Research

Author for correspondence:
Nicholas S. Peters
e-mail: n.peters@imperial.ac.uk

[†]These authors contributed equally to this work.

Machine learning methods for locating re-entrant drivers from electrograms in a model of atrial fibrillation

Max Falkenberg McGillivray[1,2,†], William Cheng[1,2,†], Nicholas S. Peters[3] and Kim Christensen[1,2,3]

[1]The Blackett Laboratory, and [2]Centre for Complexity Science, Imperial College London, London SW7 2AZ, UK
[3]ElectroCardioMaths Programme, Imperial Centre for Cardiac Engineering, Imperial College London, London W12 0NN, UK

MFM, 0000-0002-2986-2494

Mapping resolution has recently been identified as a key limitation in successfully locating the drivers of atrial fibrillation (AF). Using a simple cellular automata model of AF, we demonstrate a method by which re-entrant drivers can be located quickly and accurately using a collection of indirect electrogram measurements. The method proposed employs simple, out-of-the-box machine learning algorithms to correlate characteristic electrogram gradients with the displacement of an electrogram recording from a re-entrant driver. Such a method is less sensitive to local fluctuations in electrical activity. As a result, the method successfully locates 95.4% of drivers in tissues containing a single driver, and 95.1% (92.6%) for the first (second) driver in tissues containing two drivers of AF. Additionally, we demonstrate how the technique can be applied to tissues with an arbitrary number of drivers. In its current form, the techniques presented are not refined enough for a clinical setting. However, the methods proposed offer a promising path for future investigations aimed at improving targeted ablation for AF.

1. Introduction

Atrial fibrillation (AF) is the most common cardiac arrhythmia in clinical practice and is getting more prevalent in the general population due to the ageing demographic. However, the mechanistic origin of AF is still poorly understood. As a result, the success rates of treatment options remain limited with future improvements requiring a better understanding of how AF emerges from the underlying properties of the myocardium.

A variety of possible mechanisms have been proposed to explain the origin of AF. These include circus movement re-entry, the leading circle theory, spiral wave re-entry (otherwise known as rotors), of which micro-anatomical re-entry can be thought of as a subset, and the multiple wavelet hypothesis [1,2]. However, there are many contradictory findings, and no one mechanism explains all observations [3,4]. This suggests new techniques are needed both in clinical practice and research, with numerous researchers highlighting the potential of computational simulations and machine learning [5–10]. An issue of particular importance is that of limited mapping resolution when detecting the drivers of AF. Errors in the accuracy of imaging data limit the efficacy of treatment by ablation, causing disagreement when interpreted by the research community [11–13].

In this paper, we present a method whereby electrograms are extracted from a number of independent locations in the atria and these are cross-referenced to triangulate the position of a re-entrant circuit. By using multiple measurements, noise and local fluctuations are minimized and a prediction can be reached without being overly reliant on the imaging resolution at any one given location. The procedure applies machine learning methods to maximize the prediction accuracy of the algorithm. The work here should be considered a proof of concept for the efficacy of these methods in simple models of AF. The work does not consider some of the details that would be necessary in a realistic clinical implementation or that are found in more detailed, continuous models of AF [6]. Nevertheless, it presents a clear path for future investigations aiming for a potential improvement in ablation success rates based on locating drivers from electrograms only.

The Christensen–Manani–Peters model (CMP) is a two-dimensional cellular automaton on an architecture mimicking that of the atria [14]. While the model architecture is simplified, it preserves the key features of discrete cardiomyocyte activation at the microscopic level, while ensuring macroscopic conduction appears continuous. The CMP model has a particular focus on demonstrating the role of fibrosis in initiating and maintaining AF. AF is driven through the spontaneous emergence of re-entrant circuits which generate rapid, irregular atrial activity. This is reflected in the electrograms generated from the model which are fractionated during AF [15]. This mechanism is a form of micro-anatomical re-entry. These circuits form at regions with high levels of localized fibrosis. The model also explains the wide range of AF classifications from short intermittent episodes to long-lasting permanent AF as increasing levels of fibrosis modify the myocardium [15]. In addition to the mechanistic benefits of cellular automata, their popularity has recently grown given their capacity to simulate much longer timescales of atrial activity than computationally intensive continuous models [16–18]. Cellular automata models are particularly successful for allowing detailed mathematical analysis, while preserving a very wide range of complicated phenomena [19]. Particularly in the case of AF, cellular automata are effective at reproducing behaviour due to structural discontinuities, such as electrical signals being blocked due to fibrosis. However, these benefits come at the cost of not reproducing realistic phenomena in clinical AF such as the curvature of electrical wavefronts or the complex topology of the human atria [20]. These phenomena are more typically studied using continuous models of AF [6].

Machine learning approaches are statistical techniques implemented computationally which excel at finding hidden insights in complex data without being explicitly programmed to do so. In particular, machine learning has had recent successes in the medical community in classifying skin melanoma and in genetic sequencing [21,22]. These methods often rely on having a large dataset for the models to learn correlations (patterns) from. In clinical medicine, the lack of good-quality data in large quantities often makes such an approach difficult. However, when working with computer simulations such as those in the CMP model, data can be generated in sufficient quantities, and hence, the statistical insights these models provide can offer significant improvements when analysing data. The computational complexity of continuous models make these unsuited to generating large quantities of data [6,16]. Despite this, continuous models are still often considered preferable when studying AF computationally. However, two additional challenges specific to this work warrant the use of a simple cellular automaton over more realistic models. Firstly, generating electrograms from a model of AF using modern graphics processing unit (gpu) optimized methods is significantly faster for cellular automata than what would be possible for continuous models of AF on a complex topology. Secondly, the training process used by our supervised machine learning algorithm relies on accurately knowing the position of simulated re-entrant circuits. This can be trivially achieved in a cellular automata model of AF. However, it is not clear how re-entrant circuits, or rotors, could be easily generated at known locations in continuous models of AF without resorting to computationally costly adaptations.

The remainder of this paper is organized as follows. In §2, we briefly introduce the CMP cellular automata model used in this research. We describe the physiological motivation behind the model and

outline the process of simulating electrograms. We discuss a novel visualization of electrograms and use insights from this method to inform statistical analysis. Random forests, a machine learning technique with consistently strong results across a number of domains, are described in §3 [23,24]. These are then implemented in a search algorithm for locating AF drivers. Our results are presented and discussed in §4 and §5. Finally, we outline potential future work and discuss the CMP model's relevance to current clinical research.

2. The Christensen–Manani–Peters model

The Christensen–Manani–Peters (CMP) model is a two-dimensional cellular automaton. Each atrial muscle cell in the CMP model is represented by a single square in a larger square lattice of side length L. All cells are connected to their nearest neighbours in one dimension (the longitudinal direction) and in the orthogonal dimension with probability ν (stochastic connections, the transverse direction). This construction is a simplified computational implementation of the real myocardium, but it preserves the essential myocardial architecture which ensures that electrical impulses travel preferentially along muscle fibres rather than laterally between fibres. The strong coupling in the longitudinal direction represents individual muscle cells forming a single, long muscle fibre. The reduced coupling in the transverse direction represents the reduced electrical conductivity between neighbouring muscle fibres. The parameter ν controls the strength of the transverse coupling. This is manifested in the anisotropy of electrical conduction velocities. There are periodic boundary conditions in the transverse direction (across different muscle fibres) and open boundaries in the longitudinal direction (along a single muscle fibre) representing a simplified cylindrical geometry of the atria [14]. The cells along the left boundary are the pacemaker cells which are excited every T time steps.

In the CMP model, the electrical cycle of a given cell is as follows: Cells are updated in discrete time steps. Initially a cell is at rest. In this state, the resting cell can become excited at the next time step by a neighbouring excited cell. The cell is in the excited state for one time step after which the cell is in the refractory state for the next τ time steps. In the refractory state, a cell is unexcitable after which the cell re-enters the resting state. This cycle is indicative of the action potential of real cardiomyocytes. The beating of the atria is simulated by exciting all the cells along the left boundary of the tissue. This signal can then propagate across the tissue until it is dissipated at the opposite open boundary. The pacemaker cells are activated every T time steps. The full excitation rules of the CMP model are summarized in the box below.


CMP algorithm

1. All cells are connected longitudinally forming a muscle fibre.
2. A fraction ν of the transverse connections between cells are filled linking two muscle fibres.
3. A resting (excitable) cell connected to an excited cell at time step t will become excited at time step $t + 1$.
4. A refractory (unexcited) cell connected to an excited cell at time step t will not become excited at time step $t + 1$.
5. An excited cell at time step t becomes refractory at time step $t + 1$, and remains refractory for a duration of τ time steps before returning to the resting state.
6. Every T time steps, the pacemaker cells along one of the open boundaries of the muscle tissue are excited.
7. A small fraction δ of cells are dysfunctional and have a small probability ϵ of not getting excited.

Dysfunctional cells can block propagation along the cell strands and leave an opening for the propagating wavefront to turn back in on itself forming a circuit (re-entry). If this circuit is long enough (which occurs when ν is sufficiently small), the signal can continuously propagate around the circuit forming a persistent driver of AF (figure 1). Figure 2 shows the emergence of AF on a tissue-wide scale. In the real myocardium, the transverse decoupling of muscle fibres is associated with the build-up of fibrosis [25–27]. Hence, the parameter ν can be thought of as a control for the degree of fibrosis in the myocardium. We demonstrate this explicitly in previous work on the CMP model, where we have shown how the increased prevalence of fibrosis (decreasing ν) increases the risk of AF persisting by increasing the average path length of re-entry circuits, in agreement with clinical observations [14,15].
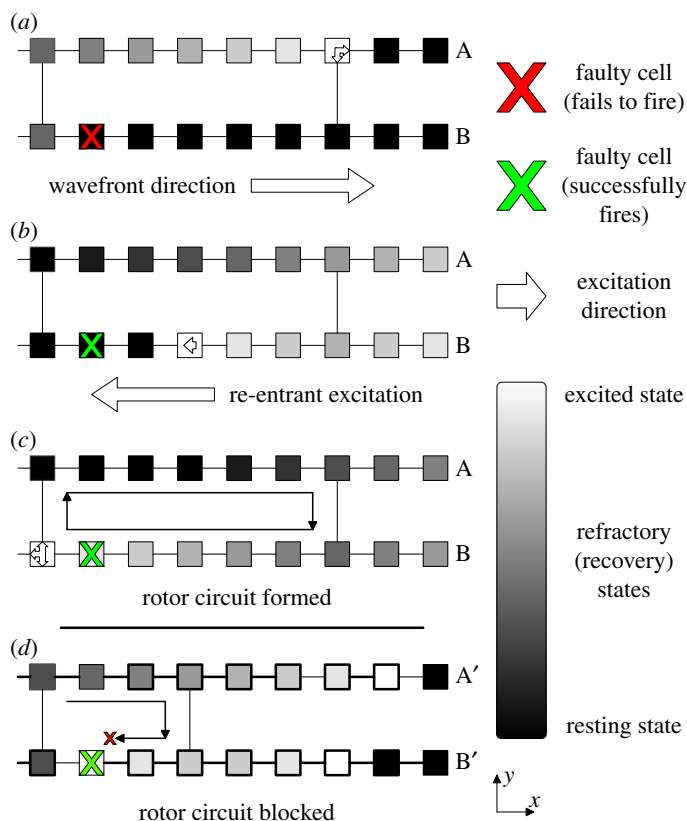
Figure 1. The formation of a simple re-entrant circuit at the cellular level. An excited cell is shown in white and resting (excitable) cells are black. Cells shown in grey are refractory (unexcitable) for the duration of the refractory period. All cells are coupled to their nearest neighbours longitudinally—this reflects the strong coupling of cells along muscle fibres. Transverse couplings exist with probability ν—reducing ν reflects the decoupling between adjacent muscle fibres caused by fibrosis. An excitation is initiated along the left wall of the heart tissue and propagates left to right. When the excitation reaches a dysfunctional cell which fails to fire (marked by a red cross), (*a*) the excitation in fibre B is blocked but excitations continue in fibre A above. When a coupling between the excited and blocked strands is reached, the excited cell can send a signal propagating backwards from right to left down the blocked strand (*b*). If the path length of the re-excitation is sufficiently long, the re-entrant excitation can excite tissue behind the main wavefront. This signal can then move to the adjacent strands forming a continuously re-excited circuit. The simple re-entrant circuit shown here is rectangular in shape and is formed from two fibres (A & B) and a single dysfunctional cell. More complicated re-entrant circuits can consist of multiple fibres and multiple dysfunctional cells. The extract shown in (*d*) is for a different region of heart tissue. Here, the same mechanism as above attempts to form a re-entrant circuit. However, the circuit path length is insufficiently long such that the re-entrant excitation is blocked by the refractory cell to the left of the dysfunctional cell marked with a cross. Hence, a continuously excited circuit cannot form in this tissue segment. This behaviour explicitly links the emergence of re-entrant circuits with regions of high fibrosis (low ν). Note that the mechanism described here is for spontaneously generated circuits. The work in this paper artificially inserts circuits into the substrate such as the one shown in (*c*). This ensures that we can keep track of the location of all circuits in the model for analysis.

It is helpful to understand the considerations behind the dysfunctional cell mechanism. The important detail is that the AF mechanism in the CMP model is possible using any mechanism of unidirectional conduction block—the stochastic failure of cells to depolarize is a simple way to include this effect in the model. In the original CMP model, the 1000×1000 grid is coarse grained into a 200×200 grid for computational ease. A 1000×1000 grid would approximately account for the total number of cardiomyocytes on the epicardial surface of the atrium. The model then treats each cell in the coarse-grained grid as a single individual muscle cell. For single cardiomyocytes in isolation, there is little clinical evidence for the cell stochastically failing to depolarize. However, instead of considering the CMP model as a microscopic depiction of atrial conduction, we can think about a mesoscopic picture where each cell in the 200×200 grid represents the average dynamics of the underlying 5×5 block of cells. These coarse-grained blocks can still follow the same branching/connectivity rules between cells as originally formulated for the CMP model.
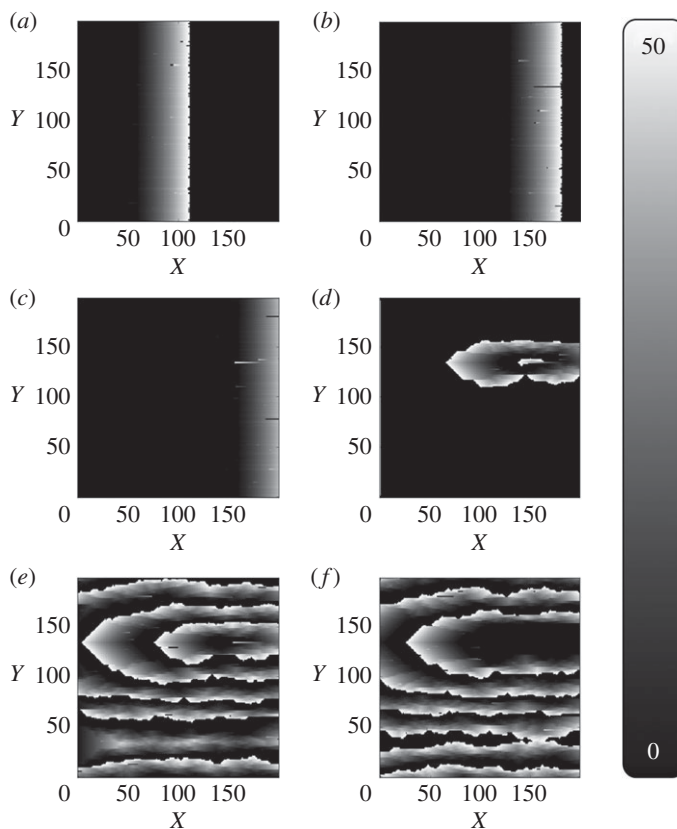
Figure 2. The formation of fibrillatory re-entrant circuits in the Christensen–Manani–Peters model for a refractory period of $\tau = 50$. Excited cells are shown in white, resting (excitable) cells are black and cells which are refractory (unexcitable) are shown in grey. (*a*) A planar wave initiated by the pacemaker cells at the left boundary of the cylindrical heart tissue propagates from left to right along muscle fibres. (*b*) A dysfunctional cell fails to fire, blocking the propagation of the signal along that fibre. Excited cells in the adjacent muscle fibre can re-excite the blocked fibre at the next vertical coupling (figure 1). (*c*) If the path length of the re-excitation is sufficiently long, the re-entrant excitation can escape the unexcitable region. (*d*) This escaped excitation initiates a continuously activated circuit from which chaotic waves propagate—this circuit is the driver of AF. (*e*) Fibrillatory waves spread across the whole tissue and prevent the formation of new planar wavefronts at the pacemaker cells on the left boundary. (*f*) If a dysfunctional cell in the re-entrant circuit fails to fire, the circuit is stopped and fibrillatory behaviour dissipates.

Within this 5×5 cell block, we can consider the possibility of unidirectional conduction block due to the imbalance between current sources and sinks—the possibility of such a mechanism has been shown by Bub *et al.* [28] in a theoretical model and has also been supported by clinical evidence [29–31]. This explains why we might expect any given block of cells to display unidirectional conduction block with some small probability, ϵ, given a suitable geometric arrangement of cells with a small probability, δ. It is also sensible to consider such an effect to be stochastic because leaking current over a number of activation cycles can push a previously blocked group of cells over the depolarization threshold.

The parameters used in this work are $\nu = 0.2$, $\tau = 50$ and $T = 220$. We fix the coupling fraction to be $\nu = 0.2$, because this is the largest coupling fraction at which we typically observe paroxysmal AF [14,15]. It is important to note that all analysis in this paper is carried out at this fixed coupling fraction. However, the model dynamics do not suggest that we would deduce largely different results for lower coupling fractions [14]. AF is not typically observed at higher coupling fractions.

The shortest timescale in the CMP model is associated with the excitation timescale of a single cardiomyocyte of $\Delta t = 0.6$ ms. In the course-grained tissue, this corresponds to an excitation timescale for each cell of $\Delta t^* = 5 \times 0.6$ ms $= 3.0$ ms. The time step of 3 ms approximately corresponds to the time taken for an electrical signal to cross a 5×5 block of cardiomyocytes in the real atrium along a muscle fibre. The refractory period is chosen to be $\tau = 50$ time steps where the unit of time is Δt^* such that the action potential of a cardiomyocyte is 150 ms in accordance with the values seen for human atrial cardiomyocytes. In sinus rhythm the CMP model therefore beats approximately once every 660 ms. Note that AF can easily be induced with a different set of parameters, but these are chosen as a realistic

reflection of atrial activity. In the original CMP model, the fraction of dysfunctional cells is $\delta = 0.05$. This is an arbitrary choice which allows AF to spontaneously emerge. In this work, we set $\delta = 0$ and instead artificially insert re-entrant circuits into the tissue. This is a reasonable concession for computational ease because this work is not concerned with the spontaneous emergence of AF but rather the diagnosis of AF after it has emerged. More specifically, the search algorithm presented in §3 only analyses two cycles of electrical activity during AF. This corresponds to less than 400 ms in a realistic clinical setting. Over this short timescale, it is reasonable to approximate the dynamics of transient rotors as stable.

3. Material and methods

A benefit of the approach studied here is that instead of relying on a single accurate measurement to locate a re-entrant circuit, a collection can be used to determine the circuit's location. This means that the electrograms can be measured at a lower signal to noise ratio as no one measurement is critical to the success of the locating algorithm. Sample electrograms from the CMP model can be seen in figure 3. These are generated as outlined in appendix A. To locate AF drivers in the CMP tissue, machine learning models are used. These models are trained with electrogram feature data gathered from a large number of simulations of varying AF instances in CMP tissue. Given the preliminary nature of the work, a simple set of features are used. These are outlined in appendix B.

For the purpose of recording training data to build the machine learning models, a 9-electrode probe is used, arranged in a 3×3 array. This represents the multi-contact mapping catheters used for ablation. Each electrode is unipolar, but electrograms can be compared across the probe as an analogue of bipolar electrograms. The 3×3 probes are distributed uniformly in a 7×7 cell grid of the CMP tissue as seen in figure 4. The distance between each probe and its nearest neighbours is 3 cells. This gives a resolution comparable to clinically used ablation mapping catheters where electrodes are approximately 1 mm long and are spaced 4 mm apart [32]. We define the probe to be on the re-entrant circuit if any part of the re-entrant circuit is within the multiprobe's 7×7 cell region.

A first look at the electrograms shows clear distinctions depending on the electrode's relative position from the re-entrant circuit, as shown in figure 3b. Features gathered from these individual electrodes include: maximum voltage, minimum voltage, first stationary point position, mean voltage, waveform skewness and other common statistical measures. A full list can be found in appendix B. With the 9-probe array, the wavefront from the AF circuit will reach each individual electrode at a different time. Therefore, gradients of the features can be measured across the whole electrode array in different directions giving more effective information on the position of the AF circuit (figure 4). Note that local gradients are quite robust measures as they do not rely on any global calibration procedure.

Visualizing the features from electrogram data can give significant insight into the electrical dynamics of the CMP model without extensive statistical analysis. This can be done using a visualization we have coined the vector feature map (VFM). Their creation, analysis and general features are described in detail in appendix B. A VFM shows the average value of a feature at a given vector relative to the centre of the re-entrant circuit in many different instances of the CMP model. An example is shown for the magnitude of the dominant Fourier transform frequency of the electrogram signal in figure 5.

To locate AF drivers using electrogram features, a supervised machine learning technique known as the random forest model is used [33]. The model is capable of giving quantitative (the distance between the probe and the re-entrant circuit) and qualitative (the probe is/is not currently on the re-entrant circuit) responses when given a set of electrogram feature information [34]. The random forest model was chosen due to its relative effectiveness compared to other machine learning models. The method has recently been used for problems such as tissue segmentation in the brain and the classification of heart failure subtypes [35,36]. Caruana & Niculescu-Mizil [23] and Caruana *et al.* [24] note that the random forest model has one of the highest average performances of any machine learning method across a wide range of different problems. Preliminary testing on our data showed considerably higher success rates for random forests than for other simple machine learning algorithms. The mathematical background for random forests is described in appendix C.

The training data for the random forests were gathered from 5000 randomly generated CMP tissues with one randomly placed AF circuit.[1] The fraction of transverse connections was chosen to be $\nu = 0.2$ as this is the critical point where instances of paroxysmal AF emerge in the CMP model [14]. Each tissue had 64 multiprobe electrodes uniformly placed giving a total of 2 880 000 electrogram recordings.

---

[1]The choice of 5000 samples is somewhat arbitrary; early trials showed 1000 samples as the minimum number of samples required for good results and 5000 samples were found to be consistently sufficient without taking excessively long to generate.
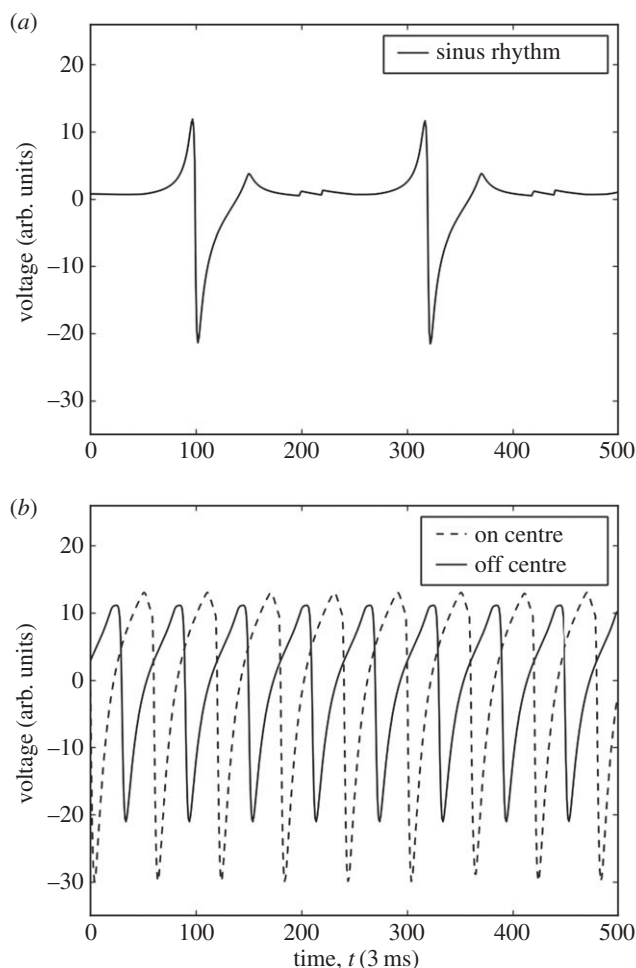
Figure 3. Electrograms simulated on perfect isotropic tissue ($\nu = 1.0$) in the CMP model using equation (A 1). Voltage is measured in arbitrary units. Time is measured in steps of 3 ms such that at sinus rhythm the pacemaker cells activate every 660 ms (220 time steps). (*a*) An electrogram recorded during regular sinus rhythm. The large depolarizations at approximately $t = 100$ and $t = 320$ correspond to the planar wavefront crossing the electrogram recording probe. The small fluctuations at approximately $t = 200$ and $t = 420$ correspond to the activation of a new planar wavefront at the left boundary by the pacemaker cells and the dissipation of electrical activity at the open boundary on the right after the wave propagates through the tissue—this is a small discretization effect in the CMP model. (*b*) Electrograms recorded during the rapid pacing of the heart where electrical wavefronts originate from a single point in the tissue. The dashed and solid lines show electrograms recorded at the centre of the electrical activity and 30 cells displaced from the centre, respectively. There are clear differences between electrograms recorded at different locations relative to the centre of electrical activity. Note that the difference between electrograms is significantly less pronounced in imperfect, anisotropic tissue ($\nu < 1.0$)—this warrants the use of statistical techniques for analysis. These fractionated electrograms are presented in [15].

3.1. Algorithm for locating re-entrant circuits

The goal of our algorithm is to demonstrate a proof of concept where re-entrant circuits driving AF can be located using solely electrogram information. The aim is not to create a perfect model which could be directly transferred to more complicated scenarios, but rather, the aim is to show the feasibility of these methods for electrical mapping in a system with large local fluctuations. As part of this approach, a small number of simplifications are applied to the CMP model to simplify simulations and analysis.

All tissue instances are generated at a fixed transverse coupling fraction of $\nu = 0.2$, which is approximately the degree of fibrosis at which we observe paroxysmal AF. We also work in the low noise limit where $\delta = 0$. As a result, temporary critical circuits cannot spontaneously form in the tissue as shown in figure 1*c*. The limitations of this decision are discussed in §5.1. The low noise limit inhibits the formation of new, stable re-entrant circuits. Instead, a single circuit is artificially constructed by picking a random point in the CMP tissue and removing the coupling to the muscle fibre above and below of the 28 cells to the right of the randomly chosen cell. This gives a rectangular circuit of path length 60—this is an
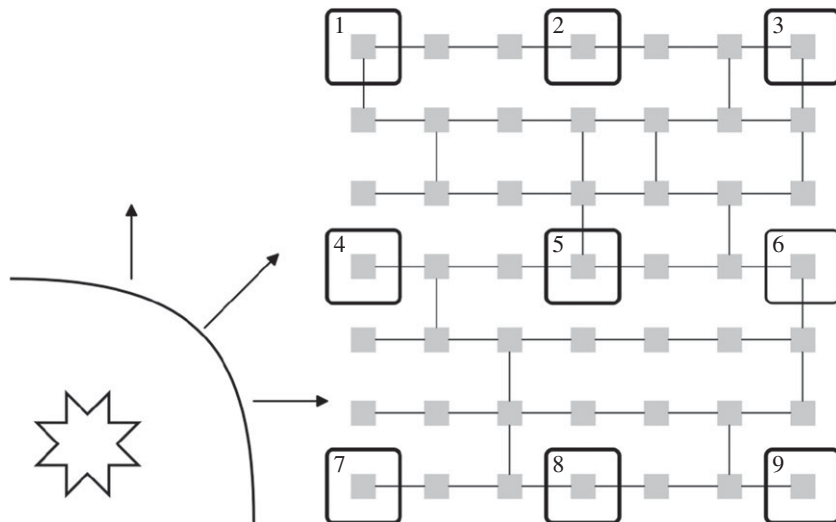
Figure 4. Electrical wavefronts propagate in all directions from an AF driver (shown as the star in the bottom left). A 3 × 3 electrode grid is used to record 9 simultaneous electrograms (large boxes numbered 1–9) across a region of 7 × 7 cells in the CMP model (small squares shown in grey). By spacing out the electrodes, the approaching electrical wavefront will cross each of the electrodes in a slightly different order and in a slightly different direction, affecting the feature behaviour seen at each electrode. Hence, taking the gradient of features across the electrode grid can give detailed information about the wavefront flow across the grid, enhancing our knowledge of the driver's position. The local zig-zagging of the electrical wavefront due to the discrete structure of the CMP models can occasionally result in errors when calculating the direction of global wavefront propagation. However, by taking the mean value of the gradient of features for a fixed unit vector across different parts of the electrode grid these errors are minimized. Any remaining discrepancies are accounted for using the constraint mechanisms described in §3.1 and appendix D. The intermediate electrode spacing used here is reflective of clinically used mapping catheters, and offers a balance between maximizing local information while ensuring that the global wavefront direction can be accurately calculated.
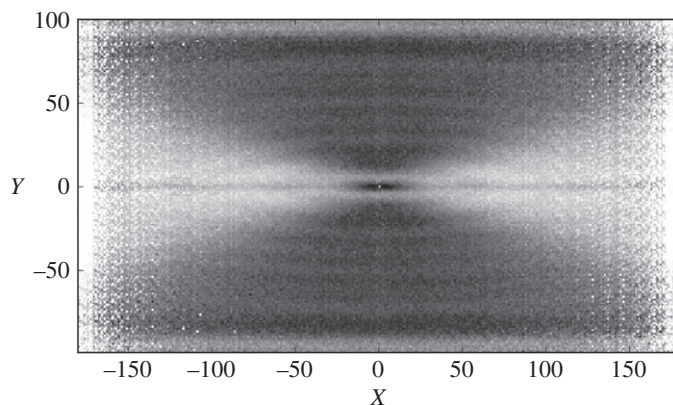
Figure 5. The vector feature map (VFM) for the dominant frequency of the electrogram's Fourier transform on an arbitrary scale, shown by the greyscale. The image is generated from tissues with a single ectopic cell beating every 60 time steps placed at a random location in the tissue with transverse coupling fraction $\nu = 0.2$. The image shows strong separation between regions of unidirectional wavefront propagation (light grey) and bidirectional wavefront propagation (dark grey). There is also a strong indicator of the driver's centre, black region at (0,0). Notice that the image is symmetric across the X- and Y-axes and that different feature maps highlight different wavefront dynamics. There is significant differentiation between bulk regions but little differentiation in small, localized regions. Combining multiple vector feature maps can improve local differentiation. More details can be found in appendix B.

arbitrary choice which is slightly larger than the refractory period of $\tau = 50$. The circuit is then artificially forced to start driving AF by the introduction of a single ectopic beat in the circuit. This is achieved by presetting the cell states in the circuit as is shown in figure 1b without any dysfunctional cells—the circuit is then allowed to evolve naturally as in the original CMP model.
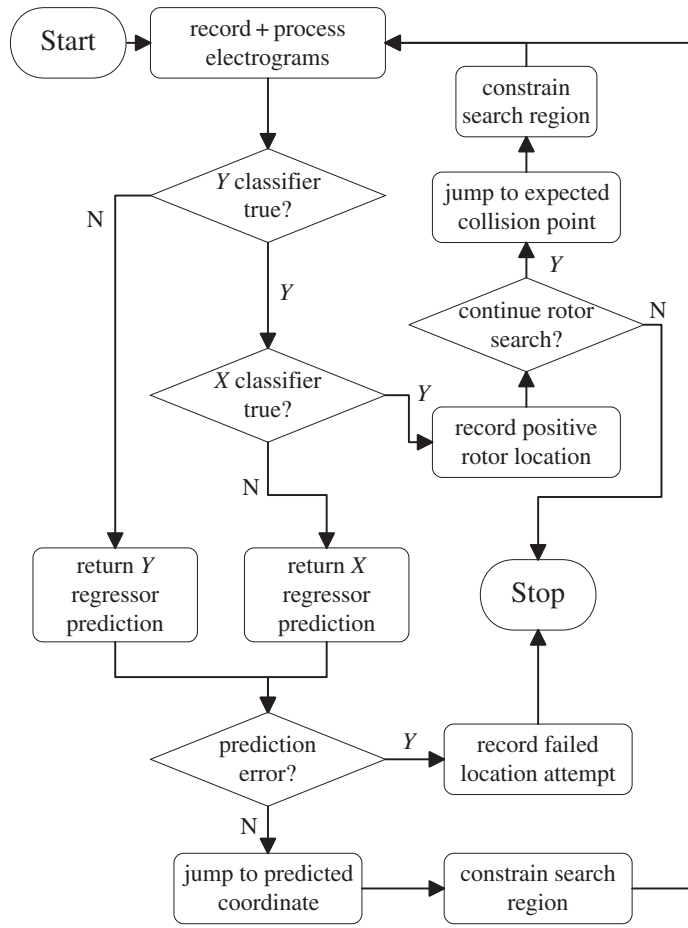
Figure 6. A flow chart describing the basic AF driver location algorithm. When starting the algorithm, the simplified CMP model is run with one or two drivers placed at random locations—the algorithm could easily be extended to more drivers if desired as outlined in figure 7. A random position is chosen to record the first set of electrograms. Search regions are constrained as described in appendix D. Predictions from the X and Y regressors are forced to lie within search constraints. 'Prediction error' initiates if the constraint region becomes too small or predictions loop between repeated coordinates.

The procedure to be tested is split into two sections. First, the CMP tissue must be initialized to generate fibrillatory behaviour. Once initialized, the data aquisition and processing cycle can be initiated. The initialization phase is outlined as follows and corresponds to the 'Start' block in figure 6.

1. Generate an instance of the CMP model at $\nu = 0.2$ of linear size $L = 200$. Do not create any dysfunctional cells which could generate unexpected re-entrant circuits (i.e. $\delta = 0$).
2. At a random location in the tissue, insert a single simple re-entrant circuit consisting of two cell strands. Preset cell states as shown in figure 1b excluding the dysfunctional cell.
3. Allow fibrillation to commence and continue until the whole tissue has been reached by the electrical activity from the re-entrant circuit. Fibrillation is defined as when more than 220 cells are active in the substrate ($1.1 \times L = 220$)—this definition effectively distinguishes between fibrillation and sinus rhythm as discussed by Manani *et al.* [15].

After the CMP model is initialized, the data acquisition and processing cycle can be used to generate predictions of the expected location of re-entrant circuits in the tissue. This is illustrated in figure 6 and broadly follows the steps below. As can be seen from the lack of local differentiation in the electrogram feature in figure 5, it is a non-trivial problem whether or not electrogram dynamics can be used to infer the position of re-entrant circuits drivers which motivates the use of a recursive method involving multiple measurements.

4. At a random location in the tissue, place a 3×3 array of electrode probes and generate electrograms at each position.

5. Extract statistical information from the electrograms and preprocess data for compatibility with required machine learning data structures.

6. Process data using machine learning models to output a prediction for the expected electrogram location.

7. Calculate the constraints of possible prediction locations based on the mechanism shown in appendix D.

8. Post-process prediction data to abide by the calculated constraints for the final prediction using the mechanism shown in appendix D.

9. Record a new set of electrograms at the predicted position of the re-entrant circuit. If the electrogram behaviour is consistent with the expected statistical features at a re-entrant circuit, either end the algorithm or proceed to searching for any remaining circuits. If the position is not consistent with the expected behaviour of a re-entrant circuit, return to step 6 and repeat the prediction process.

The algorithm uses four random forest models, two classification and two probabilistic. The classification models are used to check if the probe is positioned on the driver's *X*- and *Y*-axes. The responses for the two probabilistic models are the probabilities of the driver lying on each transverse cell column on the *X*-axis and each longitudinal cell strand on the *Y*-axis—this is an adaptation of the regression style models typical in random forests in which a continuous scale is broken into discrete ranges and each of the ranges is considered its own class. The benefits of the probabilistic approach over typical regression models is that it is easier to implement additional levels of post-processing in any predictions. The probability of the re-entrant circuit lying in each class can then be processed to infer the predicted position of the re-entrant circuit in a given direction. The probabilistic models can be made from classification (qualitative) models where instead of the majority rule, the response is given by

$$P_i = \frac{k_i}{k}, \quad \text{with } k = \sum_{i=1}^{K} k_i, \tag{3.1}$$

where P_i is the probability for the particular response i, k_i is the number of data samples that have response i and k is the total number of responses. These probabilities do not consider that electrogram measurements may have been taken previously which gave information as to the direction of the re-entrant circuit from particular regions in the tissue. The post-processing procedures described in appendix D account for this.

In a general case, it may occur that there are multiple re-entrant circuits present in a single tissue. To test our methods for this scenario, we repeated the procedure outlined at the start of this section with a simple change that two re-entrant circuits are randomly placed in the tissue (with a minimum separation of 10 cells vertically to avoid overlap). We then adjust the search algorithm to start looking for a second re-entrant circuit after the first is found. However, note that the machine learning models used are still only trained on simulations with a single re-entrant circuit—the flexibility of our model to search for multiple drivers without explicitly learning to do so is one of the major benefits of our approach. The adaptation to multiple drivers is possible because of the limited interference in the electrical activity between competing drivers. In a system with two stable drivers, tissue closer to the first driver exhibits electrical activity which can be closely approximated by the activity one would expect if the second driver was not present. This is demonstrated in figure 11 in appendix B. This principle can be implemented in our algorithm as outlined in figure 7 where the extension to tissues with multiple drivers is also described.

4. Results

The features used by the machine learning models are described in appendices B and C. All features described have some discriminatory capacity for the classifier models. Not all features are strictly necessary for generating good results (as determined using the Gini index, which is discussed in appendix C); however, the low computational cost associated with the different statistical measures warrants the inclusion of the full set of features. The most discriminatory features include the wavefront direction, the maximum and minimum electrogram voltages, the intensity of the electrogram signal, the maximum value of the voltage gradient and the magnitude of the largest Fourier transform frequency. For an example model, the feature ranking is shown in figure 13 in appendix C. As noted in §3, local noise restricts the efficacy of any individual classifier. Hence, the results presented here are based on the rotor
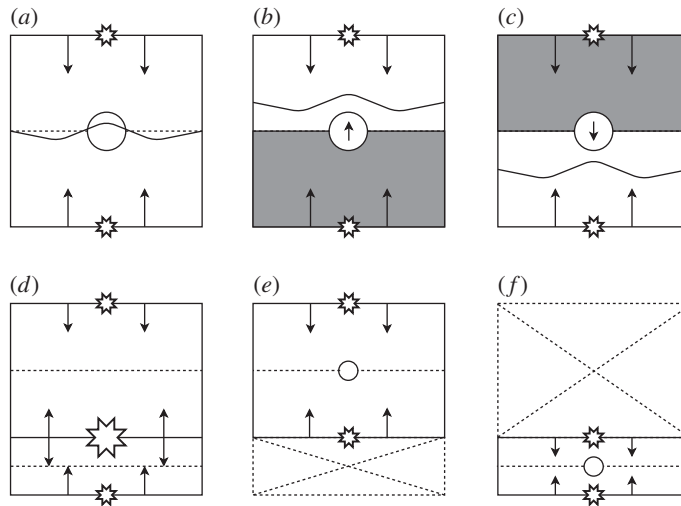
Figure 7. A diagram highlighting the key mechanism behind the extension of a single driver location algorithm to an arbitrary number of drivers. (*a*) A single driver tissue where the driver has been placed at the circular boundary to indicate wavefronts propagating from above and below (small stars). Wavefront collisions are expected (dotted line) and observed (wavy line) at the centre of the tissue. (*b/c*) The grey regions indicate the search regions for a second driver inferred by observing the absence of wavefront collisions at the expected position. (*d*) A second driver is found (big star). This two-driver system can be conceptually split into two single driver tissues, (*e/f*). These can be searched using the same recursive strategy starting at (*a*). Note that, in the current work, the mechanism for finding multiple drivers is simplified to only predict the position of expected collisions in one dimension; however, this could easily be extended to the full geometry in future work.

Table 1. The success rate of the algorithm described in figure 6. The results are generated from 1500 simulations each of the single and double driver tissues. The prediction success corresponds to the classifier models outputting a final positive prediction which lies on a position which if ablated would terminate the re-entrant circuit. The mean number of jumps refers to the total number of electrogram grid recordings required before reaching a final prediction—i.e. for each jump, one electrogram from each point on the 3×3 grid is generated over two fibrillatory cycles. Features are then extracted from these electrograms, processed and used to predict a displacement from a re-entrant circuit as outlined in §3.1. Note that, for the second circuit, the number of jumps required is in addition to those required to find the first circuit.

target driver	prediction probability	mean jumps
one circuit	95.4%	5.0 ± 1.7
two circuits, circuit 1	95.1%	5.2 ± 2.2
two circuits, circuit 2	92.6%	6.1 ± 4.2

search algorithm which combines multiple separate machine learning models and a series of constraint mechanisms to make predictions.

Table 1 shows the results from the driver search algorithm. The prediction success is consistently high both for the single-circuit and two-circuit scenarios. Additionally, the prediction is reached after a small number of electrogram recordings, typically approximately 5–6. For comparison, the probability of recording an electrogram on the re-entrant circuit at any given point in the tissue is just under 2%.

The results indicate that the algorithm is capable of locating simple re-entrant circuits on randomly generated CMP tissues with excellent prediction rates from a small number of recording locations. These results support the possibility of using statistical learning techniques for locating AF drivers on heart tissue using indirect feature measurements, derived from the accessible electrograms.

5. Discussion

It is clear from the results in §4 that a statistical analysis of electrograms can be used to extract information as to the location of re-entrant drivers in the CMP model, an idealized mathematical model of AF focusing on the discretized structure of the atrial myocardium and the electrophysiological action of fibrosis.

In recent work, we have reviewed issues concerning the limited resolution of mapping technologies used during ablation, and have found that these limitations make it hard to identify distinct mechanisms of AF in clinical practice [11]. If re-entrant circuits are to be ablated, terminating AF, these issues must be addressed. However, the methods proposed here suggest a promising line of inquiry for future work in more realistic models of AF, with a final aim that these approaches may mitigate some of the mapping resolution issues currently observed during catheter ablation. Such an approach maximizes accuracy by minimizing errors due to local noise.

5.1. Future work

In its current form, our driver-mapping approach is not refined enough for use in clinical practice—the simplifications in the CMP model do not accurately represent real electrical behaviour in the heart.

Limitations in the current work come under three categories: (a) intrinsic limitations in the original CMP model, (b) limitations in our simplified implementation of the original CMP model, and (c) limitations arising in the data analysis and machine learning procedures.

(a) The CMP model is not a fully realistic model of the atria. It represents a simplified two-dimensional cylindrical topology of the myocardium and has all cells arranged in a consistent square lattice. The refractory period of cells are uniform and fixed, and the cell-to-cell conduction velocities at the microscopic level are constant, traversing one cell-to-cell coupling per time step. In this current form, the CMP model is not suited to studying some of the rate-dependent effects typically studied in continuous models of AF. However, the CMP model's explicit focus on tissue anisotropy and the presence of fibrosis demonstrates that the key features of AF can spontaneously arise without the need for the extra details studied in other computationally intensive models [6]. Furthermore, inhomogeneities in refractory period can be implemented with small adjustments to the model. Future work could consider testing the proposed search mechanisms in this adjusted model.

(b) In our implementation of the CMP model, we worked in the low-noise limit at $\delta = 0$ where new re-entrant circuits could not form after the tissue had been initialized with one or two artificial circuits in the tissue. This was done for computational simplicity and is a small source of noise compared to the effects of stochastic coupling on wavefront dynamics. The exception is that non-zero δ does allow for small local fluctuations on short timescales as shown by figure 1d. In a simple re-entrant circuit-searching algorithm, these fluctuations can occasionally be mistaken for circuits. However, these fluctuations dissipate after a single cycle of electrical activity. Therefore, a full implementation of the search algorithm could easily account for these fluctuations by recording electrical dynamics over two or more periods. This implementation of the CMP model only analysed tissues with simple, two fibre re-entrant circuits typically formed in the real CMP model from a single dysfunctional cell. In the full CMP model, re-entrant circuits can form consisting of multiple fibres and multiple dysfunctional cells. On a local level, this can lead to different electrical dynamics. However, the global electrical dynamics are very similar across the different types of re-entrant circuit. It is also rare for complicated re-entrant circuits to form. Therefore, the simple circuit implementation represents the typical case that would be relevant in a realistic setting. Note also that predictions at the local level should be possible regardless of circuit morphology because we expect all circuits to a be a source of electrical wavefronts.

(c) Finally, there are some limitations that arise during data analysis of the electrograms. Electrograms generated from the CMP model are typically cleaner than real atrial electrograms where consistent measurements across the tissue are not always possible. Despite this, the most important features such as the direction of electrical propagation across a set of electrodes can still be found accurately using clinical electrograms [37–40]. By taking the gradient of feature values across the electrode probes in clinical practice, much of the same statistical information that has been used in this work can be calculated and analysed. Note that the local zig-zagging of wavefronts due to the cellular automata can result in errors when taking the gradient of feature values across the probe. However, by taking the mean value of multiple gradients with the same unit vector, these errors are minimized. Any remaining errors are mitigated through the inclusion of a constraint mechanism. Although there could be concerns that a discretized wavefront inhibits the convergence of our algorithm, our high success rate suggests that this is not a major factor.

As electrograms were analysed using supervised machine learning methods, it is also clear that the amount of accurate clinical data required to train models will be severely restricted. In the short term, this might appear to be a major limitation. However, other medical studies have recently shown the efficacy of using simulations to pre-train machine learning models for clinical use before refining the models using clinical data. This process of 'pre-training' before the models are slowly improved through the

acquisition of real data has had recent success applied to gene splicing by Rosenberg *et al.* [41]. It is also important to consider whether any artefacts in the dynamics of the CMP model may have affected the success of our algorithm in a way that could not be applied in a more realistic setting. Of particular note is the perfect isotropy along muscle fibres which may lead to wavefronts travelling in the longitudinal direction being unrealistically uniform. Despite this concern, this particular artefact is more likely to hinder the success of machine learning algorithms than be a benefit since wavefront uniformity can limit the model's capacity to distinguish between different regions of the CMP tissue. On the whole, this should be seen as a secondary concern when adapting this work to more realistic settings.

As a final note, it is interesting to highlight surprising parallels between the CMP model and recent developments in the mechanistic understanding of AF. Hansen *et al.* [12] have recently demonstrated that micro-anatomical re-entrant circuits forming in the transmural region with variable orientation, between the epicardium and the endocardium, can result in a complex variety of breakthrough patterns observed at the surface of the epicardium. The surface of the epicardium is typically what is mapped during ablation and for most electrophysiological studies of the atria. These breakthrough patterns can appear as full rotors, partial rotors or concentric activity—a direct explanation of observations that were previously seen as in direct conflict but which can now be explained from a single mechanism. Nattel *et al.* [3] has suggested this may act as a unifying mechanism for the understanding of AF. Interestingly, the work by Hansen *et al.* specifically associates the emergence of AF with the formation of micro-anatomical re-entrant circuits at the edges of regions with high fibrosis, in direct agreement with the mechanism described by the CMP model. Although the two-dimensional formulation of the CMP model is currently unable to investigate the formation of varying breakthrough patterns at the epicardial surface, a three-dimensional implementation should be able to reproduce the key features of these observations.

6. Conclusion

Despite major research efforts focusing on the theoretical background and clinical understanding of AF in recent years, ablation success rates have remained disappointing since the 1990s. Major improvements will only be possible given significant developments in our mechanistic understanding of AF and new technological approaches to ablation. The methods demonstrated here apply the benefits of machine learning and couple it with the capacity for large-scale simulations of cardiac electrical activity in cellular automata to locate re-entrant circuit for targeted ablation. The work presented here is a first step in a simplified theoretical model that demonstrates a clear potential for locating re-entrant circuits with high success from electrograms alone. It is clear that the current work ignores many of the details which will be crucial when applying these methods in a more realistic setting, particularly with regard to the lack of topological complexity in the CMP model and the inability to reproduce realistic wavefront curvature— these issues should be considered for any future work. Given the efficiency of the model, extending the two-dimensional CMP tissue to a three-dimensional structure should be computationally easy and investigated as a priority.

Appendix A. Electrogram simulation

For the CMP model, electrograms can be simulated using the following equation:

$$V(X', Y') = \sum_{X,Y} \frac{\Delta X(\nabla_x V) + \Delta Y(\nabla_y V)}{(\Delta X^2 + \Delta Y^2 + \Delta Z^2)^{3/2}}, \tag{A 1}$$
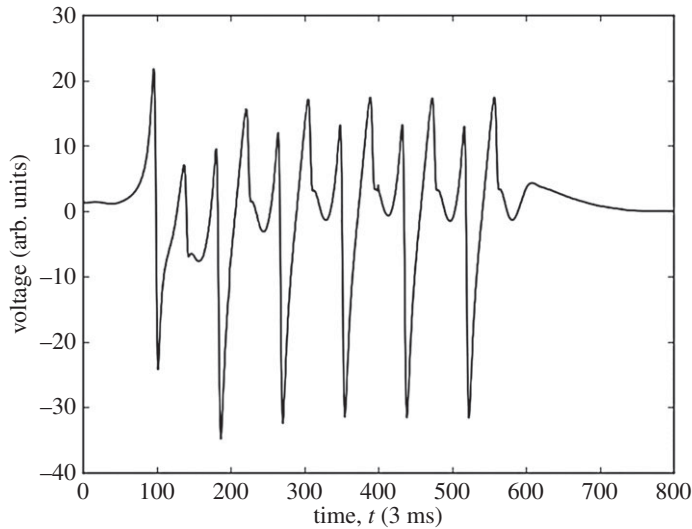
Figure 8. An example of an electrogram showing a short period of fibrillation driven by a re-entrant circuit. The electrogram has been recorded on the point of re-entry. More details can be found in [15].

where (X, Y) are the cell positions on the CMP tissue, $\nabla_x V = V(X, Y) - V(X - 1, Y)$ and $\nabla_y V = V(X, Y) - V(X, Y - 1)$ are the discretized gradients, (X', Y') is the position of the probe, $\Delta X = X - X'$, $\Delta Y = Y - Y'$ and ΔZ is the distance between the probe and tissue surface. $V(X, Y)$ is set to be the state of the cell, where a resting cell has a state value of 0 and an excited cell has a state value of 50 in arbitrary units. For a refractory cell, the state value is a linear interpolation between these two extremities. This action potential was chosen as it simplifies the process but it can be mapped back to a more realistic representation—doing so has a negligible effect on analysis.

Examples of electrograms in perfect isotropic tissue with fibrillation generated by a point source are shown in figure 3 of the main text. An example of an electrogram generated at the point of re-entry is shown in figure 8.

Appendix B. Electrogram visualization and analysis

The features generated from the electrogram data are listed in table 2. For analysis, the gradient of these features is calculated across the multi-electrode grid described in figure 4. The relative time at which the sampling of raw electrode data starts can be exploited to calculate an approximate direction of the wavefront across the electrode grid.

Feature visualization is an exceedingly useful tool in enhancing the understanding of our system, informing the design of machine learning models and at various levels of code verification.

There are a vast array of standard feature visualizations available in machine learning libraries. However, these do not account for the spatial dependence of our electrograms. Therefore, we have principally relied on a custom visualization which we have coined the 'VFM'.

In the simplified CMP model with a single artificial re-entrant circuit, each simulated heart tissue has the AF driver positioned at a different random position. Therefore, an electrogram recorded at the tissue centre is likely to show different behaviour in different heart instances. However, at a fixed vector away from the AF driver we would expect the general behaviour of features to be largely consistent independent of where in the heart tissue the driver is located. By simulating thousands of different heart instances, we can map the behaviour of features to a single image where each coordinate corresponds to the feature behaviour at that given vector away from the driver. This process is visualized in figure 9. The value shown in a VFM is the average over many instances. As such, these behaviours are not necessarily seen in individual tissue instances—the visualizations only indicate general trends.

The general regions on a VFM showing distinct flow dynamics are shown in figure 10. The key conclusions from VFM analysis are the following:

— There is poor feature separation in the bulk. Although there are some indicators differentiating between regions of unidirectional flow and bidirectional flow, this transition is not sharp across tissue instances.

Table 2. The features used for electrogram analysis. For an in-depth description of individual functions see the `scipy` and `numpy` documentation.[a] The features were primarily chosen for ease of use, descriptive capacity and speed at which they could be generated for large quantities of data. Before being processed, the raw electrogram data were sampled to ensure consistent preprocessing. This was done by cropping the raw data between the first two maxima of the dominant Fourier frequency. This ensured the output sample was a consistent length (one driver cycle period) and the sampling was started at the same point in the cycle. The time at which the crop is started relative to the initial recording time corresponds to feature 24 in the table. Note that as the initial recording time is arbitrary, this cannot be used as a feature for an individual electrode probe, but the gradient of the initial crop time across a multi-electrode probe can be measured. The sample data are represented by the label 'X' in the table. The gradient of the sample data is denoted by 'g(X)'. Functions described by 'f# - f#*' correspond to an operation involving other features listed in the table. For the Fourier features labelled with (*), the largest 9 frequencies and their corresponding amplitudes are calculated.

	feature name	scipy function
1.	maximum value	numpy.max(X)
2.	minimum value	numpy.min(X)
3.	amplitude	f1 - f2
4.	intensity	numpy.sum(numpy.absolute(X))
5.	maximum gradient	numpy.max(g(X))
6.	minimum gradient	numpy.min(g(X))
7.	amplitude gradient	f6 - f5
8.	max. gradient time	numpy.argmax(g(X))
9.	min. gradient time	numpy.argmin(g(X))
10.	amplitude gradient time	f9 - f8
11.	number of 0 V crossovers	numpy.argwhere (g(X)[i] * g(X)[i+1] < 0)
12.	first 0 V crossover time	numpy.min(f11)
13.	largest Fourier frequencies*	numpy.fft.rfftfreq(X)
14.	largest Fourier amplitudes*	numpy.absolute(numpy.fft.rfft(X))
15.	Fourier sum	numpy.sum(f14)[:10]
16.	relative Fourier amplitudes*	f14[i] / f15
17.	mean	scipy.stats.describe(X)[2]
18.	skewness	scipy.stats.describe(X)[4]
19.	kurtosis	scipy.stats.describe(X)[5]
20.	maximum time	numpy.argmax(X)
21.	minimum time	numpy.argmin(X)
22.	amplitude time	f20 - f21
23.	variance post minimum	numpy.std(X[f21:])
24.	sample start index	(see caption)

[a]Documentation for both `scipy` and `numpy` can be found at https://docs.scipy.org/doc/.

— There is strong feature separation between both the bulk and the ectopic beat's X-axis, and between the bulk and the axis along which wavefronts collide. This means these regions will be most susceptible to detection via statistical methods.

— Single-electrode features are symmetric across both the X- and Y-axes of the driver. However, this symmetry can be broken by gradient features from the multi-electrode grid. The sharp transition in Y flow direction when crossing the X-axis is visible in these features and can be used to constrain the search region for driver detection.

— Even with gradient-based features, regions in the bulk close to the driver are hard to distinguish from the bulk far from the driver. Hence, finding the driver from a single set of electrogram recordings will be difficult.

— The stochastic formation of vertical couplings in the CMP model means that the transition in X flow direction when crossing the Y-axis of the driver is only sharp at the driver (as opposed to other positions along the Y-axis which do not coincide with the driver).
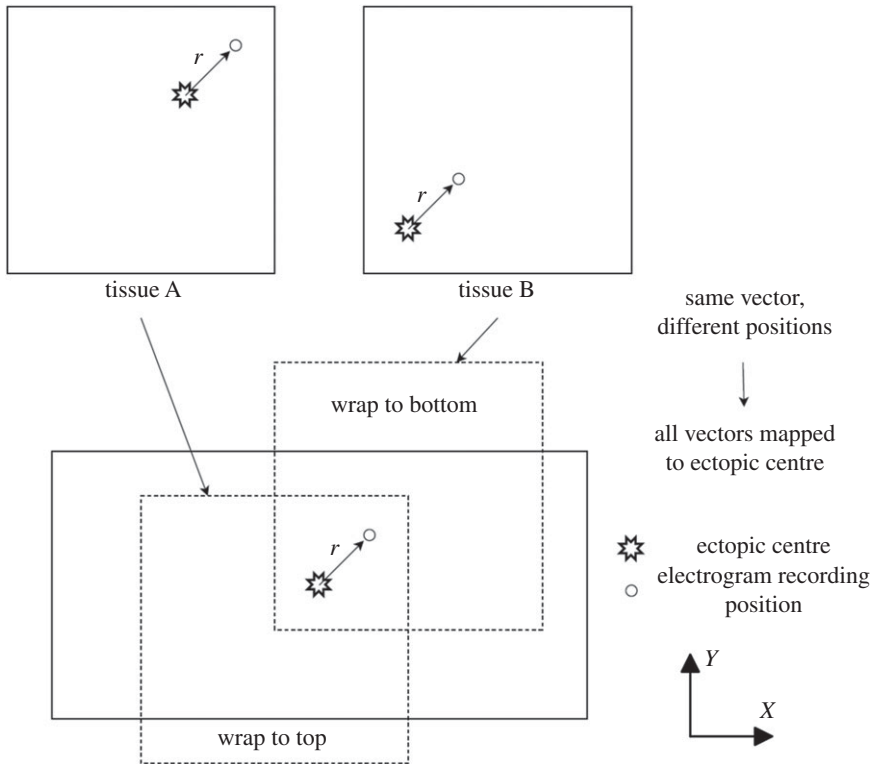
Figure 9. Consider tissues A & B, where the drivers are located at coordinate (150, 150) and (50, 50), and electrograms are recorded at (170, 170) and (70, 70), respectively. Despite the electrograms being at different coordinates in the heart tissue, the relative vector from the driver to the electrogram is the same in both tissues, $r = (20, 20)$. Hence, both electrograms typically record very similar wavefront behaviour. Note that not every tissue has electrograms recorded at every vector relative to the driver. The periodic boundary conditions of the CMP model in the Y (transverse) direction mean that the relative vector from driver to recording position is always taken as the shortest distance between the two points—this is expressed by the individual tissues being wrapped. This method used to generate vector feature maps. Thousands of instances (tissues) of the CMP model are generated with a single re-entrant circuit randomly positioned in each tissue—two such tissues, A & B, are shown in this image. At a fixed vector away from the driver, **r**, we expect largely similar wavefront dynamics, and hence largely consistent electrogram features. These vectors can be mapped onto a single image by centring each tissue on an arbitrarily chosen fixed point of the tissue's re-entrant circuit. Each coordinate in the vector feature map then corresponds to the average value of a feature recorded at that particular vector away from the re-entrant circuit in thousands of different heart instances.

— The spheres of influences between neighbouring re-entrant circuits are largely self-contained—moving between the first circuit and the second, feature behaviour is almost perfectly approximated by the feature behaviour from the nearest circuit with the exception of a very small transition region where the wavefronts from the two circuits collide. This is shown in figure 11.

Our aim is not to find the driver in a single step but rather collate a small number of measurements to find the driver to a high degree of accuracy. Bearing that objective in mind, the considerations above suggest the following approach for designing a driver locating algorithm for the simplified CMP model.
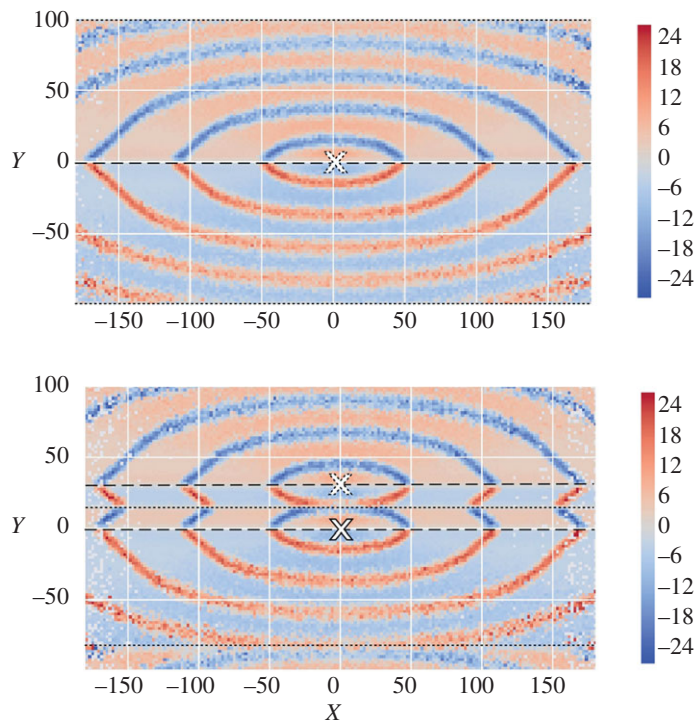
Appendix C. Random forests

The random forest model is built using a set of decision trees. Each decision tree is created via the use of labelled example responses with associated electrogram features. For a general tree, the tree's response space is split into J non-overlapping regions distributed according to a certain metric. For a quantitative tree, the metric is the residual sum of squares given by
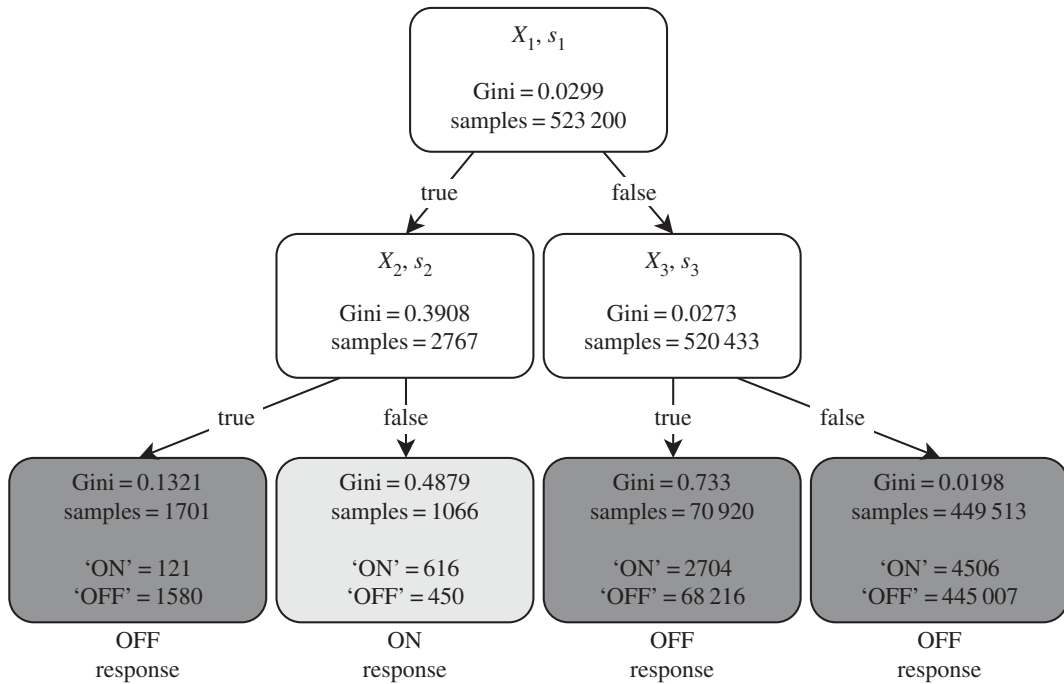
$$\sum_{j=1}^{J}\sum_{i \in R_j}(Y_i - \hat{Y}_{R_j})^2,$$ (C 1)

Figure 12. Simple qualitative decision tree with a maximum depth of two. The responses are whether or not the probe is on or off the AF driver. Three features (X_1, X_2 and X_3) with their corresponding thresholds (s_1, s_2 and s_3) are used to create the response regions.

where Y_i are the individual response values over all response regions R_j, and $\hat{Y}_{R_j}$ is the average response in the region R_j, given by

$$\hat{Y}_{R_j} = \frac{1}{|R_j|} \sum_{R_j} Y_{R_j}. \tag{C 2}$$

For a qualitative tree, the metric is instead the Gini index given by

$$G = \sum_{k=1}^{K} \hat{p}_{jk}(1 - \hat{p}_{jk}), \tag{C 3}$$

where K is the number of possible responses and $\hat{p}_{jk}$ is the fraction of responses k in region j.

The J response regions are distributed by choosing an electrogram feature X_n with threshold value s, which splits a region into the complementary regions $R_1(n, s) = \{X | X_n \geq s\}$ and $R_2(n, s) = \{X | X_n < s\}$. The feature and threshold are chosen to give the largest reduction in either of the chosen metrics at each split. This is a greedy top-down process as the feature and threshold are chosen without considering future splits to simplify and speed up training (figure 12).

These metrics can also be used to determine relative feature importances. In the case of a classification random forest, a smaller Gini index indicates that the response regions are purer. This means that features which greatly reduce the Gini index can be seen as more important (figure 13). This measure can act as a rough summary of the feature importance and help determine future feature sets.

One of the drawbacks of decision trees is overfitting. When the tree is complex (large J), more features are used which reduces the bias but increases the variance of the response. This makes the decision tree overly specific to the training data, making the model very sensitive to small perturbations in the test data. Random forests get around this through a process known as bagging. This process averages all the individual tree's responses, thus reducing the variance but keeping the bias low—this can be thought of as the mathematical equivalent of the expression 'the wisdom of the crowd'. For qualitative trees, a majority rule is used as the overall response. This effect is amplified by de-corelating the individual decision trees by choosing the best feature from a random subset of the total features. The size of the random subset is $\sqrt{m}$, where m is the total number of features processed from the electrogram waveforms.

The random forests built for this research were generated using the scikit-learn package in python with 15 decision trees [43]. This number of decision trees was sufficient to give good performance for locating the X and Y positions of the re-entrant circuit.
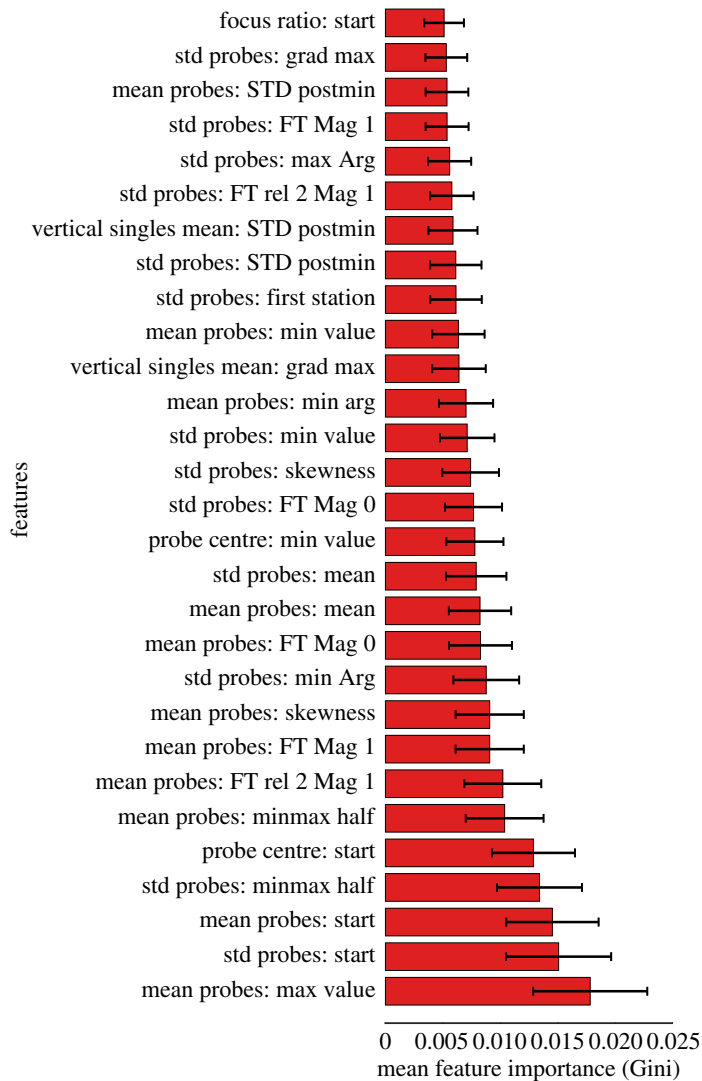
Figure 13. Figure showing the top 29 features by importance of an example random forest model using the Gini index. The prefix refers to a particular gradient/measure across the electrode grid. The suffix refers to a single electrode feature. The feature 'start' can be manipulated to calculate the wavefront direction. As the random forest model is trained, the reduction of the Gini index caused from each feature is measured and averaged across the whole decision tree ensemble. From this particular random forest, the mean maximum voltage and standard deviation of the start index are important features. It can also be seen that from the magnitude of the Gini values, no individual feature is suitable just by itself. These plots help determine suitable feature sets for future models.

All features used in the final random forest classifiers had some predictive capacity according to their Gini index. Of the features described in appendix B, the most discriminatory features are shown in figure 13. Similar quality results could have been achieved using a reduced set of features; however, the selection of features chosen can be generated with high efficiency such that there is no computational downside to using the full set of features.

Appendix D. Constraints and post-processing of predictions

The individual random forest models trained can make independent predictions about the displacement of the current electrogram position from the driver. However, the full driver location algorithm must use these in synergy and restrict their predictions to known constraints. The algorithm is split into two phases, the Y regression stage followed by the X regression stage; it is shown in its full form in figure 6.

As a proof of concept, the constraints imposed here are kept reasonably simple. However, their importance should not be underestimated in ensuring the algorithms predictions converge to the true driver position and avoid infinite cycles of repeated prediction mistakes.
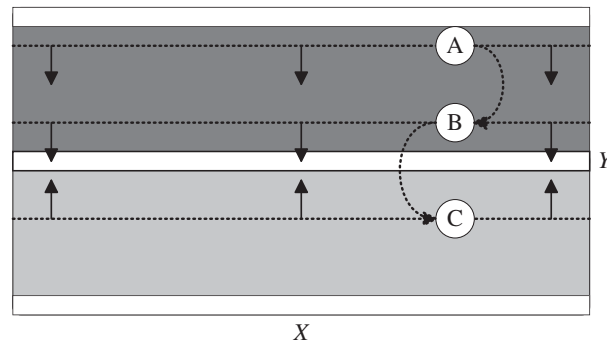
Figure 14. A schematic showing the method used to dynamically constrain the region searched for driver location by cross-referencing flow information against the flow direction maps shown in figure 10. This schematic uses the Y flow map to constrain the Y predictions. An equivalent (but less accurate) flow map can be generated for the X direction. At each point in the CMP model, the net flow across the electrode grid can be calculated. This flow is strongly dependent on position relative to the recording position's closest driver. White regions indicate regions of zero net flow in the Y direction corresponding to the driver axis (central white region) and the axis of wavefront collision (top and bottom white regions). The dark and light grey regions indicate net flow in the $+Y$ and $-Y$ directions, respectively. For each set of electrogram recordings, the flow can be calculated. If $+Y$ flow is detected, the driver must be in the $-Y$ direction from the current position (cases A & B), and vice versa (case C). With each new electrogram recording the search area is restricted towards the driver's X axis.
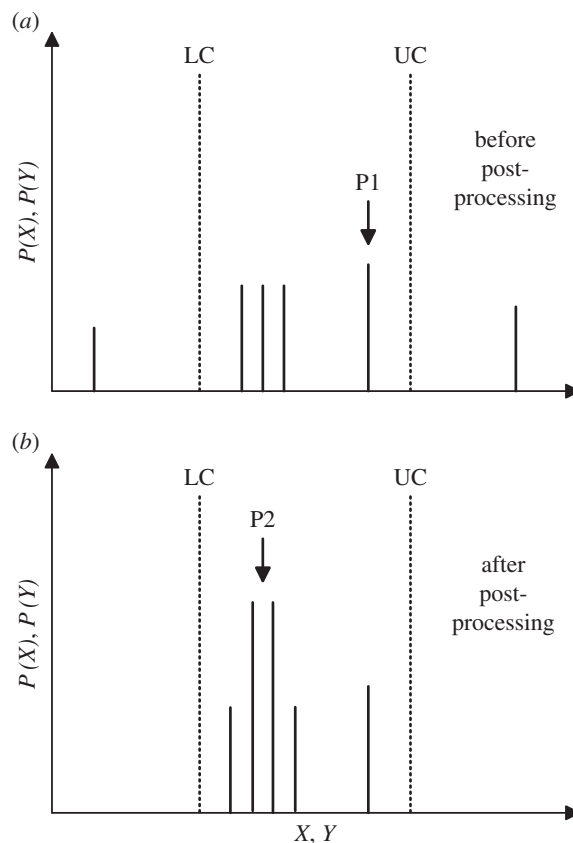
Figure 15. A schematic showing the post-processing of regression model predictions. Models output a probability distribution across the full range of X or Y. The dotted lines represent upper and lower constraints (UC/LC). Probabilities outside the constraints are set to zero. A square wave convolution is then applied which sums the probabilities of predictions over a small range in X or Y. This ensures that large clusters of small probabilities are used for predictions (P2) instead of individual outliers with only marginally larger probabilities (P1). The square wave window is of width 2 and increases in size until a peak is found with $P > 0.5$. If the square wave reaches a width of 8 cells, the process is stopped and the output prediction is taken as the midpoint of the upper and lower constraints. Ignoring constraints, the convolution process improves regression predictions from an average error of 16.2 to 13.3 cells for Y and from 6.4 to 4.7 cells for X.

The first set of constraints use details in the change of local flow direction to dynamically restrict the search area. The mechanism is described in figure 14. The predictions from individual regression models may not lie within the calculated constraints. Therefore, regression models are adapted to output a probability map for likely driver locations. The probabilities outside the constraint region are ignored and a square wave convolution is used to improve the predictions within the constraint window by giving a greater weight to a large cluster of slightly smaller probabilities rather than a single outlying larger probability. This process is illustrated in figure 15.

In a multiple-driver system, after the first driver has been found, the search for the second driver is initiated as described in figure 7. The deviation of the wavefront collision point is used to constrain the search area to the areas marked in grey. For computational ease, our proof of concept insists that the algorithm cannot re-enter the region in which the previous driver was found. This is a small simplification and should not hinder the success of the algorithm in more complicated systems.

References

1. Nattel S. 2002 New ideas about atrial fibrillation 50 years on. Nature 415, 219–226. (doi:10.1038/415219a)
2. Schotten U, Verheule S, Kirchhof P, Goette A. 2011 Pathophysiological mechanisms of atrial fibrillation: a translational appraisal. Physiol. Rev. 91, 265–325. (doi:10.1152/physrev.00031.2009)
3. Nattel S, Xiong F, Aguilar M. 2017 Demystifying rotors and their place in clinical translation of atrial fibrillation mechanisms. Nat. Rev. Cardiol. 14, 509–520. (doi:10.1038/nrcardio.2017.37)
4. Lee S, Sahadevan J, Khrestian CM, Cakulev I, Markowitz A, Waldo AL. 2015 Simultaneous bi-atrial high density (510–12 electrodes) epicardial mapping of persistent and long-standing persistent atrial fibrillation in patients: new insights into the mechanism of its maintenance. Circulation 115, 2108–2117.
5. Trayanova NA, Pashakhanloo F, Wu KC, Halperin HR. 2017 Imaging-based simulations for predicting sudden death and guiding ventricular tachycardia ablation. Circ.: Arrhythmia Electrophysiol. 10, e004743. (doi:10.1161/CIRCEP.117.004743)
6. Clayton R et al. 2011 Models of cardiac tissue electrophysiology: progress, challenges and open questions. Prog. Biophys. Mol. Biol. 104, 22–48. (doi:10.1016/j.pbiomolbio.2010.05.008)
7. Hood L, Heath JR, Phelps ME, Lin B. 2004 Systems biology and new technologies enable predictive and preventative medicine. Science 306, 640–643. (doi:10.1126/science.1104635)
8. Winslow RL, Trayanova N, Geman D, Miller MI. 2012 Computational medicine: translating models to clinical care. Sci. Transl. Med. 4, 158rv11. (doi:10.1126/scitranslmed.3003528)
9. Deo RC. 2015 Machine learning in medicine. Circulation 132, 1920–1930. (doi:10.1161/CIRCULATIONAHA.115.001593)
10. Li Q, Rajagopalan C, Clifford GD. 2014 Ventricular fibrillation and tachycardia classification using a machine learning approach. IEEE Trans. Biomed. Eng. 61, 1607–1613. (doi:10.1109/TBME.2013.2275000)
11. Roney CH et al. 2017 Spatial resolution requirements for accurate identification of drivers of atrial fibrillation. Circ.: Arrhythmia Electrophysiol. 10, e004899. (doi:10.1161/CIRCEP.116.004899)
12. Hansen BJ et al. 2015 Atrial fibrillation driven by micro-anatomic intramural re-entry revealed by simultaneous sub-epicardial and sub-endocardial optical mapping in explanted human hearts. Eur. Heart J. 36, 2390–2401. (doi:10.1093/eurheartj/ehv233)
13. Cherry EM, Evans SJ. 2008 Properties of two human atrial cell models in tissue: restitution, memory, propagation, and reentry. J. Theor. Biol. 254, 674–690. (doi:10.1016/j.jtbi.2008.06.030)
14. Christensen K, Manani KA, Peters NS. 2015 Simple model for identifying critical regions in atrial fibrillation. Phys. Rev. Lett. 114, 028104. (doi:10.1103/PhysRevLett.114.028104)
15. Manani KA, Christensen K, Peters NS. 2016 Myocardial architecture and patient variability in clinical patterns of atrial fibrillation. Phys. Rev. E. 94, 042401. (doi:10.1103/PhysRevE.94.042401)
16. Lin YT, Chang ET, Eatock J, Galla T, Clayton RH. 2017 Mechanisms of stochastic onset and termination of atrial fibrillation studied with a cellular automaton model. J. R. Soc. Interface 14, 20160968. (doi:10.1098/rsif.2016.0968)
17. Lord J et al. 2013 Economic modelling of diagnostic and treatment pathways in National Institute for Health and Care Excellence clinical guidelines: the Modelling Algorithm Pathways in Guidelines (MAPGuide) project. Health technology assessment 17, 1–192.
18. Ciaccio EJ, Biviano AB, Wan EY, Peters NS, Garan H. 2017 Development of an automaton model of rotational activity driving atrial fibrillation. Comput. Biol. Med. 83, 166–181. (doi:10.1016/j.compbiomed.2017.02.008)
19. Wolfram S. 1986 Theory and applications of cellular automata, vol. 1. Singapore: World scientific.
20. Alonso Atienza F, Requena Carrión J, García Alberola A, Rojo Álvarez JL, Sánchez Muñoz JJ, Martínez Sánchez J, Valdés Chávarri M. 2005 A probabilistic model of cardiac electrical activity based on a cellular automata system. Rev. Esp. Cardiol. (English Edition) 58, 41–47. (doi:10.1016/S1885-5857(06)60233-8)
21. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. 2017 Dermatologist-level classification of skin cancer with deep neural networks. Nature 542, 115–118. (doi:10.1038/nature21056)
22. Libbrecht MW, Noble WS. 2015 Machine learning applications in genetics and genomics. Nat. Rev. Genet. 16, 321–332. (doi:10.1038/nrg3920)
23. Caruana R, Niculescu-Mizil A. 2006 An empirical comparison of supervised learning algorithms. In Proc. of the 23rd Int. Conf. on Machine Learning. ICML '06, Pittsburgh, PA, 25–29 June, pp. 161–168. New York, NY, USA: ACM. See http://www.acm.org/10.1145/1143844.1143865.
24. Caruana R, Karampatziakis N, Yessenalina A. 2008 An empirical evaluation of supervised learning in high dimensions. In Proc. of the 25th Int. Conf. on Machine learning, Helsinki, Finland, pp. 96–103. New York, NY: ACM.
25. Zahid S et al. 2016 Patient-derived models link re-entrant driver localization in atrial fibrillation to fibrosis spatial pattern. Cardiovasc. Res. 110, 443–454. (doi:10.1093/cvr/cvw073)
26. Haissaguerre M, Shah AJ, Cochet H, Hocini M, Dubois R, Efimov I, Vigmond E, Bernus O, Trayanova N. 2016 Intermittent drivers anchoring to structural heterogeneities as a major pathophysiological mechanism of human persistent atrial fibrillation. J. Physiol. 594, 2387–2398. (doi:10.1113/JP270617)
27. Matsuo S et al. 2009 Clinical predictors of termination and clinical outcome of catheter ablation for persistent atrial fibrillation. J. Am. Coll. Cardiol. 54, 788–795. (doi:10.1016/j.jacc.2009.01.081)
28. Bub G, Shrier A, Glass L. 2002 Spiral wave generation in heterogeneous excitable media. Phys. Rev. Lett. 88, 058101. (doi:10.1103/PhysRevLett.88.058101)
29. Fast VG, Kléber AG. 1995 Cardiac tissue geometry as a determinant of unidirectional conduction block: assessment of microscopic excitation spread by optical mapping in patterned cell cultures and in a computer model. Cardiovasc. Res. 29, 697–707. (doi:10.1016/S0008-6363(96)88643-3)
30. Ciaccio EJ, Coromilas J, Wit AL, Peters NS, Garan H. 2016 Formation of functional conduction block during the onset of reentrant ventricular tachycardia. Circ.: Arrhythmia Electrophysiol. 9, e004462. (doi:10.1161/CIRCEP.116.004462)
31. Ciaccio EJ, Coromilas J, Wit AL, Peters NS, Garan H. 2018 Source-sink mismatch causing functional conduction block in re-entrant ventricular tachycardia. JACC: Clin. Electrophysiol. 4, 1–16. (doi:10.1016/j.jacep.2017.08.019)
32. Medtronic Inc. 2011 Achieve Technical Manual.
33. Breiman L. 2001 Random forests. Mach. Learn. 45, 5–32. (doi:10.1023/A:1010933404324)
34. James G, Witten D, Hastie T, Tibshirani R. 2013 An introduction to statistical learning, vol. 6. Berlin, Germany: Springer.

35. Pereira S, Pinto A, Oliveira J, Mendrik AM, Correia JH, Silva CA. 2016 Automatic brain tissue segmentation in MR images using random forests and conditional random fields. *J. Neurosci. Methods* **270**, 111–123. (doi:10.1016/j.jneumeth.2016.06.017)

36. Austin PC, Tu JV, Ho JE, Levy D, Lee DS. 2013 Using methods from the data mining and machine learning literature for disease classification and prediction: a case study examining classification of heart failure sub-types. *J. Clin. Epidemiol.* **66**, 398–407. (doi:10.1016/j.jclinepi.2012.11.008)

37. Roney CH *et al.* 2014 An automated algorithm for determining conduction velocity, wavefront direction and origin of focal cardiac arrhythmias using a multipolar catheter. In *36th Annual Int. Conf. IEEE Engineering in Medicine and Biology Society 2014, Chicago, IL, 26–30 August*, pp. 1583–1586. IEEE. (doi:10.1109/EMBC.2014.6943906)

38. Cantwell CD, Roney CH, Ng FS, Siggers JH, Sherwin SJ, Peters NS. 2015 Techniques for automated local activation time annotation and conduction velocity estimation in cardiac mapping. *Comput. Biol. Med.* **65**, 229–242. (doi:10.1016/j.compbiomed.2015.04.027)

39. Lee G *et al.* 2014 Epicardial wave mapping in human long-lasting persistent atrial fibrillation: transient rotational circuits, complex wavefronts, and disorganized activity. *Eur. Heart J.* **35**, 86–97. (doi:10.1093/eurheartj/eht267)

40. Stevenson WG, Soejima K. 2005 Recording techniques for clinical electrophysiology. *J. Cardiovasc. Electrophysiol.* **16**, 1017–1022. (doi:10.1111/j.1540-8167.2005.50155.x)

41. Rosenberg A, Patwardhan R, Shendure J, Seelig G. 2017 Learning the sequence determinants of alternative splicing from millions of random sequences. *Cells* **163**, 698–711. (doi:10.1016/j.cell.2015.09.054)

42. McGillivray MF, Cheng W, Peters NS, Christensen K. 2018 Data from: Machine learning methods for locating re-entrant drivers from electrograms in a model of atrial fibrillation. Dryad Digital Repository. (doi:10.5061/dryad.675260p)

43. Pedregosa F *et al.* 2011 Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830.