Efficient Generation of Populations of Cardiac Models

Darby I Cairns[1], Elizabeth M Cherry[1]

[1] Georgia Institute of Technology, Atlanta, GA, USA

Abstract

To model variability of cardiac action potentials (APs), a population of models (PoM) consisting of different sets of a model's parameter values can be created and calibrated to match observed variability in properties such as AP duration (APD). However, producing appropriate parameter sets for the PoM can be difficult and time-consuming. We adapted a particle swarm optimization (PSO) optimization technique to generate a population of models efficiently. Our population PSO (PPSO) algorithm discourages convergence to a local minimum, and instead guides the search to explore low-error areas of parameter space, yielding many parameter sets that can reproduce the variability of biomarkers seen in real tissue data. Using canine ventricular microelectrode recordings and a synthetic dataset, we extracted sets of APD- and voltage-based biomarkers, allowing ±10% and ±30% variations of the base biomarker values to represent variability. We created 5000- and 2500-member PoMs fitting the parameters of the Fenton-Karma (FK) and ten Tusscher-Noble-Noble-Panfilov (TNNP) models to the biomarker ranges using PPSO. Compared to a random approach, our novel PPSO method produced PoMs matching biomarkers with similar coverage of parameter space for both the FK and TNNP cases, but with greater computational efficiency, accepting up to 10 times more candidate parameter sets.

1. Introduction

Different patients exhibit considerable variability in cardiac electrophysiological properties, which can result in differing responses to cardiac interventions. For this reason, it has become important for *in silico* testbeds to consider variability within the models they include to capture the wide range of outcomes from potential treatments and to improve predictions [1]. One popular approach involves creating a population of models (PoM) by generating parameter sets that fall within ranges in biomarkers measured from data, such as APD values at different thresholds, resting membrane potential, and peak voltage, to reproduce the observed variability in real-world data [2, 3].

Although the PoM approach is beneficial, it poses a computational challenge to produce a population of parameter sets whose corresponding biomarker values are within the desired ranges. Current approaches to generating these models include starting from a baseline parameter set and randomly varying those values [4] or simply randomly generating parameter sets in their entirety [2], and then selecting those sets which meet the specified criteria. Often, a large number of candidate solutions must be considered to produce a relatively small number of acceptable candidates, such as on the order of 2% [2].

In this paper, we present a novel algorithm to generate PoMs more efficiently. Our algorithm is motivated by particle swarm optimization (PSO), a derivative-free optimization method that seeks a balance between improving solutions locally and exploring the search space broadly. PSO is an especially appealing base approach because of its computational efficiency, including the algorithm's parallelism. To generate a PoM instead of a single optimum, our population PSO (PPSO) method ensures that high-quality solutions are spread throughout the search space. This process allows the creation of a diverse population similar to one obtained through random search to be generated, but at a much lower computational cost.

2. Methods

2.1. Models

In this paper, we construct populations of models using two base models: the Fenton-Karma (FK) model [5] and the ten Tusscher-Noble-Noble-Panfilov (TNNP) model [6, 7]. The FK model is a phenomenological model consisting of three variables representing voltage and two gating variables along with fast inward, slow inward, and slow outward currents. The TNNP model is a more detailed model that represents a human ventricular cell using 19 variables and 12 transmembrane currents. In the present study, all of the FK model's 13 parameters are allowed to vary when generating a population of models, whereas for the TNNP model only the 12 conductance or other current scaling parameters are included.

The FK model was solved using forward Euler. For the TNNP model, forward Euler was used for the voltage

along with the Rush-Larsen method [8] for the gating variables and an analytical technique for the calcium buffering [9]. The timestep was $0.01\,\mathrm{ms}$ for both models.

2.2. Population PSO Algorithm

Our algorithm builds on particle swarm optimization (PSO), a derivative-free optimization algorithm that functions by creating a randomly initialized pool of candidate solutions (particles) that explore the search space [10]. Each particle has a velocity vector, which is used to update its position. The velocity of each particle is updated based on known low-error locations: the lowest-error position observed by any particle during the algorithm (the global best), and the lowest-error position the particle being updated has ever itself occupied (local best).

The PPSO algorithm begins by generating a set number of candidate parameter sets and evaluating their fitness with respect to the biomarkers, as described below. Next, the K lowest-error particles, where K is a user-selected number of "groups," are used to seed each group, and particles are permanently assigned to the nearest group based on the L^2 norm of the difference vector of normalized parameter values. PSO is then run within each group. Following each fitness evaluation, any particles that fall within the specified tolerances for each biomarker are accepted into the population. The algorithm is run until a maximum number of iterations is reached or until the desired population size has been reached. For the PoM sizes and biomarkers discussed in this paper, the algorithm normally completes within a small number of iterations, usually five or fewer.

2.3. Data, Biomarkers, and Fitness

For the FK model case, we used a dataset derived from a microelectrode recording of canine endocardial tissue APs. The voltage values were normalized to match the output of the FK model, and then the biomarkers used by PPSO were calculated from this reference data. For the TNNP model scenario, the dataset was generated from the model itself using the previously published parameter set for epicardial data. In both cases, a cycle length of $500\,\mathrm{ms}$ was used. For these preliminary results, we used four biomarkers: peak voltage (V_{peak}), maximum upstroke velocity (dV/dt_{max}), APD_{80}, and APD_{20}. For the TNNP model case, an additional biomarker, peak intracellular calcium concentration, was used. Within the PPSO algorithm, each particle's fitness was computed as the mean relative error of the biomarker values calculated using the associated parameter values with respect to these reference biomarker values. From these baseline values, we considered PoMs with high and low variability as allowing ±30% and ±10% deviation from these values, respectively, to be accepted into the population.

2.4. Implementation Details

Our PPSO algorithm utilized 4096 particles and $K = 40$ groups for all scenarios and sought to build a population of size 5000 and 2500 for the FK and TNNP cases, respectively. All biomarkers were measured from two successive APs at a single cycle length. To mitigate transient initial conditions of the models, the FK and TNNP models were paced each iteration for four and 13 APs, respectively, before using the following two APs for error calculation.

Parameter bounds were selected to allow a sizable search space, assuming limited prior knowledge. The random approach used for comparison with PPSO randomly selected parameter values within the same parameter bounds and tested whether the resulting biomarkers were within the specified tolerances; if so, the parameter set was accepted into the population.

Efficiency calculations represent averages over ten runs except for the low-variability random searches, which were both taken from single representative runs due to their high computational cost.

3. Results

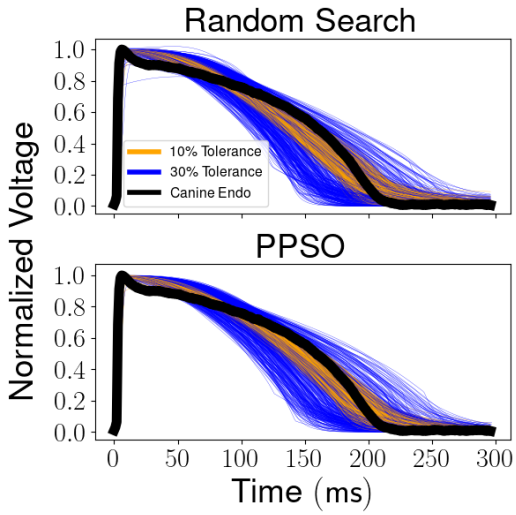

Figure 1. 500 randomly selected action potentials from 5000-member FK model PoMs generated by random search (top) and PPSO (bottom) with tolerances of 10% (orange) and 30% (blue). Black: reference canine dataset from which the biomarkers were generated.

Figure 1 shows samples from the PoMs generated using a random-search approach and the PPSO method with the FK model. Both methods produced similar distributions of action potentials for both the high- and low-variability

cases (blue and orange, respectively). Moreover, as shown in Figure 2, the parameter values of the random- and PPSO-generated PoMs have similar characteristics with comparable coverage of parameter space, so that although PPSO is derived from an optimization method, it did not produce a population that converged to a small region of the search space.

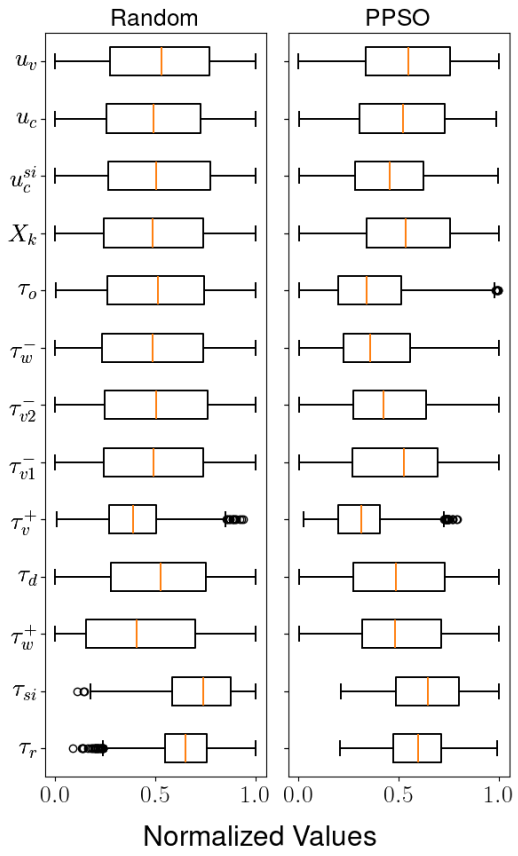

Figure 2. Spread of normalized parameter values for all members of 5000-member FK model PoMs generated via random search and PPSO for the low-variability (10% tolerance) scenario.

In terms of efficiency, for the FK model with high variability, a random search required 27,198 simulations to create a PoM with 5000 members, while PPSO required 9,181 simulations, corresponding to acceptance rates of 18% and 54%, respectively. Thus, PPSO provided a computational savings of a factor of three. For the more constrained low-variability scenario, random search required 213,001 simulations to build a PoM of the same size (2.3% acceptance), whereas PPSO required 23,292 simulations, an almost tenfold increase in efficiency (21% acceptance).

For the TNNP PoM scenario, Figure 3 shows example action potentials from random search and the PPSO method. Similar variability in action potential shapes and

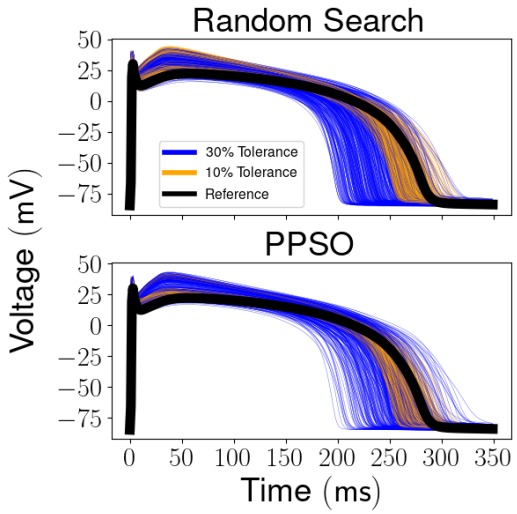

Figure 3. 500 randomly selected action potentials from 2500-member TNNP model PoMs generated by random search (top) and PPSO (bottom) with tolerances of 10% (orange) and 30% (blue). Black: reference synthetic dataset from which the biomarkers were generated.

durations can be seen for populations generated using the same biomarker variability levels (high variability, 30%, in blue and low variability, 10%, in orange). As with the FK model, the values of the parameters within the populations obtained using random search and PPSO were also comparable, as shown in Figure 4, indicating both methods achieved similar coverage of parameter space.

As for efficiency in the TNNP scenario, for the high-variability case, a random search required 16,921 simulations to create a PoM with 2500 members (14.7% acceptance), compared with 12,949 for PPSO (19.3% acceptance). In the low-variability case, random search required 247,812 simulations (1.0% acceptance), whereas PPSO required only 33,248 (7.5% acceptance), an improvement of a factor of 7.5 over the random approach.

4. Discussion

Our results demonstrate that our new PPSO algorithm is capable of generating large and diverse populations of cardiac models with variability comparable to random search with greater efficiency, thus requiring far fewer simulations, particularly under tighter biomarker constraints. Especially for mechanistic models of interest whose simulation time can be quite long, this rapid generation of PoMs could facilitate large-scale simulation studies of virtual patient cohorts with properties specified by the biomarker ranges.

While our preliminary results show that PPSO is suitable for generating FK and TNNP PoMs, and the algorithm

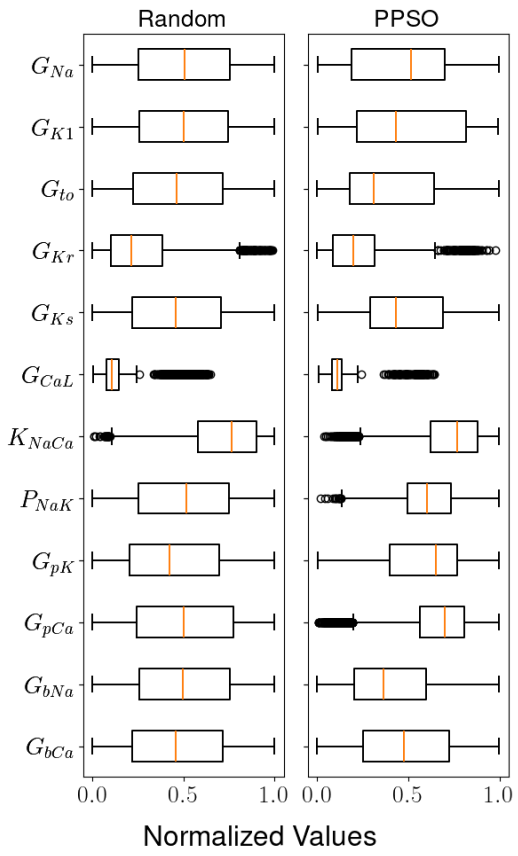

Figure 4. Spread of normalized parameter values for all members of 2500-member TNNP model PoMs generated via random search and PPSO for the low-variability (10% tolerance) scenario.

is not model-specific, further study is required to verify PPSO's suitability for other models, as well as its performance with more biomarker constraints.

PPSO also relies on a good selection of initial groups, and more in-depth initial optimization strategies for that selection may be necessary for broader applications. For example, neither the original group assignments nor the acceptance criteria contain any mechanism preventing very similar parameter sets from being generated. While further investigation is likely warranted, we have not observed this to be an issue in practice. In addition, it is possible that correlations across parameters exist in PPSO PoMs that are absent from those generated by random search. Although we have not observed such correlations thus far, further study may be necessary.

Overall, PPSO offers an efficient and robust alternative to previous methods of cardiac PoM generation under specific biomarker constraints with potential applications in research as well as clinical modeling.

Acknowledgments

This study was supported by NIH grant 2R01HL143450.

References

[1] Sarkar AX, Christini DJ, Sobie EA. Exploiting mathematical models to illuminate electrophysiological variability between individuals. The Journal of Physiology 2012; 590(11):2555–2567.

[2] Britton OJ, Bueno-Orovio A, Van Ammel K, Lu HR, Towart R, Gallacher DJ, Rodriguez B. Experimentally calibrated population of models predicts and explains intersubject variability in cardiac cellular electrophysiology. Proceedings of the National Academy of Sciences 2013; 110(23):E2098–2105.

[3] Sánchez C, Bueno-Orovio A, Wettwer E, Loose S, Simon J, Ravens U, Pueyo E, Rodriguez B. Inter-subject variability in human atrial action potential in sinus rhythm versus chronic atrial fibrillation. PloS One 2014;9(8):e105897.

[4] Llopis-Lorente J, Trenor B, Saiz J. Considering population variability of electrophysiological models improves the in silico assessment of drug-induced torsadogenic risk. Computer Methods and Programs in Biomedicine 2022; 221:106934.

[5] Fenton F, Karma A. Vortex dynamics in three-dimensional continuous myocardium with fiber rotation: Filament instability and fibrillation. Chaos 1998;8(1):20–47.

[6] Ten Tusscher KH, Noble D, Noble PJ, Panfilov AV. A model for human ventricular tissue. American Journal of Physiology Heart and Circulatory Physiology 2004; 286(4):H1573–H1589.

[7] Ten Tusscher KH, Panfilov AV. Alternans and spiral breakup in a human ventricular tissue model. American Journal of Physiology Heart and Circulatory Physiology 2006;291(3):H1088–H1100.

[8] Rush S, Larsen H. A practical algorithm for solving dynamic membrane equations. IEEE Transactions on Biomedical Engineering 2007;25(4):389–392.

[9] Zeng J, Laurita KR, Rosenbaum DS, Rudy Y. Two components of the delayed rectifier K+ current in ventricular myocytes of the guinea pig type: Theoretical formulation and their role in repolarization. Circulation Research 1995; 77(1):140–152.

[10] Cairns DI, Comstock MR, Fenton FH, Cherry EM. CardioFit: a WebGL-based tool for fast and efficient parametrization of cardiac action potential models to fit user-provided data. Royal Society Open Science 2025; 12(8):250048.

Address for correspondence:

Elizabeth M. Cherry
756 West Peachtree Street NW, Atlanta, GA 30332
elizabeth.cherry@gatech.edu