record the boxes these fixed steps visit, as shown in Figure 4.

We could do one large optimisation for all the protocol steps we wish to put in the design, but we are aiming for 51 steps, which implies 102 parameters when considering the voltage and duration of each step. This would be a very high-dimensional global optimisation problem, taking a long time to converge, and having a low chance of terminating at the global optimum. Luckily, it is not necessary to find the global minimum for a design to be “good” (in terms of achieving many of its aims and visiting a large if not-quite-maximal number of boxes). So to simplify the process substantially we have adopted a sequential design process, where we optimise three steps

at once, and repeat this 17 times to build up a complete protocol of $3 \times 17 = 51$ voltage steps. The overall scheme is illustrated in Figure 5. Whilst this sequential 3-step design process is not guaranteed to find the same number of boxes as a design that parameterised all the voltage steps at once, it results in much lower computation time and results in designs that visit much of the space.

Three steps ($s = 1, 2, 3$) are parameterised with durations (t_s in ms) and voltages (V_s in mV) encapsulated in a design parameter vector:

$$\theta = \{t_1, t_2, t_3, V_1, V_2, V_3\}.$$

The objective function that we minimise, $f(\theta)$, is evaluated by simulating the model states forward under the steps proposed by the design vector θ . Output time samples are decided by the adaptive timestep ODE solver (see below), and analysed to see which “box” the simulation trajectory is in at each time point, as we illustrated for the fixed 2.4 s initial steps in Figure 4.

N_{new} is the number of new boxes that the proposed 3-steps visit that have not been visited before. Our objective function to minimise is

$$f(\theta) = -1000 \cdot N_{\text{new}} + \underbrace{\sum_{i=1}^3 \theta_i}_{\text{duration of steps}} + \underbrace{50000 \cdot g(\theta)}_{\text{penalty term}}. \quad (10)$$

The second term of $f(\theta)$ simply keeps the protocol short. There are numerous advantages to short patch clamp experiments, the main one being that there is a limited time available before the patch clamp seal deteriorates – typically 10–30 min. It is also preferable to record some additional independent validation experiments in the same cell, and to repeat all protocols after applying a full specific pharmacological block to isolate the current of interest. Experimental conditions (such as temperature, concentrations and the magnitude of leak currents and other artefacts) can also drift over time, and this is minimised by keeping experiments short. Another benefit of very short experimental protocols is that there is time available to purposefully alter experimental conditions of interest and re-measure currents in the same cell (e.g. to examine the effect of pharmacological compounds or ionic concentration changes). The weight of 1000 on the first term in $f(\theta)$ is a normalisation factor to make the first two terms approximately the same size or equally important (since the durations are in milliseconds and we are aiming for a protocol of around 10 s duration in total, then the second term will be $\mathcal{O}(10^4)$).




Figure 4. Trajectories of the model through the fixed initial steps, and the boxes visited during this part of each protocol (blue). Note that instantaneous jumps in voltage occur on steps in the protocol, so whilst the phase-voltage plot is shown as a continuous line to assist in interpretation it does not visit boxes “en-route” whilst making the vertical jumps in voltage, as we observe here on the left- and right-hand sides of the phase-voltage cube. The line colour is set according to a colourmap though time, so that the trajectory in phase-space on the right can be related to the other plots more easily.




Figure 5. A flowchart showing the main stages of the optimal design algorithm. The initial and final fixed voltage clamps are those defined in Table 3. The rounds of optimisation refers to how many times we add 3 steps, in the designs shown here this is set to 17 rounds.

Parameter constraints could be enforced using a constrained optimisation algorithm, or suitable transforms, but for convenience we used an unconstrained optimisation algorithm and added the penalty function $g(\theta)$ to the objective function, such that $t_1, t_2, t_3 \geq$

20 ms and $V_1, V_2, V_3 \in [-120, +60]$ mV. These constraints ensure that voltages stay within physiologically-relevant ranges, and that steps are not so short that their currents are dominated by artefacts – our usual practice in fitting current traces has been to discard the data in the 5 ms after voltage steps as this region can be heavily contaminated by capacitive spikes in the cell types we have worked with to date [6]. Therefore, this penalty function takes the form:

$$\begin{aligned}
 g(\theta) = & \sum_{i=1}^3 \underbrace{(1 + 20 - \theta_i)H(20 - \theta_i)}_{\text{penalty for steps under 20 ms}} \\
 & + \sum_{i=4}^6 \underbrace{(1 + \theta_i - 60)H(\theta_i - 60)}_{\text{penalty for voltages over +60 mV}} \\
 & + \sum_{i=4}^6 \underbrace{(1 + -\theta_i - 120)H(-\theta_i - 120)}_{\text{penalty for voltages under -120 mV}}, \quad (11)
 \end{aligned}$$

where $H(\cdot)$ indicates the Heaviside step function. The use of the linear terms before the Heaviside functions gives the penalty a value of at least 1 as well as a gradient which can assist a “lost” optimiser in finding its way back into the correct area of parameter space. The factor of 5×10^4 multiplying $g(\theta)$ in Equation (10) then ensures that any penalty for infringing parameter constraints is larger than the cost associated with the duration of the steps. When minimisation is complete, θ sets the new 3 steps of the design, and the a and r variable states at the end of the selected voltage steps are used as the

initial conditions for the next 3-step optimisation in the sequence. We are using a non-gradient based optimiser (details below) so discontinuities introduced by the Heaviside functions are not a problem, continuous versions of $g(\theta)$ would be possible if needed, by introducing \tanh functions for example.

2.4. Computational Implementation

For all the computation in this study we used MatLab™ R2023b and its ODE solver `ode15s` with absolute and relative tolerances of 10^{-8} . Note that relatively fine tolerances are needed to avoid an adaptive time-stepping solver “skipping over” certain boxes when time steps spacing is relaxed in the less stiff parts of the protocol. Also note that analytic solutions to Equation (7) are available for fixed voltage steps, of the form $a(t') = (a(0) - a_\infty) \exp\{-t'/\tau_a\} + a_\infty$ where t' is the time since the start of each step and $a(0)$ is the initial condition for that step, which could potentially speed up computations substantially. However we used a numerical ODE solver as the fixed sections include ramps where this analytic form of solution cannot be used, and also so that we could investigate the use of ramps within the optimised part of the design (not shown, but see Discussion section comments). A hybrid analytic/numeric solution scheme would be reasonably easy to construct – to use analytic solutions across fixed voltage steps and a numerical ODE solver elsewhere – but with the whole design process running reasonably quickly on a desktop PC this was not necessary at this time. Code to reproduce this study is openly available, please see the “Data Availability” section at the end of the paper.

To get a good initial guess for θ we uniformly sample each θ_i from range $t \in [20, 1000]$ ms or $V \in [-120, +60]$ mV as appropriate (so that the penalty term $g(\theta) = 0$), to get 1000 initial guesses for the design parameter vector θ . We then evaluate $f(\theta)$ for each of the 1000 initial guesses, and select the best of these design vectors to be the initial guess for optimisation. The computation for this step is relatively cheap with 1000 evaluations of $f(\theta)$ taking less than a second on a desktop computer.

Having a reasonable initial guess, we then use the “Covariance Matrix Adaptation – Evolutionary Strategy” (CMA-ES) global optimisation algorithm [19]. CMA-ES samples from a multi-variate Normal distribution to get a cloud of particles/parameter sets, evaluates the score at each one, then evolves the Normal distribution’s mean and covariance matrix – with the aim of moving towards the optimum before shrinking the size of the cloud to “home in” on the optimum. An important feature is that the CMA-ES covariance can get wider if points far from

the mean are getting good scores, assisting in searching more widely if needed. We used a parameter-refinement stopping criterion of `xtol` = 2 (ms or mV); that is, when the algorithm is finding that improvements in score are only happening for changes in the parameters of less than 2 ms (or 2 mV) it will terminate. This step size was picked by considering that voltages vary over 180 mV and step durations are typically on the order of 100 or 1000 ms, and refinement below this tolerance is unlikely to materially change the exploration of the phase-voltage cube. In fact, the first step of evaluating $f(\theta)$ is to round-up all the parameters to the nearest millisecond or millivolt with a “`ceil`” call, just to produce designs that are simple to communicate and possible to implement on any experimental patch-clamp hardware. Because the optimiser does not know about this step, it would cause problems (a flat objective) if the optimiser was trying to refine parameter estimates below one millisecond or millivolt, as they are all rounded to the same value, but the choice of `xtol` circumvents this.

In practice, CMA-ES does not find an optimum that is better than the initial guess on every run, as it is a stochastic algorithm and does not directly evaluate at the initial guess but instead uses it as the mean for a cloud of particles spread across the domain. Here the CMA-ES population was set to 50, and the hyperparameter for the initial population variance was set to 100 ms for step durations ($\theta_{1,2,3}$) and 20 mV for step voltages ($\theta_{4,5,6}$). All other CMA-ES settings were left at default values. We then run CMA-ES repeatedly until it does improve on the score of the best of the 1000 random guesses, in the results presented below this usually happens on the first run, rarely takes more than 2 runs, and always happened before 10 runs.

The designs we present used 17 iterations of this 3 step procedure. Following the optimisation procedure we append another fixed section of protocol that we have termed a “*reversal ramp sequence*”. This is comprised of another 6 clamp sections: a step to cause large activation of I_{Kr} followed by removal of inactivation then a ramp to detect at what voltage current reverses (see [3], Figure 10), details are in Table 3. The motivation of this reversal ramp is to observe the apparent reversal potential, with a suggestion in other recent work that the degree of any deviation from the Nernst potential might indicate the size and influence of experimental artefacts [6]. The complete protocol is then assembled from the 6 initial steps, plus $17 \times 3 = 51$ optimised steps, plus 6 reversal ramp sequence steps: a complete protocol of 63 clamp instructions which is plotted and analysed below in terms of box coverage. The automated patch machine we initially produced these designs for had a control software limitation of 64 sections in a single protocol definition. If




Figure 6. The optimal design that was found. Left: the voltage protocol, gating variables and current through time for the best design. Right: the phase-voltage space with the trajectory for this design plotted. This protocol visits 178/216 boxes (82.4%) and therefore to make this figure clearer we highlight just those boxes that are *not* visited in red. Note that the plot is rotated relative to Figure 4, again for clarity, and the un-visited boxes are clustered in the $a \approx 1, r \approx 1$ corner. This finding makes intuitive sense as it is difficult for both gating variables to be close to one at the same time for any voltage given that $a_\infty \rightarrow 1$ at high voltages and $r_\infty \rightarrow 1$ at low voltages.

future machines have a larger limit, then extra rounds of 3-step designs could be added to attempt to hit some of the remaining unvisited boxes, or the number of boxes (currently 6^3) could be increased to explore the accessible region of the phase-voltage cube to a greater degree of refinement.

3. Results

Because our optimiser uses random sampling internally for placing its cloud of particles, every run produces a different design. The optimal design we found after 100 runs of the protocol design procedure is shown in Figure 6; this protocol visits 178/216 boxes (82.4%). For context, the “sinusoidal protocol” introduced in Beattie et al. [1] visits a surprisingly-low 54 boxes (25%), whilst the “staircase protocol” introduced in Lei et al. [3] visits 60 boxes (28%).

The five top-scoring protocols out of the 100 runs are shown in Figure 7 with the details of all steps given in Table 2. We see some common properties emerge – sections of long steps, ensuring we observe certain slow activation gate states at certain voltages, and also interspersed with shorter regions of faster steps up and down which probe the inactivation processes at various voltages. As these can be difficult to see we have shown zoomed in regions of the protocol and currents it provokes on the right-hand side of Figure 7. As we ensured with the penalty term $g(\theta)$ and rounding procedure described above, none of these steps are shorter than 20 ms. Steps of 20 ms are chosen repeatedly though: this is not a surprise, given the need for short steps to provoke rapid transients to explore the inactivation/recovery

processes at various levels. Similarly, the extremes of the voltage range of -120 to $+60$ mV are commonly used to provoke fast behaviour.

Figure 8 summarises the properties of the 100 designs in terms of how many boxes are visited and how long the resulting protocols are. Protocols are all between 8 and 12 s long, and visit 73–82% of the available discretised boxes of the phase-voltage cube.

3.1. Identifiability

To test whether the protocols that result from this design process provide good parameter identifiability for a model, a worthwhile exercise is always to simulate data with some parameter set, “forget” the ion channel model parameters (p_1, \dots, p_8, g_{Kr}), and then attempt to re-infer them from just the simulated current trace with noise added. We have not shown the full results of this exercise here, but we have tested tens of the resulting designs (following the “Method 4” parameterisation procedure in Clerx et al. [2], which involves simply optimising model parameters to minimise square error between the model prediction and the full simulated current trace). In the Shuttleworth et al. [20] simulation study we used one of these phase-voltage cube filling designs (called “design d_3 ” in that paper). In that paper’s Table 3 ($\lambda = 1$ setting) results are shown of a practical identifiability assessment for the Beattie et al. [1] model that was used to generate the protocol design. A realistic level of synthetic i.i.d. noise was added to the model output to generate synthetic data, and variation in the parameter fits across repeated optimisations with 10 different instances of that noise is shown in the Table, very small standard deviations on

Table 2. Details of the 5 most optimal protocols, as shown in Figure 7.

Clamp #	Protocol 1 (best)		Protocol 2		Protocol 3		Protocol 4		Protocol 5	
	t (ms)	V (mV)	t (ms)	V (mV)	t (ms)	V (mV)	t (ms)	V (mV)	t (ms)	V (mV)
1	672	29	648	29	643	29	645	29	645	29
2	84	-101	53	-109	145	-103	53	-109	53	-109
3	52	-13	27	-7	22	-1	23	-2	23	-2
4	20	-112	24	-98	22	-95	37	-116	20	-106
5	101	60	105	59	109	60	151	-1	99	60
6	195	-90	196	-90	195	-90	27	-90	195	-90
7	170	29	170	29	21	59	108	60	169	29
8	281	-87	280	-87	281	-87	194	-90	280	-87
9	20	-1	21	-3	21	-6	21	-42	22	-3
10	20	-119	20	-66	20	-118	169	29	20	-118
11	22	-36	20	-116	20	-33	20	-82	20	-31
12	20	-77	20	-33	20	-80	20	54	20	-83
13	147	58	132	60	204	60	47	-80	184	57
14	306	-59	69	-87	77	-90	697	29	394	-57
15	37	-6	21	-2	53	-16	333	-57	53	-7
16	63	-90	317	-60	314	-60	21	60	41	-43
17	144	60	20	-120	20	-119	62	-120	58	-90
18	26	-90	21	-33	24	-41	29	-40	26	52
19	20	24	513	29	43	-90	239	58	29	60
20	40	-119	34	-106	144	24	78	-90	33	-120
21	24	-40	21	58	70	-52	77	-18	32	-35
22	31	-89	53	-90	177	29	514	-51	20	-77
23	288	25	80	60	42	-90	45	-120	30	29
24	167	-114	21	-60	21	-4	20	-61	20	-91
25	73	60	20	54	82	60	802	48	174	60
26	24	-60	20	-119	21	-60	20	-95	28	-90
27	20	27	20	26	20	-117	188	-16	21	13
28	24	59	21	-19	55	57	59	17	22	-120
29	20	-97	21	-116	24	-94	37	-98	23	-1
30	21	25	23	-36	412	29	41	19	20	-59
31	20	-59	34	-120	24	47	30	-120	322	58
32	20	-120	24	-2	34	-57	30	4	26	-93
33	20	-23	20	-115	35	-15	32	-45	138	-21
34	20	8	179	60	33	-92	255	-120	22	-57
35	21	-60	20	-95	232	38	48	30	21	-120
36	20	50	20	28	103	-26	206	-116	20	15
37	45	-120	21	60	35	-120	20	-66	47	60
38	31	-9	21	-78	20	12	20	-1	20	-95
39	21	-41	20	-12	20	-114	20	-75	20	-19
40	235	60	42	-118	544	58	178	60	106	60
41	22	-94	154	-18	33	-92	20	-95	20	-95
42	41	-15	63	-62	92	15	20	59	20	59
43	20	41	50	12	22	-105	20	-120	20	-120
44	20	-99	41	-54	20	-60	20	-25	20	-33
45	20	-26	34	37	20	22	20	-28	20	-42
46	88	60	152	60	20	-118	163	38	48	-120
47	20	-95	20	-95	20	-51	20	-118	23	-2
48	20	60	20	60	20	-25	685	-11	20	-51
49	40	-120	20	-101	20	60	20	-116	22	-16
50	20	-38	20	-26	20	-96	47	60	44	-90
51	20	-58	20	-33	20	-29	20	-39	25	33

Note: These steps need to have the two “bookend” sections added (see Table 3) which are identical for all designs.

Table 3. Details of the start and end clamp sections for all designs.

Clamp #	Initial steps: for leak and conductance			End: reversal ramp sequence		
	Step/Ramp	t (ms)	V (mV)	Step/Ramp	t (ms)	V (mV)
1	Step	250	-80	Step	1000	-80
2	Step	50	-120	Step	500	40
3	Ramp	400	-120 to -80	Step	10	-70
4	Step	200	-80	Ramp	100	-70 to -110
5	Step	1000	40	Step	390	-120
6	Step	500	-120	Step	500	-80

Note: “t” indicates the duration of the clamp section, and “V” the relevant voltage(s) for this clamp. Where “Ramp” is specified it is a linear ramp over time between the voltages shown, as opposed to a constant voltage clamp for a “Step”.




Figure 7. The five best scoring voltage-clamp designs out of 100 runs of the phase-voltage cube optimisation procedure, and the simulated currents they provoke. Full details of the clamps are given in Table 2. Each row is one protocol, ranked from the best at the top (also seen in Figure 6) to 5th best at the bottom which still visits 175/216 boxes (81%). The voltage protocol is shown in blue (left hand axis), and the resulting simulated I_{K_r} is shown in red (right hand axis). On the left we show the full protocol, and on the right a zoom in on the grey highlighted section from the left plot, sharing the same y axes. The horizontal grey line indicates the reversal potential E_K (left axis) and zero current (right axis).




Figure 8. Summary statistics for 100 runs of the design process, with histograms on the axes showing the marginal distribution of the cloud of points.

the true parameters are returned, and examining the data behind those tables (see “Data Availability”), the maximum absolute percentage error for any parameter in the best fits to any of these datasets was less than 0.4%.

3.2. Predictions for a different model

For our second model we use a linear structured Markov model with three closed states, open and then inactive (C-C-C-O-I) as proposed by Wang et al. [9]. In the C-C-C-O-I model, there is no longer independence between activation and inactivation processes (the inactivated state only being accessible from the open state). This model was re-calibrated in Beattie et al. [1] to the same data as the Hodgkin-Huxley model that we used for the optimal designs here. Please see the supplement of Beattie et al. [1] for a table of parameters.

Shuttleworth et al. [20] tested whether space-filling designs still provide good identifiability by repeating the




Figure 9. Predictions of the C-C-C-O-I model by Wang et al. [9] re-calibrated to data gathered under the sinusoidal protocol (done by Beattie et al. [1] using their “Cell #5” data, which is the same data used to calibrate the Hodgkin-Huxley model used for designs here). **Top row:** voltage clamp protocol, **Bottom row:** simulated current. **Left:** sinusoidal protocol – both models display similar currents. **Middle:** under the best design – the two models that closely agreed for almost all the sinusoidal protocol show greater divergence under the space-filling design. **Right:** a zoomed in view of part of the space-filling design predictions (the shaded region from the middle plots).

process discussed in the section above with this C-C-C-O-I model providing the synthetic data. They found that there was very good parameter identifiability for this model as well, with a mean absolute percentage error in parameters of 0.4% and a maximum absolute percentage error for any of the 15 model parameters across all 10 synthetic datasets of 2.9%.

Taking the optimal design discussed above (shown in Figure 6), generated based upon the Hodgkin-Huxley model phase-voltage cube, we now examine currents that are predicted under this protocol with a different model. In Figure 9 we show how the newly designed protocol separates the model predictions well in various parts of the protocol, the violation of Hodgkin-Huxley model assumptions or its simplifications of the C-C-C-O-I model appear to be highlighted well, even though both respond to the sinusoidal protocol used in their calibration very similarly.

4. Discussion

In summary, these new protocol designs visit a wide range of channel gating states, at a wide variety of voltages, and we therefore expect them to be very useful in parameterising models, as well as being a challenging test of the assumptions we make about ion channel gating when constructing ion current models.

The question “how much box coverage is desirable?” is difficult to answer, as it depends on how much reality deviates from the assumptions behind the gating rates and the gate independence. As mentioned in the methods, if the assumptions hold across the whole phase-voltage cube, then exploring only two perpendicular sides or slices would be sufficient to characterise activation and inactivation independently, and space-filling designs would simply provide extra confirmatory information as we explore the cube. If we doubt any of the assumptions however, then exploring more of the cube may highlight more model inadequacy and be very useful data to collect, guiding future refinements of a model.

We are in the process of testing these designs with real experimental hERG1a data at present, and intend to publish a comparison of the model parameterisations that result.

4.1. Limitations

The designs here were *local*, that is based upon a single parameter set for a single candidate model (described by Equations (1)–(6)). So there is a need to understand something about the channel of interest and to have a reasonable candidate model before running this design process. That said, we believe such designs will be extremely useful as part of a model refinement pipeline, particularly

when moving from the simplest Hodgkin-Huxley gating variable representations to Markov models with more parameters and flexibility.

Some of the boxes are visited for only very short amounts of time. As an example, downsampling the output from using all the ODE solver's timesteps (as we do in the above results), to regular 1 ms steps results in a drop in box coverage for the 100 designs from 76–82% to 61–70%, implying that around 15% of the boxes we visit are occupied for less than 1 ms as we pass through them on very fast trajectories. This is not a surprise given that the fastest time constant in the model for this voltage range is the inactivation's $\tau_r = 3.9$ ms at -120 mV. The implication is that on a step to -120 mV we will get $1 - 1/e = 63.2\%$ of the way from the initial r to r_∞ in just 3.9 ms, so it is easy to see why we might spend less than 1 ms in many of the boxes en-route if we start more than a couple of boxes away. Some additional criteria within the objective to spend a good amount of time in each box could be desirable. But extra time in certain boxes might be difficult to achieve without just revisiting the box many times: inactivation is just a fast process at certain voltages – a biological feature of I_{Kr} rather than a quirk of this model or approach.

The 3-step-at-a-time approach is a limitation, in that an all-at-once optimisation should be able to hit more boxes and potentially reach boxes that cannot be hit in only 3 steps. However we re-ran our analysis with a 6-step design, allowing $6 \times 9 = 54$ steps in total (slightly more than the 51 shown here). We achieved exactly the same optimum of 178 boxes in 100 runs, suggesting the 3-step design is flexible enough to find a similar optimum to those we would get with more steps fitted at once. We also tried a sequential design comprised of step-ramp-step-ramp-step (with ramps continuous with the step voltages either side), as we imagined that ramps, that visit all the voltages en-route between two steps, would visit a lot of extra boxes in the voltage dimension. However, this time only ten sequential optimisations could occur with the limited number of clamp commands (5 clamps \times 10 = 50 steps) which led to only 68% of boxes being visited. Related to the paragraph above, primarily this is thought to be due the removal of a lot of the fast vertical transitions in voltage which mean that it is harder to get the inactivation state away from r_∞ when using ramps rather than steps. Code and results for both of these exercises are provided in a subfolder of the code (see Data Availability).

4.2. Future Work

Various adaptations would be easy to include as the proposed framework is very flexible. For a reasonable

amount of time during the optimised protocols there are relatively small currents (Figure 7), although overall they compare favourably with past protocols in terms of current magnitude. One idea to increase the proportion of the protocol with large currents would be a simple penalty term for small currents (based on $|a \cdot r \cdot (V - E_K)|$), which could encourage larger currents and a better signal-to-noise ratio. There would be caveats here, as larger currents in patch-clamp experiments imply larger artefacts related to voltage drop over series resistance; bigger signal with lower signal-to-noise is not always better in patch clamp experiments. Consideration of artefacts during the fitting process may allow us to overcome this problem [6]. Such an alteration would also be at the expense of somewhat arbitrarily enhancing the weighting given to exploring certain parts of the phase-voltage cube: it may be very useful information to observe that there really are small currents when you expect that there should be small currents.

Extension to a higher-dimensional phase-voltage hypercube for Markov models with more than three states, or additional Hodgkin-Huxley gates such as a 3-gate sodium m^3hj model, is very straightforward – at least mathematically and computationally, if not in terms of visualisation. For Markov models, this extension may require more than 3 steps to be optimised at a time to get into hidden deeper states, and the number of boxes in each dimension (here 6) might need to be reduced to give a good exploration across the whole space with the limited number of steps available in a practical setting. There may also be substantially more “inaccessible space” in a larger Markov model – for instance, it is difficult to imagine in a model with a long chain of closed states that a large probability can be produced in one of the middle states with the rest being close to zero.

5. Conclusions

These new experimental designs for voltage-clamp experiments are short and provide plentiful information on model kinetic parameters, as well as providing data that should probe the validity of key assumptions behind the ion channel modelling approaches such as independence of gating processes and transition rate dependencies on voltage.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the Wellcome Trust [grant no. 212203/Z/18/Z], the EPSRC [grant no. EP/R014604/1], the

University of Macau [reference no. SRG2024-00014-FHS], and the Science and Technology Development Fund, Macao SAR (FDCT) [reference no. 0155/2023/RIA3 and 0048/2022/A]. GRM acknowledges support from the Wellcome Trust via a Wellcome Trust Senior Research Fellowship. CLL acknowledges support from the FDCT and the University of Macau. The authors would like to thank the Isaac Newton Institute for Mathematical Sciences for support and hospitality during the programme “The Fickle Heart” when some work on this paper was undertaken. This research was funded in whole, or in part, by the Wellcome Trust [212203/Z/18/Z]. For the purpose of open access, the authors have applied a CC-BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

Data availability statement

Code associated with this paper is openly available under a BSD licence at <https://github.com/CardiacModelling/space-filling-designs/>, a permanently archived version is available on Zenodo [21].

ORCID

Gary R. Mirams  <http://orcid.org/0000-0002-4569-4312>
 Michael Clerx  <http://orcid.org/0000-0003-4062-3061>
 Dominic G. Whittaker  <http://orcid.org/0000-0002-2757-5491>
 Chon Lok Lei  <http://orcid.org/0000-0003-0904-554X>

References

- [1] Beattie KA, Hill AP, Bardenet R, et al. Sinusoidal voltage protocols for rapid characterisation of ion channel kinetics. *J Physiol (Lond)*. 2018;596(10):1813–1828. ISSN 1469-7793. doi: [10.1113/JP275733](https://doi.org/10.1113/JP275733)
- [2] Clerx M, Beattie KA, Gavaghan DJ, et al. Four ways to fit an ion channel model. *Biophys J*. 2019;117(12):2420–2437. doi: [10.1016/j.bpj.2019.08.001](https://doi.org/10.1016/j.bpj.2019.08.001)
- [3] Lei CL, Clerx M, Gavaghan DJ, et al. Rapid characterization of hERG channel kinetics I: Using an automated high-throughput system. *Biophys J*. 2019;117(12):2438–2454. ISSN 0006-3495. doi: [10.1016/j.bpj.2019.07.029](https://doi.org/10.1016/j.bpj.2019.07.029)
- [4] Fink M, Noble D. Markov models for ion channels: versatility versus identifiability and speed. *Philos Trans R Soc A Math Phys Eng Sci*. 2009;367(1896):2161–2179. doi: [10.1098/rsta.2008.0301](https://doi.org/10.1098/rsta.2008.0301)
- [5] Whittaker DG, Clerx M, Lei CL, et al. Calibration of ionic and cellular cardiac electrophysiology models. *Wiley Interdiscip Rev Syst Biol Med*. 2020;12(4):e1482. doi: [10.1002/wsbm.v12.4](https://doi.org/10.1002/wsbm.v12.4)
- [6] Lei CL, Clerx M, Whittaker DG, et al. Accounting for variability in ion current recordings using a mathematical model of artefacts in voltage-clamp experiments. *Philos Trans R Soc A*. 2020;378(2173):20190348. doi: [10.1098/rsta.2019.0348](https://doi.org/10.1098/rsta.2019.0348)
- [7] Hille B. *Ion Channels of Excitable Membranes*. 3rd ed. Sunderland (MA): Sinauer Associates Inc.; 2001.
- [8] Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol*. 1952;117(4):500–544. doi: [10.1113/jphysiol.1952.sp004764](https://doi.org/10.1113/jphysiol.1952.sp004764)
- [9] Wang S, Liu S, Morales MJ, et al. A quantitative analysis of the activation and inactivation kinetics of HERG expressed in *Xenopus* oocytes. *J Physiol (Lond)*. 1997;502(1):45–60. ISSN 0022-3751 1469-7793. doi: [10.1111/j.1469-7793.1997.045bl.x](https://doi.org/10.1111/j.1469-7793.1997.045bl.x)
- [10] Clancy CE, Rudy Y. Cellular consequences of HERG mutations in the long QT syndrome: precursors to sudden cardiac death. *Cardiovasc Res*. 2001;50(2):301–313. doi: [10.1016/S0008-6363\(00\)00293-5](https://doi.org/10.1016/S0008-6363(00)00293-5)
- [11] Mazhari R, Greenstein JL, Winslow RL, et al. Molecular interactions between two long-QT syndrome gene products, HERG and KCNE2, rationalized by in vitro and in silico analysis. *Circ Res*. 2001;89(1):33–38. doi: [10.1161/hh1301.093633](https://doi.org/10.1161/hh1301.093633)
- [12] Fink M, Noble D, Virag L, et al. Contributions of herg k⁺ current to repolarization of the human ventricular action potential. *Prog Biophys Mol Biol*. 2008;96(1–3):357–376. doi: [10.1016/j.pbiomolbio.2007.07.011](https://doi.org/10.1016/j.pbiomolbio.2007.07.011)
- [13] Kemp JM, Whittaker DG, Venkateshappa R, et al. Electrophysiological characterization of the hERG R56Q LQTS variant and targeted rescue by the activator RPR260243. *J Gen Physiol*. 2021;153(10):e202112923. doi: [10.1085/jgp.202112923](https://doi.org/10.1085/jgp.202112923)
- [14] Lei CL, Clerx M, Beattie KA, et al. Rapid characterization of hERG channel kinetics II: temperature dependence. *Biophys J*. 2019;117(12):2455–2470. doi: [10.1016/j.bpj.2019.07.030](https://doi.org/10.1016/j.bpj.2019.07.030)
- [15] Zeng J, Laurita KR, Rosenbaum DS, et al. Two components of the delayed rectifier K⁺ current in ventricular myocytes of the guinea pig type: theoretical formulation and their role in repolarization. *Circ Res*. 1995;77(1):140–152. doi: [10.1161/01.RES.77.1.140](https://doi.org/10.1161/01.RES.77.1.140)
- [16] Winslow RL, Rice J, Jafri S, et al. Mechanisms of altered excitation-contraction coupling in canine tachycardia-induced heart failure, II: model studies. *Circ Res*. 1999;84(5):571–586. doi: [10.1161/01.RES.84.5.571](https://doi.org/10.1161/01.RES.84.5.571)
- [17] Garfinkel A, Shevtsov J, Guo Y. *Modeling life: the mathematics of biological systems*. Cham: Springer; 2017.
- [18] Lee J, Smaill B, Smith N. Hodgkin–Huxley type ion channel characterization: an improved method of voltage clamp experiment parameter estimation. *J Theor Biol*. 2006;242(1):123–134. doi: [10.1016/j.jtbi.2006.02.006](https://doi.org/10.1016/j.jtbi.2006.02.006)
- [19] Hansen N, Müller SD, Koumoutsakos P. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evol Comput*. 2003;11(1):1–18. doi: [10.1162/106365603321828970](https://doi.org/10.1162/106365603321828970)
- [20] Shuttleworth JG, Lei CL, Whittaker DG, et al. Empirical quantification of predictive uncertainty due to model discrepancy by training with an ensemble of experimental designs: an application to ion channel kinetics. *Bull Math Biol*. 2024;86(1):2. doi: [10.1007/s11538-023-01224-6](https://doi.org/10.1007/s11538-023-01224-6)
- [21] Mirams GR. *Cardiacmodelling/space-filling-designs*: Version 1.0, 2024. Available from: <https://doi.org/10.5281/zenodo.11103980>