

# Ccoffinn: Automated Wave Tracking in Cultured Cardiac Monolayers

Jakub Tomek,<sup>1,\*</sup> Rebecca A. B. Burton,<sup>2</sup> and Gil Bub<sup>3</sup>

<sup>1</sup>Department of Physiology, Anatomy, and Genetics and <sup>2</sup>Department of Pharmacology, University of Oxford, Oxford, United Kingdom; and <sup>3</sup>Department of Physiology, McGill university, Montreal, Quebec, Canada

**ABSTRACT** Cardiac arrhythmias are one of the most frequent causes of death worldwide. A popular biological model used to study arrhythmogenesis is the cultured cardiac cell monolayer, which provides a good trade-off between physiological relevance and experimental access. Excitation wave patterns are imaged using high-bandwidth detectors, producing large data sets that are typically analyzed manually. To make such analysis less time consuming and less subjective, we have designed and implemented a toolkit for segmentation and tracking of cardiac waves in optical mapping recordings. The toolkit is optimized for high-resolution detectors to accommodate the growing availability of inexpensive high-resolution detectors for life science imaging applications (e.g., scientific CMOS cameras). The software extracts key features of propagating waves, such as wavefront speed and entropy. The methods have been validated using synthetic data, and real-world examples are provided, showing a difference in conduction velocity between two different types of cardiac cell cultures.

## INTRODUCTION

Cardiac arrhythmias such as ventricular or atrial fibrillation are a major factor in the occurrence of cardiac arrest, one of the most frequent causes of death worldwide, claiming 300,000–400,000 deaths annually in the United States alone (1). Despite recent advances and decades of research, a precise insight into the mechanics of fibrillation is lacking (2). A popular experimental model of arrhythmogenesis, or, more generally, signal propagation in excitable media, is cultured cardiac monolayers (CCMs) (3). CCMs can be grown in a controlled manner, allowing for manipulation of their spatial and functional organization. For example, CCMs can give insight into the origins of reentrant waves and can be used to determine conditions that allow different wave topologies to occur (4). Despite being structurally different from *in vivo* tissue, CCMs remain popular models of cardiac conduction due to their simplicity, controllability of growth, and ease of experimental access due to very little movement of the tissue and the absence of deep 3D structure. CCMs therefore allow researchers to perform experiments that would not be practical *in vivo*.

Optical mapping techniques are used to observe wave propagation in CCMs. Two types of optical mapping, dye

based and dye free, are used. Typically, cell activity is recorded using fluorescence imaging, using either voltage-sensitive or calcium-sensitive dyes (5). Although these techniques directly measure key biophysical properties, they tend to suffer from two major drawbacks, phototoxicity and photobleaching, rendering longer recording of tissue nearly impossible. An alternative approach is based on dye-free visualization of contraction of cardiac cells, using either phase imaging (6) or off-axis illumination (7). Although dye-free techniques do not capture membrane voltage or calcium concentration, they allow direct visualization of wave patterns without the drawbacks associated with fluorescent-imaging approaches.

All the optical mapping techniques are capable of producing a large quantity of image data, but the analysis of such data is still frequently done manually (8), or largely manually (9), which is subjective and time-consuming. An automated analysis of wave propagation in optically mapped tissue would be less prone to operator bias and has the potential to increase throughput. However, most of the existing automated methods for analyzing optical mapping data focus on the properties of single cells or patches of tissue, rather than macroscopic behavior of the tissue.

There are automated approaches to the analysis of higher-level activity (e.g., waves), but their focus seems to be mainly on the analysis of data from electrode arrays (10,11); an example task being solved is keeping track of the number of waves in a recording, recording the events

---

Submitted February 22, 2016, and accepted for publication August 26, 2016.

\*Correspondence: [jakub.tomek.mff@gmail.com](mailto:jakub.tomek.mff@gmail.com)

Editor: Andrew McCulloch.

<http://dx.doi.org/10.1016/j.bpj.2016.08.049>

© 2016 Biophysical Society.



of splitting and fusion (12,13). Such electrophysiological recordings are characterized by an excellent temporal resolution of  $>1$  kHz, a low amount of noise, and relatively poor spatial resolution (e.g., the  $22 \times 23$  and  $24 \times 21$  arrays used in (10)). These techniques have been successfully applied to voltage-sensitive dye recordings collected by specialized low-resolution high-speed detectors (14).

Recently, relatively inexpensive high-resolution, low-noise sCMOS sensors with moderate frame rates have been made widely available for biomedical research (15). These detectors yield data sets with relatively poor temporal resolution of up to 100 Hz, and signals typically have a lower signal/noise ratio (SNR), but their high spatial resolution offers several advantages. Such an increase in spatial resolution may be used to measure detailed properties of wave propagation, such as the distribution of wavefront speeds for different parts of a wave, or to track small and short-lived wavelets (16), which cannot be achieved with lower-resolution devices. The gold standard for measuring conduction properties is the method by Bayly et al. (10), developed for conduction-velocity estimation in electrophysiological recordings, also used in the toolkit Rhythm (14). However, in our experience, the method is not ideally suited for processing high-resolution, heterogeneously conducting tissue; a discussion of its shortcomings and a comparison with the methods developed by us can be found in Section S7 in the Supporting Material.

Because of the lack of suitable software for analysis of the behavior of waves in high-spatial-resolution data sets, particularly for recordings where the activation is not a continuous pattern (such as obtained using conventional calcium imaging) but a discrete one, we designed the toolkit Ccoffinn (Cardiomyocyte Cultures Optically-mapped: Fast Feature extractIoN and trackiNg), presented in this text; the code (with user guide and sample data) is available from <https://ccoffinn.dpag.ox.ac.uk/>. The approach is based

on the segmentation of single cells or groups of neighboring cells, using the information of their activity to construct a representation of waves and track their movement. After these steps, features are extracted that describe the quantitative properties of the observed waves, such as wavefront speed, wavefront smoothness (or lack thereof), or several measures of order in the recorded activity. The presented methods are validated using synthetic data and used as a proof of concept to demonstrate the difference in conduction properties of myocyte cultures and cocultures of myocytes and cardiac neurons. The whole process of segmentation, tracking, and feature extraction, is implemented mainly in Matlab, and it is relatively high throughput, allowing for processing hundreds of data sets per day on a personal computer. Although the software was designed for data sets obtained using a dye-free imaging approach, we demonstrate that it also can be directly applied to data sets collected using fluorescence techniques with a continuous pattern of activation.

## MATERIALS AND METHODS

This section briefly sketches the functionality of Ccoffinn, describing the process of wave segmentation, tracking, and subsequent feature extraction.

### Wave segmentation and tracking

The main aim of the wave segmentation is to estimate which parts of tissue are firing in a synchronized manner in each frame of a source video. In addition to simple wave segmentation, we also propose a method for detection of a wavefront (the part of a wave that advances forward). Using the methods for segmenting waves, we can then track movement of wavefronts and extract features describing the waves. An example output of Ccoffinn for both dye-free and calcium imaging is shown in Fig. 1. Below is a more detailed description of the algorithms used to get from a source video to extracted features. Furthermore, Sections S3–S5 in the Supporting Material give an overview of parameters, a description of the synthetic and experimental data used, and notes on implementation, respectively.




FIGURE 1 An example output of wave segmentation and tracking. Each row contains two consecutive frames of a recording, with segmentation and annotation by Ccoffinn in the third column. In the first row is a straight wave imaged using a calcium dye at 200 fps. The second row is a case of dye-free imaged wavelets at 50 fps with spurious activity (which is the most difficult type of activity to process). The colors in the images encode the roles of blobs within the waves. Blue represents the area inside the waves, cyan the borders of waves that are not in the upstroke phase, magenta the wavefront blobs (border blobs in the upstroke phase), red the wavefront blobs that were not wavefront in the previous frame (“became-wavefront”), and yellow the tracking arrows linking border blobs from the first shown frame to those that “became wavefront” in the second shown frame.

## Blob segmentation

The aim of this step is to estimate locations of pixel clusters that are activated at the same time (given available temporal resolution). This can be single cells or clusters of cells, depending on the spatial resolution used. Using the off-axis-illumination dye-free imaging, active cells or groups of cells manifest as bright blobs; thus, the segmentation of the cell culture is reduced to blob detection. The segmentation is greatly simplified, as the blobs in monolayers do not change their location throughout the recording. Therefore, we can average all the frames of the video, obtaining an image with an excellent SNR that contains all the blobs that were active. The averaged image is then segmented using the four-connected version of the SeNeCA algorithm (17), chosen for its computational simplicity and its ability to cope with data with uneven background brightness, which is a common feature in our data sets. The parametrization of SeNeCA is described in Section S3 in the Supporting Material.

For cases of data where the imaged tissue does not manifest clear blobs (e.g., macroscopic calcium imaging), an alternative segmentation method based on filtered binning is included in Ccoffinn. Using this approach, a frame of the video is binned into squares with a chosen side length (parameter *binningSize*; see Section S3 in the Supporting Material). Using the whole source video, the SNR is measured for each of the brightness traces of the squares using the Matlab *snr* function. Optionally, the squares with SNR smaller than a given threshold (parameter *minSNR*) are discarded. The remaining squares are then considered to be segmented blobs in a manner similar to segmentation using the default approach.

## Determining times of blob activity

Once the blobs have been segmented, the next task is to determine which blobs are active at which frame of the recording. Simple thresholding of intensities performed poorly on our data, given the wide range of brightness and/or SNR of blob traces. We did not want to use shape-specific methods (such as template matching), as our aim was to keep the method flexible with respect to different imaging modalities and species-specific action-potential morphologies. Methods based on the derivative of the measured intensity, which is often used as a surrogate for the rate of change of membrane voltage (14), cannot be used solely, because the derivative of intensity only provides the time of the upstroke and not information such as action-potential duration (APD). We designed a simple spike-mining algorithm that uses the derivative of intensity to detect spikes as events and then uses temporal surroundings of the detected event to extract the start and end of the particular spike. The detailed description of the algorithm is given in Section S1.1 in the Supporting Material.

## Wave segmentation

Here, we aim to detect waves in each frame. Informally, a wave is defined as a “sufficiently large group of active blobs which are sufficiently dense.” Our approach works frame by frame, and in each frame, it takes all the blobs active in that particular frame, considers them in areas of dense-enough activation (see Section S1.2 in the Supporting Material for details), and divides them into groups using a recursive criterion: blobs *a* and *b* belong to the same wave  $\Leftrightarrow$  *a* and *b* are close enough or if there exist blobs *c* and *d* so that *a* is close enough to *c*, *b* is close enough to *d*, and *c* and *d* belong to the same wave (the case when  $a = c$  or  $b = d$  is allowed).

The blobs are then further segmented according to their “roles”: blobs within a wave, blobs at the border of the wave, and blobs that are at the border and are furthermore wavefront blobs (i.e., the wave is propagating across them, as opposed to, e.g., border blobs that are at the tail of the wave and are ceasing to be active rather than becoming active). The whole process of wave segmentation is described in detail and illustrated in Section S1.2 in the Supporting Material.

## Wavefront tracking

After the blobs are segmented into various types according to their position in waves, the next step is to track the movement of the wavefronts

throughout the recording. The tracking information may be then used to estimate wavefront-related features of interest. Unlike in traditional tracking of blobs, the blobs are stationary in our data and only the pattern of their activation is tracked. This means that the number of tracked objects (wavefront blobs) is highly variable throughout the recording and waves tend to disappear completely unless they are reentrant.

For wavefront tracking, Ccoffinn uses a novel algorithm, to our knowledge, based on solving a minimum-cost bipartite matching problem to link wavefront blobs of a wave with the border blobs of the same wave in the previous frame. This approach is optimal in the total distance between the linked cells and is sufficiently fast for recordings with <10,000 segmented blobs. The detailed explanation of the algorithm is given in Section S1.3 in the Supporting Material. Alternatively, we have implemented the method of Bayly et al. (10), which may be used for smoothly conducting recordings.

## Feature extraction

Rather than to observe wave segmentation and tracking in a video visually, it may be more convenient to extract several features that describe the data set using an automated method. Furthermore, an automated approach avoids operator bias, which can occur if researchers are free to select regions manually, possibly causing false-positive results (18). Ultimately, an automated analysis is more easily scaled to high-throughput applications (e.g., drug screening). For the purpose of the automated description of the recorded behavior, we have designed several features that either simply describe the data or may serve as a way of showing a difference between two groups of data. Most features are actually multisets, which can be further summarized, e.g., averaged over time (producing a spatial map) or over space (producing a development over time), or both (producing a single number). The features extracted are given below, with precise definitions in Section S2 in the Supporting Material.

- Conduction-oriented: wavefront speed (conduction velocity), wavefront roughness
- Organization-oriented: entropy, perimeter/area ratio
- Spike properties: APD, interspike duration, beating frequency
- Structural: number of blobs, sparseness of blobs (spatial)

## RESULTS

We made two comparisons using the available dye-free data (more details on the data are provided in Section S4.2 in the Supporting Material). First, we manually divided the recordings into organized waves (target, spiral, or straight) and disorganized waves (wavelets or chaotic activity). We then compared the features extracted from organized waves to those from disorganized waves. Second, we compared the properties of organized waves in myocyte cultures to organized waves in cocultures of myocytes and neurons. In both comparisons, we used the rank sum test to compare features of the groups of recordings. Furthermore, an evaluation of Ccoffinn using synthetic data is given in Section S6 in the Supporting Material.

The results of comparing the recordings of organized and disorganized activity are given in the upper part of Table 1. The disorganized cultures had a significantly higher perimeter/area ratio and wavefront roughness. This is an expected result, as these two features are mainly aimed at orderliness of propagation in tissue. The third feature aimed at orderliness, entropy, failed to discern between the two classes of data, supporting the opinion, based on the evaluation of

**TABLE 1** Comparisons of Recordings of Organized and Disorganized Activity and of Organized Activity in Myocyte Cultures and Myocyte-Neuron Cocultures

| Organized versus Disorganized Activity                                  |                       |                       |                |
|---|-----------------------|-----------------------|----------------|
| Feature   | Org Median (IQ range) | Dis Median (IQ range) | <i>p</i> value |
| Number of cells   | 4678 (4138–5087)      | 4890 (4318–5573)      | 0.1909         |
| Wavefront speed (mm/s) <sup>a</sup>                                     | 11.3 (7.6–20.8)       | 5.9 (5.2–7.4)         | 0.0024         |
| Interspike period (ms)  | 443.2 (331.8–564.3)   | 372.1 (341.9–466.1)   | 0.2673         |
| APD (ms)  | 138.5 (112.1–180.4)   | 138.2 (103.6–150.3)   | 0.3286         |
| Cell sparseness (m)   | 170.2 (152.8–204.6)   | 192.4 (168.5–210.7)   | 0.3453         |
| Wavefront roughness (°) <sup>a</sup>                                    | 27.1 (23.3–30.9)      | 37.7 (32.7–38.9)      | 0.0010         |
| Beat frequency (Hz)   | 0.838 (0.606–1.403)   | 1.173 (0.894–1.437)   | 0.1694         |
| Entropy   | 0.322 (0.034–2.130)   | 0.127 (0.060–0.509)   | 0.7281         |
| Perimeter/area ratio <sup>a</sup>                                       | 0.139 (0.093–0.162)   | 0.188 (0.161–0.229)   | 0.0022         |
| Organized Activity in Myocyte Cultures versus Myocyte-Neuron Cocultures |                       |                       |                |
| Feature   | Myo Median (IQ range) | Co Median (IQ range)  | <i>p</i> value |
| Number of cells   | 4956 (4751–5313)      | 4411 (4127–4835)      | 0.1471         |
| Wavefront speed (mm/s) <sup>a</sup>                                     | 6.5 (5.3–9.5)         | 17.8 (11.1–21.5)      | 0.0110         |
| Interspike period (ms) <sup>a</sup>                                     | 564.3 (432.9–643.1)   | 366.3 (307.7–464.6)   | 0.0420         |
| APD (ms)  | 180.4 (134.4–204.5)   | 133.5 (107.3–151.3)   | 0.0559         |
| Cell sparseness (m)   | 184.3 (159.5–204.8)   | 169.3 (146.1–204.4)   | 0.7925         |
| Wavefront roughness (°)   | 29.6 (24.1–38.6)      | 26.6 (23.1–29.3)      | 0.2635         |
| Beat frequency (Hz)   | 0.599 (0.516–0.929)   | 0.947 (0.815–1.891)   | 0.0559         |
| Entropy   | 0.667 (0.306–2.081)   | 0.168 (0.003–2.179)   | 0.3676         |
| Perimeter/area ratio  | 0.152 (0.135–0.192)   | 0.117 (0.084–0.162)   | 0.2198         |

For each feature (or mean of feature), the median and interquartile range are given. Org, organized recordings (target, spiral, or straight wave); Dis, disorganized activity (wavelets); IQ, interquartile; Myo, cultures of myocytes; Co, cocultures of myocytes and neurons.

<sup>a</sup>Features that differ significantly ( $p < \alpha = 0.05$ ) between the respective groups.

synthetic data, that it is not a useful feature. Wavefront speed was significantly increased in the group containing organized activity, which is an expected result as well.

The results of comparing recordings of organized activity in myocyte cultures and cocultures of myocytes and neurons are given in the lower part of [Table 1](#). The most striking result is the large increase of wavefront speed in cocultures compared to myocyte cultures. At the same time, the perimeter/area ratio is not significantly different between the two groups of data; thus, the difference in wavefront speed and wavefront roughness is unlikely to be caused by different morphology of waves. Also, we can observe a borderline-significantly faster beat rate in cocultures (this trend is further bolstered by shorter APD and interspike period, consistent with more frequent spiking).

To link the results using real data to the analysis of synthetic data ([Section S6](#) in the [Supporting Material](#)), we have estimated the noise standard deviation in the real data using the method by Immerkær (19). The mean over recordings of standard deviation of the noise was 0.1063 with a standard deviation of 0.0583, suggesting that the results presented in this section should not be affected too severely by the presence of noise.

## DISCUSSION

In this work, we present the toolkit Ccoffinn, a novel approach, to our knowledge, to segmentation and wave

tracking in optically mapped tissue, with subsequent feature extraction, focused on the analysis of high-resolution recordings with heterogeneously conducting tissue. The method was validated using synthetic data and applied to a large cardiac culture data set, where it successfully detected significant differences in properties of two types of cell cultures.

Our software toolkit has the following advantages: the ability to extract many different features from data sets, describing various properties of the imaged tissue that, to our knowledge, have not been analyzed automatically before (e.g., wavefront roughness); computational efficiency; and freely available code. The code is written in Matlab in a modular fashion using static classes and subprograms to allow users to easily incorporate different methods in Ccoffinn while still being able to use the toolkit's core functionality (e.g., the current spike-mining algorithm can be seamlessly replaced without disrupting Ccoffinn's tracking, visualization, and feature-extraction functionality).

The main limitation of Ccoffinn is the offline nature of the software; this is very convenient for blob segmentation tasks, but having a real-time implementation might be also very useful. Implementing a real-time version of the toolkit is feasible, provided that no complex visualization of the segmentation is needed. The main change needed would be an alternative algorithm for determining which cells are active and which are not. A frame-by-frame use of the SeNeCA algorithm might be an option, as this has been validated in (17).



We believe that the main future use of the toolkit will be to discover and quantify the effects of pharmacological agents on cardiac propagation, or for determining functional differences between tissue types, similar to the proof-of-concept application presented in this text. A second use might be to explore the relationships between features in a very large data set, for example, the homogeneity of distribution of cells and the propagation velocity. When the real-time version of our approach is implemented, a third possible use is for optogenetic applications (20), allowing the automated control of illumination for feedback applications that depend on wavefront location or velocity. Ccoffinn could greatly improve the reaction time to experimental events and could allow the application of patterns of stimulation that are well beyond the capabilities of human operators.

## SUPPORTING MATERIAL

Supporting Materials and Methods, Supporting Results, 11 figures, and four tables are available at [http://www.biophysj.org/biophysj/supplemental/S0006-3495\(16\)30816-5](http://www.biophysj.org/biophysj/supplemental/S0006-3495(16)30816-5).

## AUTHOR CONTRIBUTIONS

J.T. designed the new algorithms used, wrote and assembled the code, and performed the analyses using real and synthetic data. R.A.B.B. produced real dye-free data used in the article and provided consultancy and information on the data to J.T. G.B. supervised the project, wrote the initial design document (which features to extract, etc.), and consulted the algorithms used with J.T.

## ACKNOWLEDGMENTS

We thank Dr. Emilia Entcheva for providing us with her data on calcium imaging. J.T. thanks the Bakala Foundation and Engineering and Physical Sciences Research Council (EPSRC) for their support in his DPhil programme. G.B. acknowledges support from the BHF Centre of Research Excellence, Oxford (RE/08/004). R.A.B.B. holds an EPSRC Developing Leaders Grant, a Goodger award, a Winston Churchill Fellowship and a Paul Nurse Junior Research Fellowship (Linacre College, Oxford)

## REFERENCES

1. Zipes, D. P., and H. J. J. Wellens. 1998. Sudden cardiac death. *Circulation*. 98:2334–2351.

2. Israel, C. W. 2014. Mechanisms of sudden cardiac death. *Indian Heart J.* 66 (Suppl 1):S10–S17.
3. Tung, L., and Y. Zhang. 2006. Optical imaging of arrhythmias in tissue culture. *J. Electrocardiol.* 39 (4, Suppl):S2–S6.
4. Bub, G., K. Tateno, ..., L. Glass. 2003. Spontaneous initiation and termination of complex rhythms in cardiac cell culture. *J. Cardiovasc. Electrophysiol.* 14 (10, Suppl):S229–S236.
5. Herron, T. J., P. Lee, and J. Jalife. 2012. Optical imaging of voltage and calcium in cardiac cells and tissues. *Circ. Res.* 110:609–623.
6. Hwang, S. M., K. H. Yea, and K. J. Lee. 2004. Regular and alternant spiral waves of contractile motion on rat ventricle cell cultures. *Phys. Rev. Lett.* 92:198103.
7. Burton, R. A. B., A. Klimas, ..., G. Bub. 2015. Optical control of excitation waves in cardiac tissue. *Nat. Photonics.* 9:813–816.
8. Linnenbank, A. C., J. M. T. de Bakker, and R. Coronel. 2014. How to measure propagation velocity in cardiac tissue: a simulation study. *Front. Physiol.* 5:267.
9. Doshi, A. N., R. D. Walton, ..., R. Coronel. 2015. Feasibility of a semi-automated method for cardiac conduction velocity analysis of high-resolution activation maps. *Comput. Biol. Med.* 65:177–183.
10. Bayly, P. V., B. H. KenKnight, ..., W. M. Smith. 1998. Estimation of conduction velocity vector fields from epicardial mapping data. *IEEE Trans. Biomed. Eng.* 45:563–571.
11. Weber, F. M., C. Schilling, ..., O. Dössel. 2010. Wave-direction and conduction-velocity analysis from intracardiac electrograms—a single-shot technique. *IEEE Trans. Biomed. Eng.* 57:2394–2401.
12. Kay, M. W., and J. M. Rogers. 2006. Epicardial rotors in panoramic optical maps of fibrillating swine ventricles. *Conf. Proc. IEEE Eng. Med. Biol. Soc.* 1:2268–2271.
13. Rogers, J. M., M. Usui, ..., W. M. Smith. 1997. A quantitative framework for analyzing epicardial activation patterns during ventricular fibrillation. *Ann. Biomed. Eng.* 25:749–760.
14. Laughner, J. I., F. S. Ng, ..., I. R. Efimov. 2012. Processing and analysis of cardiac optical mapping data obtained with potentiometric dyes. *Am. J. Physiol. Heart Circ. Physiol.* 303:H753–H765.
15. Beier, H. T., and B. L. Ibey. 2014. Experimental comparison of the high-speed imaging performance of an EM-CCD and sCMOS camera in a dynamic live-cell imaging test case. *PLoS One.* 9:e84614.
16. Entcheva, E., and H. Bien. 2006. Macroscopic optical mapping of excitation in cardiac cell networks with ultra-high spatiotemporal resolution. *Prog. Biophys. Mol. Biol.* 92:232–257.
17. Tomek, J., O. Novak, and J. Syka. 2013. Two-Photon Processor and SeNeCA: a freely available software package to process data from two-photon calcium imaging at speeds down to several milliseconds per frame. *J. Neurophysiol.* 110:243–256.
18. Ioannidis, J. P. A. 2005. Why most published research findings are false. *PLoS Med.* 2:e124.
19. Immerkaer, J. 1996. Fast noise variance estimation. *Comput. Vis. Image Underst.* 64:300–302.
20. Entcheva, E. 2013. Cardiac optogenetics. *Am. J. Physiol. Heart Circ. Physiol.* 304:H1179–H1191.

**Biophysical Journal, Volume 111**

**Supplemental Information**

**Ccoffinn: Automated Wave Tracking in Cultured Cardiac Monolayers**

**Jakub Tomek, Rebecca A.B. Burton, and Gil Bub**

Supplementary materials:  
Ccoffinn toolkit: Automated wavefront tracking and  
feature extraction in optical mapping of cardiac cell  
cultures

J. Tomek<sup>1</sup>

R.A.B. Burton<sup>2</sup>

G. Bub<sup>3</sup>

<sup>1</sup>Corresponding author. E-mail: [jakub.tomek.mff@gmail.com](mailto:jakub.tomek.mff@gmail.com). Department of Physiology, Anatomy, and Genetics, University of Oxford, Oxford, United Kingdom

<sup>2</sup>Department of Pharmacology, University of Oxford, Oxford, United Kingdom

<sup>3</sup>Department of Physiology, McGill university, Montreal QC Canada

## S 1 Methodology description

### S1.1 Spike-mining

The pseudocode summarising the spike-mining process is given in Pseudocode S.1 with further details below.

---

**Pseudocode S.1** The algorithm for determining which blobs are active in which frames.

---

```

function blobActivity = getActivity(
    frames, {Sequence of recorded frames}
    mask {Result of blob segmentation}
    smoothingParameter {Kernel width for averaging filter}
    frameRate {Frame rate of the frame sequence}
    activityThreshold {Threshold for finding spike times}
    apdN {Level of APD-N%, on scale 0-1}
    minSpike, maxSpike {Limits of spike duration}
    isDoubleHump {1 if additional filtering is to be applied}
)

begin

diffThreshold = activityThreshold / frameRate;
{Get pixels occupied by each blob.}
blobPositions = getCP(mask);
{Getting and filtering intensity traces}
bI = getBI(frames, blobPositions); {Blob intensities}
bIfilt = filterBI(bI, frameRate, isDoubleHump);
[upstrokeTimes, zeroTimes] = processDVDT(bIfilt, frameRate,
                                         diffThreshold);
blobActivity = findSpikes(bIfilt, upstrokeTimes, zeroTimes,
                         minSpike, maxSpike, frameRate);

end;

```

---

First, *diffThreshold* is computed as *activityThreshold/frameRate*, which scales the *activityThreshold* by an element proportional to the duration between consecutive frames, *dt*. When we later compare discrete derivative of intensity (*dV*) to a threshold, this scaling makes sure we are comparing *dV/dt* to an element proportional to a fixed-value *activityThreshold*, which makes the threshold invariant to sampling frequency used.

Then, a list of pixels belonging to each blob is extracted from the segmentation mask in `getCP`. In the second step (`getBI`), the average brightness of a blob's pixels is measured for all the blobs in all frames and stored in the matrix *bI*, so that:



$$bI(i, j) = \frac{\sum_{\forall k} b(B_k^i(j))}{||B^i||} \quad (1)$$

where  $B^i$  is the list of pixels belonging to the  $i$ -th segmented blob and  $b(B_k^i(j))$  is the brightness of the  $k$ -th pixel belonging to the  $i$ -th blob in frame  $j$ .  $||B^i||$  is the number of pixels belonging to the  $i$ -th blob.

In `filterBI`, the intensity of each blob is first scaled to the 0–1 range and its baseline is subtracted (the baseline is estimated as the intensity after median filtering with the filter size representing 2s duration, thus with the filter size depending on the frame rate, with symmetric treatment at the edges of the trace). Then, the intensity is smoothed using a median filter with filter size of `smoothingParameter`, yielding `bIfilt`. Median filtering has the advantage that it does not change the time of the action potential upstroke, unlike filters based on averaging. If `isDoubleHump` is true, zero-phase averaging of the same filter width is further applied to connect the two humps representing an action potential. Traces with two humps per action potential are commonly seen in temporally filtered traces from dye-free imaging experiments, since both excitation and contraction generate an increase in intensity. The added averaging step helps to connect the two humps so that they are segmented as a single spike for the purposes of activation detection.

In `processDVDT`, times of maximum  $dV/dt$  are determined first ( $dV$  stands for intensity in general, which traditionally represents voltage). For each blob, `dIS` is defined so that `dIS[1] = 0` and `dIS[i] = bIfilt[i] - bIfilt[i-1]`. The variable `dIS` is then thresholded using `diffThreshold` and within each block above threshold, the maximum value of `dIS` is found and saved in `upstrokeTimes`. Furthermore, all times when `dIS` is zero or crosses zero are saved in `zeroTimes`.

Ultimately, the times of activity of each blob are determined in `findSpikes`. For each blob and time of its upstroke  $t_u$ , let us consider the following variables:

- $tBaseline_{prev}$ : The last element of `zeroTimes` smaller than  $t_u$ . This gives a time when the blob is likely to be repolarised before the currently processed action potential.
- $tBaseline_{peak}$ : The first element of `zeroTimes` that is larger than  $t_u$ . This is the time of maximum intensity of the given action potential.
- $tBaseline_{2after}$ : The second element of `zeroTimes` that is larger than  $t_u$ . This gives the time when the blob is likely to be repolarised after the current action potential.
- $intensity_{prev} = bIfilt[intensity_{prev}]$
- $intensity_{peak} = bIfilt[intensity_{peak}]$
- $intensity_{after2} = bIfilt[intensity_{after2}]$

- $baseline = \frac{intensity_{after} + intensity_{before}}{2}$

Then, we define a threshold for the given spike as follows:

$$spikeThreshold = baseline + (1 - apdN)(intensity_{peak} - baseline)$$

The particular blob is then considered active in the interval of frames represented by the segment of  $bIfilt > spikeThreshold$  which contains frame  $t_u$  (note that this happens for each upstroke detected in the particular blob and the times of activity are combined, rather than overwriting one another). Also, if the duration of the interval in ms (obtained using the number of frames taken and  $frameRate$ ) does not fit within  $minAPD$  and  $maxAPD$ , the interval is discarded as an artifact.

## S1.2 Wave segmentation

As mentioned in the main text, the first step of wave segmentation is to split blobs active in a particular frame to separate waves. Then, each wave is divided into border blobs ( $\mathcal{B}$  blobs) and inside blobs ( $\mathcal{I}$  blobs). The  $\mathcal{B}$  blobs are labeled as wavefront blobs ( $\mathcal{W}$  blobs) if they are in the phase of an action potential when voltage rises, and the subset of these that were not wavefront blobs in the previous frame are labelled new-wavefront blobs ( $\mathcal{NW}$  blobs). The whole process is sketched in Pseudocode S.2; described below in depth for a single frame:

Using the output of `getActiveBlobs`, we first extract the locations of blobs that are active in the given frame. The function `filterActiveBlobs` filters the active blobs according to their centroids to retain as active only the ones that have at least  $minDensity$  neighbours within  $distanceThreshold$  pixels, while the rest are rejected (This helps to remove spurious action potentials and helps the later tracking steps focus on the propagation of organised activity.).

Then, a graph  $G=(V,E)$  is created. The vertices  $V$  are the centroids of active blobs in this frame. In `getEdges`, near blobs are linked with an edge:  $\forall v_i, v_j \in V : (v_i, v_j) \in E \Leftrightarrow d(v_i, v_j) \leq distanceThreshold$ , where  $d$  is Euclidean distance and  $distanceThreshold$  a parameter. The graph  $G$  is then decomposed into connected components which represent separate waves.

The next step is to determine which blobs are at the border of the waves. To do this, a *wave mask* is created (in `cleanMask`):

$$wave\ mask(i, j) = \begin{cases} 1, & \text{If } \exists k \text{ such that } (i, j) \in P_k \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where  $P_k$  is the  $k$ -th blob active in the current frame.

---

**Pseudocode S.2** The algorithm for segmenting blobs active in a given frame into waves. The blobs are also labelled according to their position in their respective wave.

---

```

procedure waveSegmentation(
    frames, {Sequence of recorded frames}
    frameRate, {Rate of acquiring frames}
    mask, {a matrix with the same size as a frame,
           with 1 where blobs are, 0 elsewhere}
    blobActivities, {output of getActivity}
    bISmoothed, {as computed within getActivity}
    minDensity, {Minimum density of blobs}
    distanceThreshold) {The minimum allowed distance of waves}
begin
[centroids, dist] = getInfo(mask) {Get centroids and distances of blobs
                                   in mask}
for frame in frames do
begin
    {Part a) Segmenting blobs into waves}
    activeBlobs = getActiveBlobs(frame, blobActivities);
    activeBlobsFiltered = filterActiveBlobs(centroids[activeBlobs],
                                           distanceThreshold, minDensity);
    V = centroids[activeBlobs];
    E = getEdges(activeBlobsFiltered, distanceThreshold);
    waves = getConnectedComponents(V,E, dist);

    {Part b) Determining which active blobs are are border blobs
    and which are inside blobs}
    waveMask = cleanMask(mask, waves);
    wH = getWrapperHull(waves, waveMask, distanceThreshold);
    [insideBlobs, borderBlobs] = splitIB(activeBlobs, wH);

    {Part c) Further segmenting border blobs}
    [wfBlobs, nonwfBlobs] = splitWavefront(borderBlobs, frameRate,
                                           blobBrightnessBS);
end;
end;

```

---




Figure S.1: Four stages of determining the border blobs. In (a), a part of the binary mask of blobs active in a given wave is shown, with the result of its morphological dilation in (b). The wrapper hull, morphological erosion of (b) is shown in (c). Ultimately, (d) shows the pixels that were 1 in (a), but not in (c): these represent the blobs at the border of the given wave.

Then, we compute a *wrapper hull* as a binary mask obtained using the following steps, illustrated in Figure S.1. The name 'wrapper hull' is used due to similarities with a convex hull, However, since the waves are not always convex objects, a convex hull would not work sufficiently well in this application.

1. Morphologically dilate the wave mask with a circular structural element of radius  $\lfloor \frac{distanceThreshold-1}{2} \rfloor$ . Note that due to the waves being at least  $distanceThreshold$  apart by their definition, this operation cannot merge separate waves into a single object.
2. Morphologically erode the wave mask with a circular structural element of radius  $\lfloor \frac{distanceThreshold-1}{2} \rfloor + 2$ .

To obtain  $\mathcal{B}$  blobs (in `splitIB`), the wave mask is taken and all 1s in it that are also 1s in the wrapper hull are set to 0. Only the blobs that were not fully removed from the wave mask are then labelled as  $\mathcal{B}$  blobs: their existence is a consequence of the morphological erosion being performed with a slightly larger structural element than the dilation.

Ultimately (`splitWavefront`), the  $\mathcal{B}$  blobs are further separated according to the slope of  $cISmoothed$  at the given frame; the slope is estimated as the difference between consecutive elements. The border blobs with positive slope at the given frame (i.e., becoming activated) are labelled as  $\mathcal{W}$  blobs; the rest of the blobs stays labelled as  $\mathcal{B}$  blobs.

### S1.3 Wavefront Tracking

The wavefront tracking is performed in a hierarchical way. First, separate waves (connected components of active blobs) in frame  $f$  are linked to separate waves in frame  $f-1$ . For each




Figure S.2: In (a) is an example of tracking problem in a section of a spiral or target wave. In blue are the  $\mathcal{B}$  blobs in the previous frame, in red are the  $\mathcal{NW}$  blobs in the current frame. The resulting tracking arrows are shown in green. Due to the fact that there are more  $\mathcal{NW}$  blobs than  $\mathcal{B}$  blobs, two  $\mathcal{NW}$  are not linked to anything via an arrow. The subfigure (b) contains an illustration of a graph used in computation of minimum-cost maximum-flow.

wave  $w_i$  in frame  $f$ , a wave  $w_j$  in frame  $f-1$  is chosen so that the number of active blobs that are present both in  $w_i$  and  $w_j$  is the maximum possible. If the size of the largest found intersection is 0, no link is formed. These steps ensure that only blobs belonging to the same wave are tracked.

When the high-level links  $(w_i, w_j)$  are formed between waves,  $\mathcal{W}$  blobs in  $w_i$  and  $\mathcal{B}$  blobs in  $w_j$  are then tracked as described below. An illustration of the tracking output of our software is given in Figure S.2a.

Two groups of tracking arrows are defined:

- The trivial group is “zero” arrows: if a  $\mathcal{W}$  blob in frame  $f$  is also a  $\mathcal{B}$  blob in frame  $f-1$ , it means that the wave is static there, not propagating further.
- The second group of tracking arrows contains the arrows representing an actual movement of a wavefront. However, rather than linking full wavefronts at frames  $f$  and  $f-1$ , we link all the border blobs in frame  $f-1$  to the  $\mathcal{NW}$  blobs ( $\mathcal{W}$  blobs in frame  $f$  that were not  $\mathcal{W}$  in frame  $f-1$ ) to obtain a smaller integer programming problem later. Also for increased computational performance, we disregard the  $\mathcal{B}$  blobs in frame  $f-1$  with no  $\mathcal{NW}$  blob in frame  $f$  within  $maxSpeed$ ; these removed border blobs are too far from the wavefront to be linkable to the new wavefront. Finally, tracking arrows longer than  $maxSpeed$  are removed.

To obtain the nontrivial tracking arrows, we form a complete oriented bipartite graph

with  $G = (V_1 \cup V_2, E)$  where vertices in  $V_1$ , the first partite, represent centroids of eligible  $\mathcal{B}$  blobs in frame  $f - 1$  and vertices in  $V_2$  represent  $\mathcal{NW}$  blobs in frame  $f$ , with  $(u, v) \in E \Leftrightarrow u \in V_1, v \in V_2$ . Furthermore, we set  $\forall (u, v) \in E : c(u, v) = 1, a(u, v) = d(u, v)$ , where  $c$  is capacity,  $a$  is cost and  $d$  is Euclidean distance of the vertices. We can then formulate the problem of finding the best tracking arrows as finding the maximum-size matching of the two partites in  $G$  with minimum total cost. Such a matching may be found, for example, by maximising network flow. Using this approach, the source and sink vertices are added to the graph (also with unitary capacities, but zero costs) and linked to their respective partites as shown in Figure S.2b. Then, we aim to find the cheapest flow  $f$  of size  $\phi = \min(|V_1|, |V_2|)$ . In order to make sure that an edge is either fully used in the flow or not at all (i.e., a tracking arrow is present there or not), we formulate the task as an integer programming problem with variables  $x_{(u,v)} = f(u, v)$ :

$$\text{minimize:} \quad \sum_{u \in V_1, v \in V_2} x_{(u,v)} a(u, v) \quad (3.1)$$

$$\text{subject to:} \quad x_{(u,v)} \leq 1 \quad \forall (u, v) \in E \quad (3.2)$$

$$x_{(u,v)} \geq 0 \quad \forall (u, v) \in E \quad (3.3)$$

$$\sum_{(u,v) \in E} x_{(u,v)} - \sum_{(v,w) \in E} x_{(v,w)} = 0 \quad \forall v \in V_1 \cup V_2 \quad (3.4)$$

$$\sum_{(source,v) \in E} x_{(source,v)} = \phi \quad \forall v \in V_1 \quad (3.5)$$

$$\sum_{(v,sink) \in E} x_{(v,sink)} = \phi \quad \forall v \in V_2 \quad (3.6)$$

$$x_{(u,v)} \in \mathbb{Z} \quad \forall (u, v) \in E \quad (3.7)$$

After the optimum of this system is found via the Branch and Bound method, the arrows are filtered. The arrows longer than  $maxSpeed$  are discarded, as well as arrows that are estimated to pass outside wrapper hull in the frame containing the blobs at the end of the arrows; this heuristic aims to discard artifactory arrows between unrelated wavefronts which typically pass through zones outside waves. The exact criterion for discarding arrows is as follows: First, for the  $i$ -th arrow considered,  $L_i$  is the set of pixels under the  $i$ -th tracking arrow, determined using Bresenham's line-drawing algorithm between the arrow's start and end points. Then, let  $\|wh(L_i)\|$  be the number of pixels of  $L_i$  belonging to wrapper hull and  $\|L_i\|$  is the total number of pixels belonging to  $L_i$ . The  $i$ -th arrow is then discarded if  $2\|wh(L_i)\| > \|L_i\|$  and  $\|wh(L_i)\| > 2$ . The first condition serves to discard arrows predominantly traversing outside the wrapper hull, the second condition then protects very short arrows in slow conducting zones: due to the nature of wrapper hull and the fact that erosion is larger than dilation by two pixels, each arrow traversing less than four pixels would be



discarded automatically.

The arrows passing the filtering that are not longer than *maxSpeed* are drawn in places of oriented edges  $(u, v) : u \in V_1, v \in V_2$  where  $x_{(u,v)} = 1$  and the coordinates of their start and end points are saved for future use in feature extraction, along with the “trivial” zero-length arrows representing static wavefronts.

Note that the integer programming formulation above is fairly restrictive, allowing only a single arrow leaving a blob and a single arrow pointing to a blob. This prevents the artifactory behaviour of nearest-neighbour tracking which tends to label several blobs with many outgoing or incoming arrows. The integer programming framework allows relaxing the “at most one in, at most one out” restriction in an elegant way: the capacities of the edges from *source* to  $V_1$  directly correspond to the maximum number of allowed outgoing arrows and the capacities of the edges from  $V_2$  to *sink* directly correspond to the maximum number of allowed incoming arrows. In the following subsection, the required generalisation is described in detail.

For convenience, Ccoffinn contains an implementation of the frequently used method by Bayly et al. (1), although in general case, we recommend using the above-described minimum-cost-maximum-flow approach, as discussed in section S7. For cases when high number of blobs is present and the default approach would be too computationally demanding, Ccoffinn also contains a common and computationally cheap algorithm that can be used in place of the default one: the nearest-neighbour tracking (2). There a wavefront blob in frame  $f$  is linked to the nearest border blob in frame time  $f-1$ . According to our experience, however, the method performs poorly in the case of waves that are not very simple (e.g., in case of spiral or wavelet waves).

### S1.3.1 Generalisation of the MCMF tracking

The tracking algorithm described previously may suffer from the rule that allows a maximum of one arrow exiting and one arrow entering a blob. An example of a potential problem is a target wave spreading faster in one direction than in others. In such a case, conduction velocity and potentially direction of conduction would be under-tracked; a concrete example of this problem is given in section S7.5.

The integer programming may be modified to accommodate for more arrows entering or leaving a blob. The parameter *maxArrows* determines the maximum allowed number of arrows leaving/entering a blob. It can be set to 'auto', which then uses the smallest integer value that is sufficient to cover all the blobs from the larger partite of the tracking graph; e.g., if movement from 4 to 10 blobs is tracked, *maxArrows*='auto' is considered 3. The exact value is computed separately for each instance of the tracking problem. The integer programming may be then modified as follows:

- The upper bound on flow in the inequality 4.2 is increased to *maxArrows* for all edges from the source to the first partite and from the second partite to the sink. Then, e.g.,

when flow 3 enters a vertex in the first partite, this then forces the tracking to send three distinct arrows from the given vertex (the maximum capacities of edges between the partites are still capped at 1, preventing duplicate arrows).

- $\phi$  is redefined from  $\phi = \min(|V_1|, |V_2|)$  to  $\phi = \min(\max(|V_1|, |V_2|), \maxArrows \cdot \min(|V_1|, |V_2|))$ , where  $V_1, V_2$  are the first and second partite respectively. The first component of the minimum makes sure that more tracking arrows are never required than is the size of the larger partite. The second component then makes sure that more arrows than available are never requested of the tracking algorithm. It can be clearly seen that for  $\maxArrows = 1$ , the original algorithm is obtained, while for  $\maxArrows = \text{'auto'}$ , the number of tracking arrows is precisely the size of larger partite.

## S 2 Feature description

Below are given definitions of features, followed by a brief discussion of their physiological relevance.

### S2.1 Feature definitions

#### S2.1.1 Number of blobs

A scalar value giving the total number of blobs segmented in the dataset.

#### S2.1.2 Sparseness of blobs

A multiset containing the distances from each blob to their nearest non-self neighbour, providing information on how evenly are blobs spread in the culture.

#### S2.1.3 Wavefront speed

A multiset of lengths of tracking arrows found throughout the recording, including or excluding the zero-length arrows, based on the value of *useZeroArrows*. The mean value may be used as an indicator of overall conduction velocity, however, note that unlike manual operators who generally focus on the propagation of the fastest wavefronts, mean *wavefront speed* takes into account the speed of wave propagation in all directions where wavefronts are propagating. According to our observations, the 90th or 95th percentile can be used in our datasets to obtain the conduction velocity in the fastest-propagating wavefronts. Using only the longest arrow per dataset is problematic as in complex wave patterns, such as wavelets, long artefactual arrows may appear, linking two parts of the observed wave, that should not be considered representative.

The choice of *useZeroArrows* depends largely on the data available. The motivation for using zero arrows is that in slow-conducting tissue (or very quickly sampled one), a motion may not be captured between consecutive frames, even though it is present. E.g., a straight wave moving one layer of blobs per two frames has a speed 0.5 layer per frame, but the nonzero arrows will be only of length one; these would be captured every other frame, while in the complementary frames, only zero-length arrows are recorded. Thus, averaging all the arrow lengths, one gets the correct answer of 0.5 when considering zero-length arrows. On the other hand, the zero-length arrows may introduce artefactual drop in wavefront speed estimation in, e.g., a well conducting tissue where a dish edge is present - when the wave reaches the dish edge, it will contain zero-length arrows there for several consecutive frames, contributing low CV information to the overall pool of tracking arrows. Thus, the switch *useZeroArrows* should be set to 1 in tissues with low propagation velocity in pixels/frame, while being set to 0 in other cases.

#### S2.1.4 Interspike duration

For each blob, let us consider its vector of activity ( $i$ -th row of *blobActivities*, as obtained in the Section SS1.1). Then, for  $i$ -th blob, let  $IS_i$  be the multiset of lengths of 2-connected sequences of **0**s that are not leading, nor trailing, that cannot be further extended (i.e., bordering with **1**s on both sides of the interval of **0**s). The leading and trailing intervals of **0**s are excluded as they may not represent a full period between two action potentials. Then,

$$interspike\ duration = \uplus_{i=1,\dots,n} IS_i \quad (4)$$

where  $n$  is the number of blobs and  $\uplus$  stands for multiset union (i.e., union with allowed repeats of elements). Ultimately, all the values are scaled by  $1000/frameRate$  to convert the interspike duration in frames to milliseconds.

#### S2.1.5 APD

This feature is an analogy of interspike duration, except that lengths of non-trailing and non-leading intervals of **1**s are computed. Note, that while feature is a correlate of physiological APD, it may not represent it exactly and the level of APD-N it represents may depend on the nature of the data. This is chiefly due the fact that the APD-N is estimated using the median-filtered signal and sharp peaks in the signal (as seen in the synthetic data) are removed. Thus, estimating, for example, APD70 of the filtered signal may not be equivalent to APD70 in the original signal as the thresholds used for measuring action potential shape depend on the signals peak value.

### S2.1.6 Wavefront roughness

*Wavefront roughness* serves to describe the lack of smoothness in signal propagation. A tracking arrow from location  $u$  to location  $v$  in frame  $f$  is called *suitable* if there is another tracking arrow in the frame  $f$  with the origin closer than  $usm$  (a parameter) micrometers from  $u$ . Informally, suitable arrows are these with an arrow nearby in the same frame. For the  $i$ -th suitable arrow  $\alpha$ , we define  $\beta_i$  as the set of tracking arrows in the same frame with their origin closer than  $usm$  to the origin of  $\alpha$  (including  $\alpha$ ). Then, we can define:

$$\text{wavefront roughness} = \bigoplus_{i=1, \dots, m} \overline{S}(\varphi(\beta_i)) \quad (5)$$

where  $\varphi(\beta_i)$  is the set of angular components of polar coordinates (in degrees) of all the arrows in  $\beta_i$ ,  $\overline{S}$  is the circular standard deviation and  $m$  is the number of suitable tracking arrows.

It is expected that when a wavefront is propagating smoothly, such as in straight waves, most of its tracking arrows will point in the same direction, contributing low values to *wavefront roughness*, while chaotic activity with many wavebreaks will display lack of local smoothness (i.e., high  $\varphi(\beta)$ ), contributing high values of *wavefront roughness*.

### S2.1.7 Beat frequency

*Beat frequency* of a recording can be used to estimate the beating rate of blobs in the recording. For a single blob and its brightness trace, we define its beat frequency as the number of spikes per second. The *beat frequency* of the whole recording is then simply a multiset union of the beat frequencies of single blobs.

### S2.1.8 Entropy

This feature is a variant of the spatiotemporal entropy described previously (3), using the representation of a wave as a set of blobs, rather than a set of pixels (which seems to be the approach used by Jung). Let us construct a 3D space  $S$  where  $(x, y, f) \in S \Leftrightarrow \exists \text{ blob } c$  so that  $(x, y)$  is the centroid of blob  $c$  and  $\text{blobActivitiesFiltered}(c, f) = 1$  (*blobActivitiesFiltered* is obtained in Section SS1.2). Informally, a point is present in the space if it represents a centroid of an active blob in a certain frame.

We then define a distance  $d_S$  in the space  $S$ :

$$d_S((x_1, x_2, x_3), (y_1, y_2, y_3)) = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2 + \zeta(x_3 - y_3)^2} \quad (6)$$

where  $\zeta$  is the mean of *Sparseness of blobs* feature. This distance function is essentially an euclidean distance with weighted distance in the z-axis (this is equivalent to an ordinary euclidean distance in a space where frames are spaced  $\zeta$  pixels apart).

Let us define a graph  $G'(V', E')$ , where  $V'$  contains vertices representing all points in space  $S$  and  $(u, v) \in E' \Leftrightarrow d_S(u, v) < distanceThreshold$ . A *spatiotemporal wave* is then defined as a connected component of  $G'$  with the number of the points in the connected component giving its size.

Using notation from the work by Jung (3), spatiotemporal volume of class of size  $s$  is defined as  $V_s = s \cdot n_s$ , where  $s$  is the size of connected component and  $n_s$  is the number of such connected components in the data. Then, relative coverage of a size class  $s$  is  $v_s = \frac{V_s}{\sum_s V_s}$ . Spatiotemporal entropy is then defined as follows:

$$S = - \sum_s v_s \ln v_s \quad (7)$$

Even though entropy is a popular measure of order, providing a single number for the whole dataset, it has its limitations. We can observe, for example that a completely random activity of isolated blobs, yielding many connected components of size 1, has the same (zero) entropy as when there is a single large wave present in the recording. Also, a very heterogeneous wave, such as a group of wavelets, may still be a single connected object in the 3D space, depending on the exact shape of the waves and quantity of noise; this feature then becomes fairly unstable, with small changes to the recording causing large difference in entropy.

Efficient computation of entropy in our context, namely the decomposition of  $G'$  into connected components is not entirely trivial. The only difficulty in computing the entropy is the decomposition of the graph  $G$  into connected components due to large size of the graph itself. With hundreds of active blobs per frame and a video that is thousand frames long (which are both possible in our data), the number of edges present in the graph can be too large to fit into memory of a desktop computer (especially using memory-inefficient representations, such as adjacency matrix).

The first observation that helps is that the whole graph does not have to be represented explicitly at any time for us to obtain the decomposition into connected components. The standard DFS-based algorithm for finding connected components can be adapted to not require the whole graph. Let us recall that an edge is present between vertices in  $G$  if the two vertices are closer than *distanceThreshold*. Thus, the step when the search expands into a vertex's neighbours (usually determined directly from the graph representation) can be modified to instead query all the other unlabelled vertices (candidate neighbours), expanding into these that are near enough to the vertex expanded. While this greatly reduces the memory complexity from  $O(|V|)^2$  to  $O(|V|)$ , it also increases the time complexity from  $O(|V|)$  to  $O(|V|)^2$ , making the computation impractically long.

A second improvement to the basic algorithm is restricting the search of candidate neighbours so that not all vertices are searched. We can divide the 3D spatiotemporal space occupied by the vertices of  $G$  into cubes with side length of *distanceThreshold* and for each cube, we record the list of vertices present in it. Then, when looking for neighbours of a vertex to

be expanded  $v$ , instead of testing the distance to all the vertices in the space, we test only the distance to the same cube and the 26 neighbouring ones. In a standard case of our data, 256x256 pixels and *distanceThreshold* of ca. 16, we search only at most 27 cubes instead of 4096 there are, which is a considerable time-saver. This function has been implemented in C to further improve performance and connected to the main code via MEX interface.

### S2.1.9 Perimeter-to-area ratio

In order to counteract the deficiencies of *entropy*, we have designed a different measure of the degree of order of the recorded activity. For each frame, we obtain the mask of active blobs in the frame (after filtering out blobs in sparse regions, i.e., using *blobActivityFiltered*). Morphological closing with circular structural element and radius *distanceThreshold*/2 is performed to fill holes between blobs belonging to a same wave (but not mixing separate waves together). Then, we define the *perimeter-to-area ratio* in the  $i$ -th frame as:

$$\textit{perimeter-to-area}(i) = \frac{\#sidePixels(m_i)}{\#pixels(m_i)} \quad (8)$$

where  $m_i$  is the processed mask of the  $i$ -th frame,  $\#sidePixels(m_i)$  is the number of one-containing pixels in  $m_i$  that border with a zero (“perimeter”), and  $\#pixels(m_i)$  is the total number of ones in  $m_i$  (“area”).

The whole feature *perimeter-to-area ratio* is then defined as a multiset union of all the  $\textit{perimeter-to-area}(i)$  such that  $\#pixels(m_i) > 0$ .

We can clearly see that a wavelet activity will have overall high *perimeter-to-area ratio* due to large perimeter and low area, unlike, e.g., round target waves, which have large areas and small perimeters.

A slightly different approach would be to compute the surface and volume of a 3D voxel, instead of computing perimeter-to-area ratio for each frame separately. However, such approach would be then very sensitive to wave speed: an organised wave that moves very fast would have larger surface than a wave of the same type that travels more slowly due to larger number of facets between two frames. We consider this effect undesirable, as a measure of wavefront structure should not depend on propagation speed, which is why we have adopted the frame-by-frame approach instead.

A limitation of this approach is that waves entering the field of view (FOV) may have a fairly high *perimeter-to-area ratio* at the given frame, as most of the true area of the wave is not seen. However, such cases are highly unlikely to affect the feature value in a major way, as waves tend to be present in FOV for multiple frames. If this limitation affects the results severely for some unexpected reasons, the code for this feature’s extraction could be modified to ignore several leading and trailing of activity in each wave when *perimeter-to-area ratio* is determined, thus discarding the frames when a wave enters or leaves the FOV altogether.



## S2.2 Physiological relevance of features

*Wavefront speed*, *APD*, and *firing frequency* are widely reported measures of cardiac wave dynamics that yield information on connectivity, upstroke velocity, automaticity, and ion channel kinetics, e.g. calcium dynamics (1, 4, 5). *Interspike duration* is a measure of cardiac recovery time, which is often used to describe the presence or absence of dynamic changes in action potential shape, e.g., alternans (6). The other measures described in the paper (*wavefront roughness*, *entropy*, and *perimeter to area ratio*) attempt to quantify the relative complexity of propagating waves in tissue. Since complex dynamics can be caused by a wide variety of different causes, it is often impossible to assign a one to one electrophysiological or structural correlation to these measures. However, these measures are directly relevant to the understanding wavefront interactions and arrhythmogenesis in our preparations.

*Wavefront roughness* is an indirect measure of the connectivity of the tissue as well as the number and size of local structural heterogeneities. Well connected tissue will generate smooth wavefronts whereas locally disconnected tissue, or tissue in the presence of spatially distributed heterogeneities, tend to generate wavefronts that appear fractured (7–11). *Entropy* is a measure of spatiotemporal complexity (3, 12, 13), and can be used to quantify variability in wavefront structure. Tissue experiencing changes in electrophysiology, or connectivity (14) may be expected to display higher entropy values. However, *entropy* may be difficult to interpret as it does not necessarily map onto an intuitive understanding of complex cardiac dynamics: for example, many multiple interacting wavelets (e.g. fibrillation) may have low entropy if the wavefronts have similar sizes. Entropy calculation is also sensitive to noise and the threshold for activation. In contrast, wave *perimeter to area ratio* directly maps to an intuitive measure of spatial complexity in our data sets. Wave *perimeter to area ratio* gives information on the shape of a wave (larger values map to smaller waves and waves with non-circular shapes). As with the *wavefront roughness* descriptor, *perimeter to area ratio* is a function of local heterogeneity, but also a measure of dynamic instability in homogeneous tissue. For example multiple wavelets caused by steep restitution can have a low roughness value and a high perimeter to area value, whereas multiple wavelets caused by local heterogeneity are expected to have high values for both descriptors. Taken together, these measures can be used to differentiate different experimental preparations (e.g. Table 1).

## S 3 An overview of parameters

Table S.1 shows a list of constants given by physical conditions of the experiment, Table S.2 contains the free parameters of the methods used in our toolkit, and Table S.3 contains the computed parameters, values of which are based on constants and free parameters. The table S.4 contains an overview of switches that control the behaviour of the software.

| <b>Name</b>             | <b>Description</b>  |
|-------------------------|---|
| <i>frameRate</i>        | The frame rate of the recording in frames per second.   |
| <i>mimps</i>            | The upper limit of speed of waves in micrometres per ms.  |
| <i>pixelMicroMetres</i> | How many micrometres does a side of a pixel measure (based on the resolution of the recording). |

Table S.1: An overview of constants that need to be specified as a part of input. These numbers can be different for each video processed in a batch (e.g., all videos of cell cultures created and imaged using a fixed protocol, before and after addition of a drug).

| Name                             | Description  | Value       |
|----------------------------------|--|-------------|
| <i>hlt, llt, cws</i>             | The parameters of the SeNeCA segmentation algorithm for the case when blob-based segmentation is used: High light threshold, low light threshold, and contrast window size respectively (these are used only if SeNeCA is used for blob segmentation). Fourth parameter described in the article on SeNeCA, <i>minLight</i> , is not used, i.e., set to 0.         | 1.4, 1.3, 8 |
| <i>binningfactor, minimumSNR</i> | The parameters of the filtered binning segmentation algorithm: the video is binned into squares with side length <i>binningfactor</i> , ignoring squares traces of which have SNR lower than <i>minimumSNR</i> (these are used only if filtered binning is used for blob segmentation). <i>minimumSNR</i> is 0 by default, meaning that no filtering is performed. | 3, 0        |
| <i>smoothingFactor</i>           | Used to compute <i>smoothingParameter</i> ; the larger it is, the shorter the filter. E.g., for the default value and 100fps, smoothing width is 100ms   | 10          |
| <i>activityThreshold</i>         | Used to determine which segments of <i>dIS</i> (as obtained in Section S1.1) are labelled as containing an AP upstroke   | 2           |
| <i>apdN</i>                      | The N determines the APD level measured; scaled to 0-1. I.e., 0.5 represents APD50. The default value used is suited for data with relatively high amount of noise and disorder.   | 0.3         |
| <i>minAPD, maxAPD</i>            | The minimum/maximum duration of an action potential in ms. If a candidate action potential does not fit within this range, it is discarded.  | 25, 800     |
| <i>distanceThreshold</i>         | If two blobs are closer than this (in pixels), they are considered to belong to the same wave. This parameter may be either given exactly as a number, or as 'auto', being estimated heuristically as 4· median distance of each blob to its nearest non-self neighbours.  | 'auto'      |
| <i>minDensity</i>                | If a blob considered active based on <i>activityThreshold</i> does not contain at least this many other blobs considered active within <i>distanceThreshold</i> , it is relabelled as not active to keep focus on non-sparse activity.   | 5           |
| <i>maxArrows</i>                 | The maximum number of arrows starting and ending in each blob; used in MCMF tracking. This may be an integer value or 'auto'.  | 1           |
| <i>usm</i>                       | Used in feature extraction, determines the radius (in micrometres) around a start point of a tracking arrow in which other arrows are considered for computing roughness of wavefront.   | 750         |

Table S.2: An overview of free parameters that need to be specified as a part of input; these numbers are expected to be the same for all the videos processed in a batch. The third column contains values that we used in experiments described in this text, included in order to improve the reproducibility of our results.

| Name                          | Description   | Formula  |
|-------------------------------|---|--|
| <i>uniformitySurroundings</i> | <i>usm</i> in pixels  | $\frac{usm}{pixelMicroMetres}$                               |
| <i>maxSpeed</i>               | Maximum speed of waves in pixels per frame  | $\frac{1000 \cdot mimpms}{frameRate \cdot pixelMicroMetres}$ |
| <i>smoothingParameter</i>     | Size of averaging filter (in frames) used in smoothing intensity traces of blobsA. It set to make a single cardiac action potential smooth, but not to merge two separate ones. | $\frac{frameRate}{smoothingFactor}$                          |

Table S.3: An overview of parameters based on constants and free parameters.

| Name                      | Description   | Values                |
|---------------------------|---|-----------------------|
| <i>isDebug</i>            | If 1, traces of all the blobs segmented in the processed recording are stored in tmp subfolder, along with their smoothed version and segmentation of activity. | 0/1                   |
| <i>isDoubleHump</i>       | If 1, additional smoothing is applied in spike-mining to connect "double-humps" resulting from pre-processing of dye-free imaged data.                          | 0/1                   |
| <i>segmentationMethod</i> | This switch selects between SeNeCA cell segmentation (for recordings with blobs) and binningSNR segmentation (for continuously activated recordings)            | 'seneca'/'binningSNR' |
| <i>toPlot</i>             | If 1, the segmentation and tracking progress is visualised on-screen  | 0/1                   |
| <i>trackingMethod</i>     | This determines whether the default minimum-cost-maximum-flow approach or Bayly's method is used for tracking wavefronts.                                       | 'MCMF'/'Bayly'        |
| <i>useZeroArrows</i>      | If 1, zero-length arrows are considered in CV estimation. This is to be used in cases of low pixels-per-frame conduction.                                       | 0/1                   |
| <i>verbose</i>            | If 1, progress of computation is output to the console and progress bars are displayed to track the progress of time-consuming steps.                           | 0/1                   |

Table S.4: An overview of switches used in Ccoffinn.

## S 4 Data description

### S4.1 Synthetic data

While it is important to demonstrate usefulness of methodology using real data, such an evaluation is limited as the underlying signal is intrinsically unknown and the irregularity of the activity in real systems makes it complex to analyse. Therefore, in addition to evaluating our software on real data, we also performed an evaluation using synthetic datasets, where we can determine the correctness of the evaluated methodology. The main aim is to investigate the robustness of the method to sampling frequency and noise quantity in the data.

The cardiac activity was modelled using a cellular automaton (15), which allowed us to create various types of waves and to measure them at different sampling frequencies. Below are described three experiments: Syn1, Syn2, and Syn3, which aim to test different properties of the process of feature extraction. All created videos have the resolution of 255x255 pixels and include 2601 blobs (a grid of 51-by-51). Examples of frames of the synthetic data can be found in Figure S.3 and Figure S.5.

In experiment Syn1, three 400 frame videos sampled at 50 fps are compared, containing a spiral wave, a planar wave, and a target wave respectively. These three types of waves are the major types of waves observed in cardiac monolayers and our aim was to confirm that they are segmented correctly.

In experiment Syn2, three different types of a wave (planar, spiral, and wavelet) are visualised at three different sampling frequencies (25, 50, and 100 fps; 200, 400, and 800 frames respectively) and we observe the effect on the values of the features extracted. The recording is generated at 100fps and then subsampled to the respective lower sampling frequencies.

In experiment Syn3, the same three types of wave are observed for 400 frames at 50 fps with three levels of Gaussian additive noise with zero mean and variances 0.02, 0.04, 0.06,..., 1 (on a scale of 0-1) and the focus is on the effect of the noise on the features extracted.

### S4.2 Real data

Two types of slow-conducting cell cultures imaged using a dye-free approach were used to evaluate the effectiveness of the toolkit: a culture of ventricular myocytes from SD rat pups (n=19) and a co-culture of ventricular myocytes grown along stellate cardiac neurons of litter mates of the rats (n=16). The protocol of culturing is given in (16). A temporal filter has been applied to the raw data, replacing  $f_i$  with  $abs(f_i - f_{i-5})$ , where  $f_i$  is the  $i$ -th frame. The filtered data can be interpreted as the amount of motion (contraction or relaxation) in a given pixel. A sample script is provided with Ccoffinn, showing how to apply such filtering to a dye-free dataset obtained using phase imaging.

Additionally, several datasets of calcium-imaged cell cultures (with high CV compared to



the dye-free imaged cultures) were used to demonstrate the usefulness of Ccoffinn for such an imaging modality. The protocol used is described in (17).

## S 5 Implementation notes

The methods for segmentation and tracking of waves, as well as the feature extraction, were written in Matlab, using C code and MEX interface in performance-critical areas. Pure-Matlab alternatives to the C functions are also provided for convenience and greater platform independence. The GLPK library was used as a fast integer programming solver, connected to Matlab via GLPKMEX interface. The cellular automata used to generate the synthetic datasets is written in Java. The runtime of the segmentation and tracking depends greatly on the quantity and the pattern of activity of the blobs in the data. On the real data analysed in this report, the average number of objects segmented in a video was  $\approx 4300$  and the average runtime of segmentation and tracking was 0.11 seconds per frame (with the visualisation and console verbosity turned off). The computer used for the computation was a PC with a six-core processor Intel Xeon E5-1650 v3, 16GB of operating memory and a magnetic hard drive.

## S 6 Results: Validation using synthetic data

Below are described the results of the three experiments Syn1, Syn2, Syn3, described in Section S4.1.

### S6.0.1 Syn1: Results

Figure S.3 shows an output of the segmentation and tracking algorithms on three model waves. The output of the segmentation and tracking algorithms was visually assessed and found to be correct.

### S6.0.2 Syn2: Results

The features extracted from the three datasets (varying by their sampling frequency) show that the sampling frequency of the video generally does not seem to drastically affect the features extracted (note that the change in *number of blobs* is minuscule, as can be judged from y-axis). However, we note that while *entropy* does not seem to be overly affected by the sampling frequency, it does not capture the complexity of waves very well. Intuitively, we would expect straight waves to be the most orderly with wavelets most disorganised. However, it is clearly not so and the straight waves have by far the largest entropy. The reason for this is that in the 3D spatiotemporal space, different iterations of a straight wave




Figure S.3: Three types of synthetic waves, before (top) and after (bottom) segmentation and tracking (we note that the tracking depends on the previous frame). To the left is a straight wave, in the middle is a spiral wave, and to the right is a wavelet wave. The interpretation of colours is identical to Figure 1 in the main text.

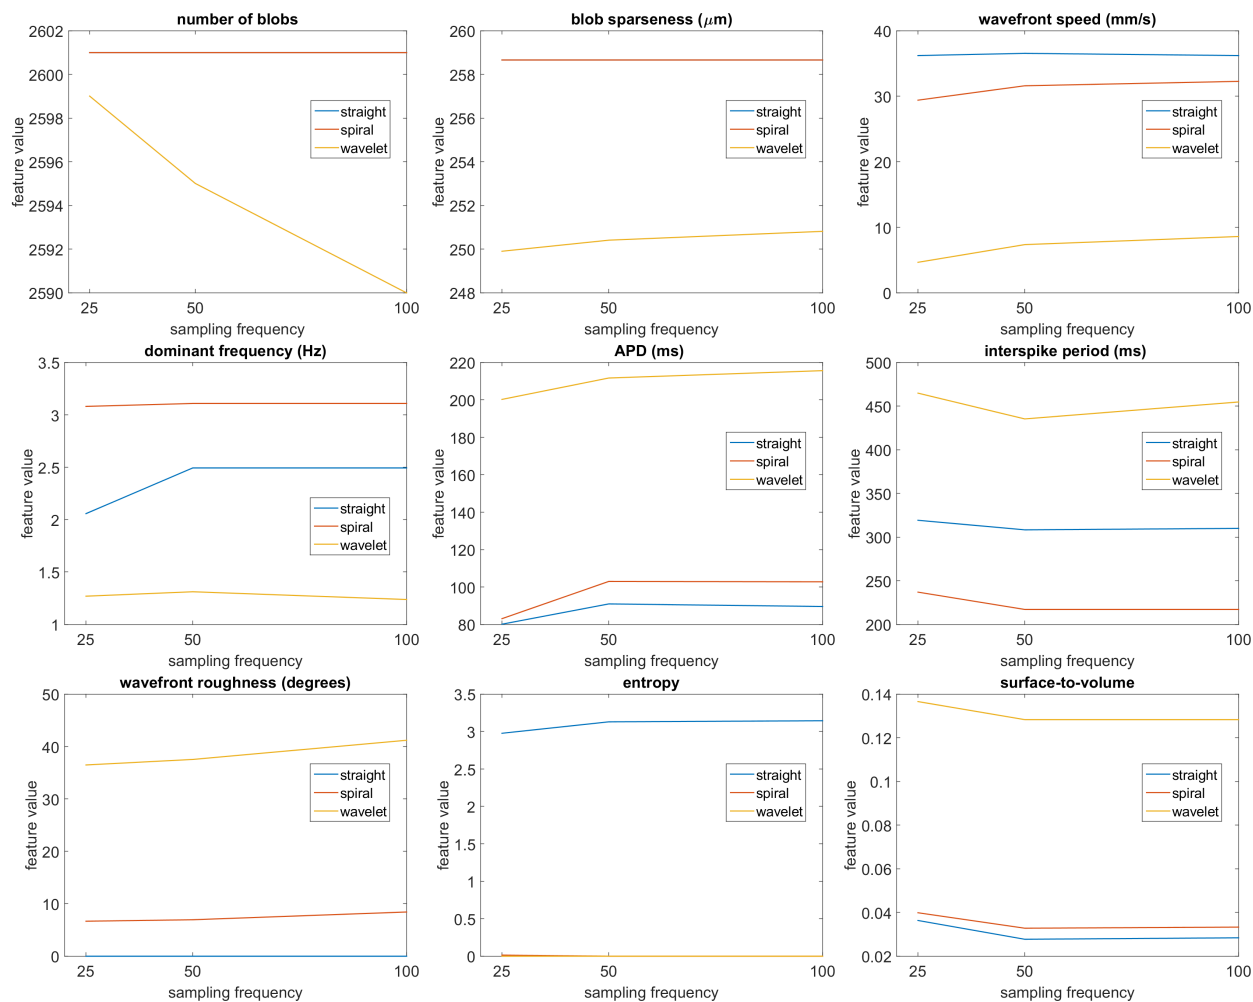


Figure S.4: The effect of sampling frequency on values of features (or means of features in case of non-scalar features). In *number of blobs* and *blob sparseness*, spiral and straight wave have the same values for all sampling frequencies; the same holds for wavelet and spiral wave values in *entropy*. In all the waves, ground truth for *number of blobs* is 2601. For straight waves, ground truth values of *wavefront speed*, *dominant frequency*, and *wavefront roughness* can be estimated; the values are 36.8 mm/s, 2.5 Hz, and 0, respectively.

are completely separate objects of possibly slightly different size. Unlike that, spiral waves (and many wavelet activities) are merely a single connected spatiotemporal object. Thus, even though *entropy* is an elegant and theoretically well-justified method, this flaw is too serious for the feature to be practically useful and alternative methods of order estimation provided in the toolkit are recommended.

### S6.0.3 Syn3: Results

An example of data with four selected levels of additive noise are given in Figure S.5. Figure S.6 shows how features are affected by increasing noise. In features based on segmentation, namely *number of blobs* and *blob sparseness* (min distance to nearest neighbour), the only case with a trend effect of noise is the case of straight waves, where at high noise levels, locations of blobs are sometimes estimated incorrectly with the centroid being slightly shifted. We wondered why there is a difference between straight and spiral waves, which are both organised. Examining the data, we found that the straight-waves pass through any given point less frequently than in the case of spiral waves (this is also manifested in the *dominant-Frequency*) and that the process of segmenting the averaged frames of the recording thus has an input with poorer signal-to-noise ratio in the case of straight waves (relatively more frames of noise for a given pixel than in the case of spiral wave), which is why straight waves are more damaged by the noise. This reveals an important observation that if Ccoffinn was to be used for recordings with very high levels of noise and very sparse waves, the recording should be summarised in a different way than the average of input frames, as this approach is sensitive to noise (e.g., an approach based on high quantiles and/or maximum might be better in such a context, as pixels containing pure noise are unlikely to manifest as high values as pixels that contain genuine activation).

The feature *wavefront speed* tends to decrease with increasing noise level. The reason behind is the fact that overly long arrows are deleted. Considering a single arrow in a noise-free environment, when noise is added, the arrow can stay the same, shorten, or prolong; however, the prolongations are limited as too long arrows are removed. For this reason, the net effect of noise perturbation is mean arrow shortening.

In *dominant frequency*, we see a decrease as the noise is added, accompanied by prolongation of both *APD* and *interspike period*. This is again due to the asymmetric filtering process: an action potential is much more likely to be shortened below the *minAPD* (or even to be completely removed by noise) than prolonged above *maxAPD*.

The *wavefront roughness* grows with noise in the organised waves, as their usually smooth wavefronts become jagged and heterogeneous, increasing the roughness of propagation. We note, however, that at all levels of noise, the three types of waves are ordered in their most natural order of smoothness of propagation. *Entropy* again fails to discern the degree of order associated with the types of waves, similarly to the result from experiment Syn2. *perimeter-to-area ratio* changes little with increasing noise and maintains the ordering of organised

versus disorganised waves at all noise levels. The separation of spiral and straight waves is relatively small, but stable.

Overall, we believe that while the whole process of data analysis and feature extraction is not immune to the quantity of noise present in the data, it is at least fairly resistant, giving good results even at high noise levels in most features (the only exception is *wavefront roughness*). Especially in the case of noise with standard deviation smaller than 0.2, most features are highly stable.




Figure S.5: A spiral wave generated using a cellular automaton with four levels of additive noise (given as the standard deviation of the noise on scale 0-1): 0, 0.2, 0.4, and 1. The interpretation of colours is identical to Figure 1 in the main article.



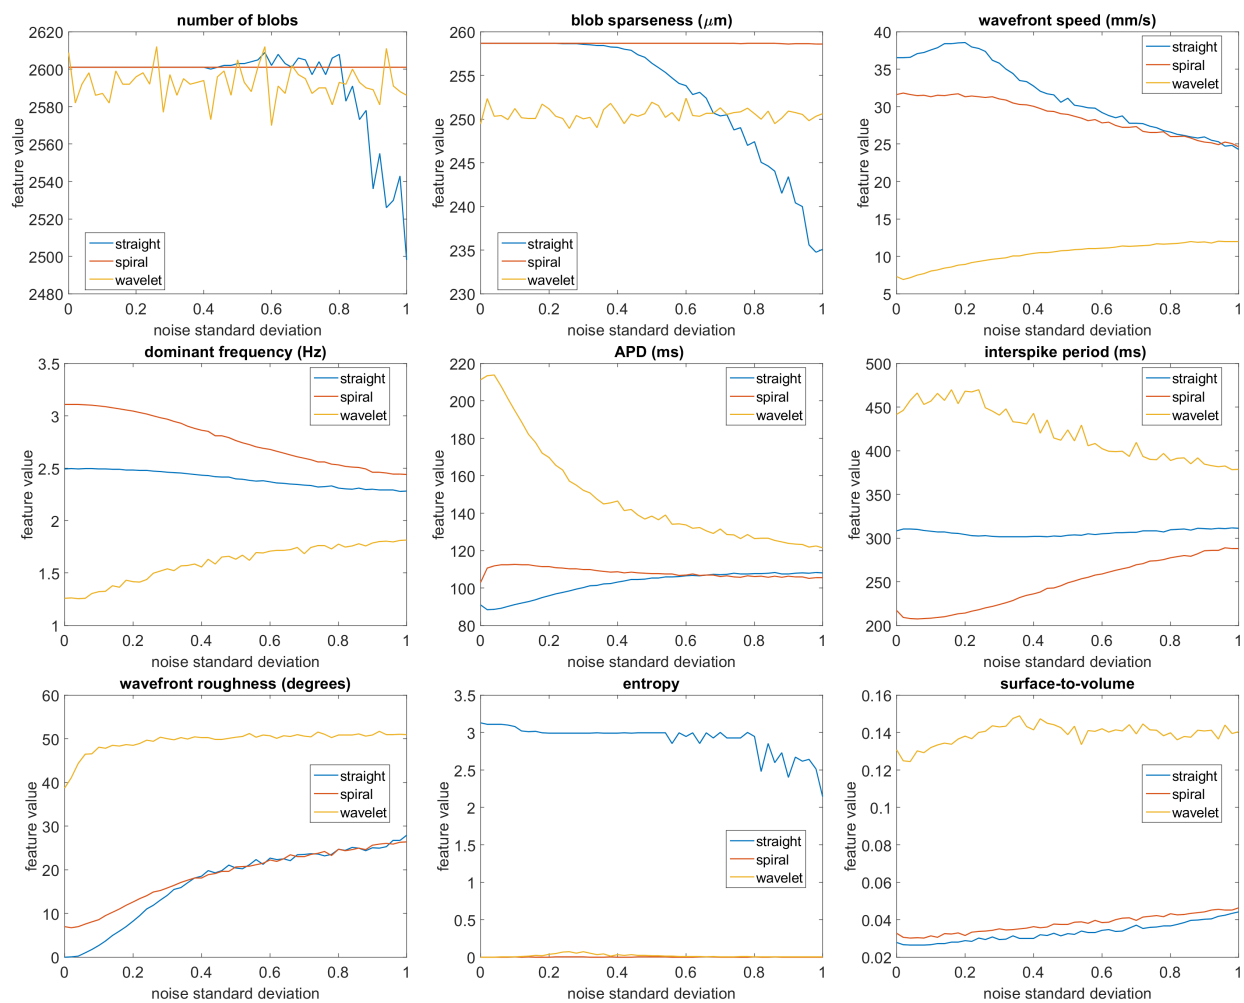


Figure S.6: The effect of increasing noise standard deviation on values of features (or means of features in case of nonscalar features). In all the waves, ground truth for *number of blobs* is 2601. For straight waves, ground truth values of *wavefront speed*, *dominant frequency*, and *wavefront roughness* can be estimated; the values are 36.8 mm/s, 2.5 Hz, and 0, respectively.

## S 7 Bayly’s wave tracking algorithm and MCMF tracking in high-resolution anisotropic data

In this section are discussed the problems associated with the algorithm for wave tracking and conduction velocity (CV) estimation by Bayly et al. (1) in high-resolution, heterogeneously conductive tissue, as well as discussing potential shortcomings of the MCMF algorithm which is the default approach of Ccoffinn.

Bayly’s method considers each “active point” in a wave (e.g., a pixel), finds points in 3D spatiotemporal space that are closer than given thresholds in the x,y,t dimensions, and fits a quadratic surface through these using least squares criterion. Partial derivatives are then used along x,y axes to find a velocity vector for the given active point. A small problem intrinsic to this approach is the fact that the problem of surface fitting may be poorly conditioned, making an estimation of CV impossible in a given point. This problem happened several times in most datasets available to us, however, this error did not occur frequently enough to significantly impact CV mapping or make the results uninformative.

In the relatively high-resolution data available to us, where a wave is represented by up to thousands of active objects and for a reasonably sized spatiotemporal neighbourhood, the basic algorithm is fairly slow, taking almost a minute per frame of data, making high-throughput analysis impossible. However, using Ccoffinn’s framework for wave segmentation, an adjustment to the basic algorithm can be made, where for an active point  $(x, y, f)$  (in frame  $f$ ), “near active points” are taken only from wavefront blobs in frames  $f - 1$ ,  $f$ , and  $f + 1$ , closer than *maxSpeed* in the plane given by first two dimensions of the space. Using only two adjacent frames along the one in which the point  $(x, y, f)$  is used, as a quadratic surface does not have enough degrees of freedom to correctly fit all the points in a larger temporal neighbourhood.

Considering only wavefront blobs, rather than all the active blobs, the algorithm becomes much faster, even though it is still an order of magnitude slower than the minimum-cost-maximum-flow (MCMF) approach described previously. We note that this holds for the data available to us at the moment; it would most likely not scale for extremely large wavefronts of hundreds of thousands blobs, where the potential exponentiality of the integer programming used to solve the MCMF problem would overcome the high-multiplicative-constant polynomial complexity of the Bayly’s algorithm. However, both algorithms would be extremely slow in such a case and an entirely new approach to tracking would have to be adopted.

The main issue associated with Bayly’s algorithm is that in heterogeneously conductive tissue (such as poorly coupled cell cultures, cultures rich in fibroblasts, or intact myocardial tissue with myocardial infarction), where the heterogeneity captured by sufficiently high resolution imaging method is genuine, the surface fitting approach can consistently yield counterintuitive and incorrect results due to its dependence on smoothing. Below are several

examples of wavefront topologies constructed to illustrate certain scenarios, where Bayly’s algorithm fails to provide correct answer. We note that the errors demonstrated below are not not a principal flaw of Bayly’s algorithm, as it was designed for low-resolution electrophysiological data where the problems tend not to occur, but rather an illustration of the algorithm’s struggles with heterogeneously conducting tissue. According to our experience, even in high-resolution data, Bayly’s method performs well as long as the tissue is isotropic/homogeneously conducting. Furthermore, in section S7.5 are also discussed potential issues associated with the non-generalised and generalised MCMF algorithm used in Ccoffinn, showing how both MCMF and Bayly’s method may struggle in certain context.

We visualise the tracking information using 2D plots containing points in the plane, representing the wavefront, where the colour further codes for the time slice in which the points are considered active; the warmer the colour, the later in the “recording”. Estimated vectors are shown using arrows of the given direction and length. When an arrow is missing, it was either a case of a poorly conditioned problem, or the arrow was rejected as too long (longer than neighbourhood radius, which is essentially a representation of maximum conduction velocity). When a particular point is discussed, it is encircled in red and the spatial neighbourhood is shown using a dashed magenta circle (however, only points that are at most 1 frame away from the encircled point may be considered, even if they were close enough in the spatial plane). Furthermore, in selected cases, we show the points in neighbourhood in 3D with the quadratic surface fitted to them.

## S7.1 Slow band in a straight wave

Figures S.7 and S7.1 show a case of a straight wave travelling from bottom to top (with central point of the wavefront slowing down over time, limping behind the rest of the wavefront), with three different radii of neighbourhood shown.

In the panel S.7a are shown velocity vectors for the case when neighbourhood radius is 30. The main systematic error is present at the centre of the wave, where tracking arrows tend to aim towards the centre, rather than purely to the top of the frame. The reason is that in frames 6 and 7, only two points are present within the magenta circle, while in frame 8, three points are present; two a part of “normal” wavefront, one a part of the slowed band. The fitted surface (shown in S.7b) passes exactly through the seven points in the neighbourhood, but is practically wrong, both in the direction, and the conduction velocity estimate. A secondary type of error is present at the edges of the frame, where the optimal fit is not unique and the fitting procedure in Matlab returns unstable results. In practical settings, these cases would be best detected automatically and discarded.

At the bottom of panel S.7c is an identical scenario, but with neighbourhood radius of 60. The problem with unstable borders is solved at the sides (where for each point at the side, 9 points are in its neighbourhood rather than 6 in the previous case, making the optimal solution unique). However, the problem with incorrect estimation of direction in the vicinity

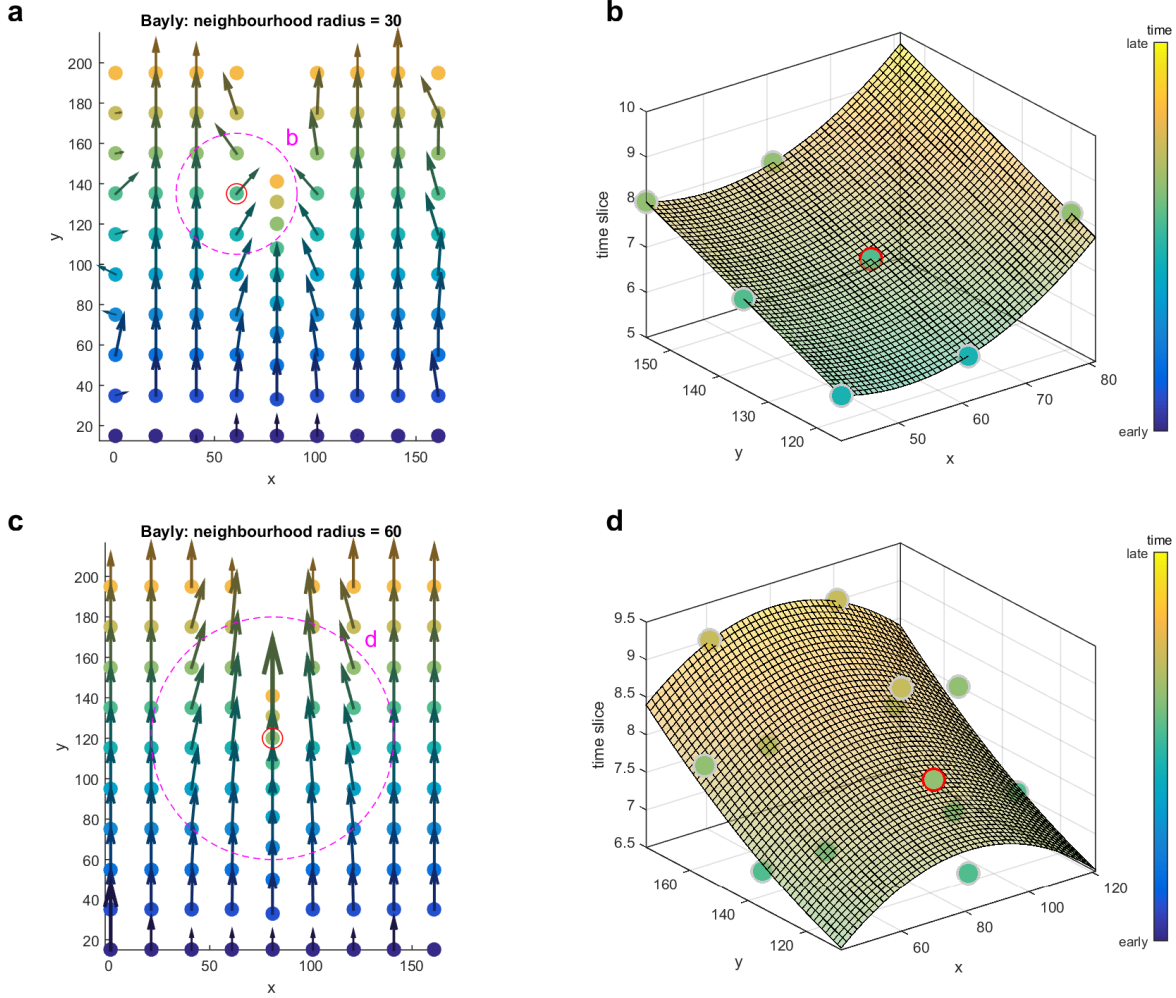


Figure S.7: In the left column are diagrams of wavefront of a straight wave travelling from bottom to top with a slowing-down segment. The colours code for time. Velocity vectors obtained using Bayly’s algorithm are superimposed on the locations of “wavefront blobs”. In a) is the tracking information for neighbourhood radius 30, in b) for neighbourhood radius 60. In the right column are shown quadratic surfaces fitted through the respective points and their surroundings shown in the left column.

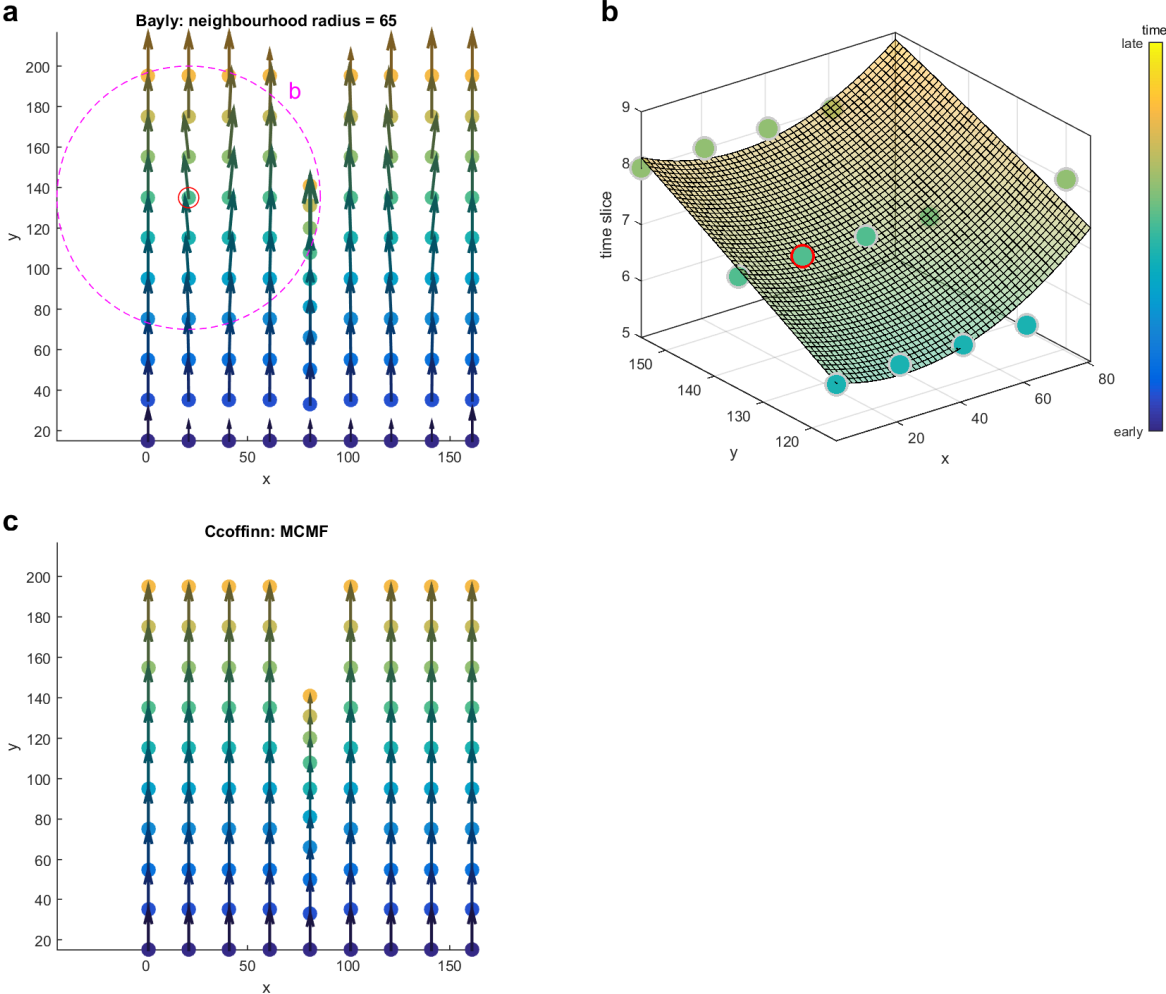


Figure S.8: a) and b) are an analogy of the previous figure, with neighbourhood radius 65. In c) is given the tracking information obtained using Ccoffinn’s MCMF tracking algorithm.

of the slowing down band is aggravated overall. Furthermore, CV becomes overestimated severely in the slowing down band, as demonstrated by the case of the encircled point (producing the longest arrow shown in the plot) and the respective fitted surface. This is due to the fact that when a surface is fitted around the encircled point, the algorithm also considers the points of the fast part of wavefront that are some distance away, which results in an incorrect high CV estimate. Indeed, this error increases as the distance between the fast and slow wavefront increases making the estimated CV excessively large.

Increasing the neighbourhood size to 65 yields the plot in S7.1a, revealing yet another type of error, e.g., in the point circled in red, where the velocity vector points away from the slowing down band. When fitting the respective surface (shown in S7.1b), the surface has to be lifted at its right part to pass near the point belonging to the slow part of the wavefront; however, as the fitted model is a quadratic surface, this forces the left part of the surface to be lifted up, making the partial derivative such that the velocity vector then aims away from the centre of both the fitted surface, and the centre of the velocity plot.

The panel c contains the wave tracking information by Ccoffinn’s MCMF method, providing correct estimation of both CV and conduction direction.

The three examples of various mistakes made by Bayly’s algorithm on the particular data also highlight the sensitivity of the method to the neighbourhood radius, as the tracking outputs in the vicinity of the slowing down segment differ between the three examples rather considerably. At the same time, there is no “correct” neighbourhood size in general. In a heterogeneously conducting excitable medium, the example given above may be only a small part of the overall activation pattern; “good” values of neighbourhood size may then differ substantially between distinct areas of the tissue. As the algorithm allows only a single parametrisation of neighbourhood size, it then follows that at least some areas will contain tracking artifacts. This issue could be probably ameliorated by extending the Bayly’s algorithm to somehow dynamically determine an appropriate neighbourhood size.

## S7.2 Near wavefronts in a single wave

A second type of problem associated with the surface-fitting approach of Bayly is that different wavefronts of a single wave can affect one another; one example is shown in the panel S.9a. The panel shows a wave spreading towards the top and bottom of the plot; it also steers towards left in its top-travelling segment. The steering then affects the direction and length of several arrows associated with the bottom-travelling segment of the wave, which is incorrect. The tracking by MCMF shown in the panel S.9b is correct.

## S7.3 Wave shape affects CV and direction estimation

The panel S.10a shows two waves of identical CV (10 arbitrary units), travelling from bottom to top; a straight wave to the left and a pointed “arrow” wave to the right. While the CV

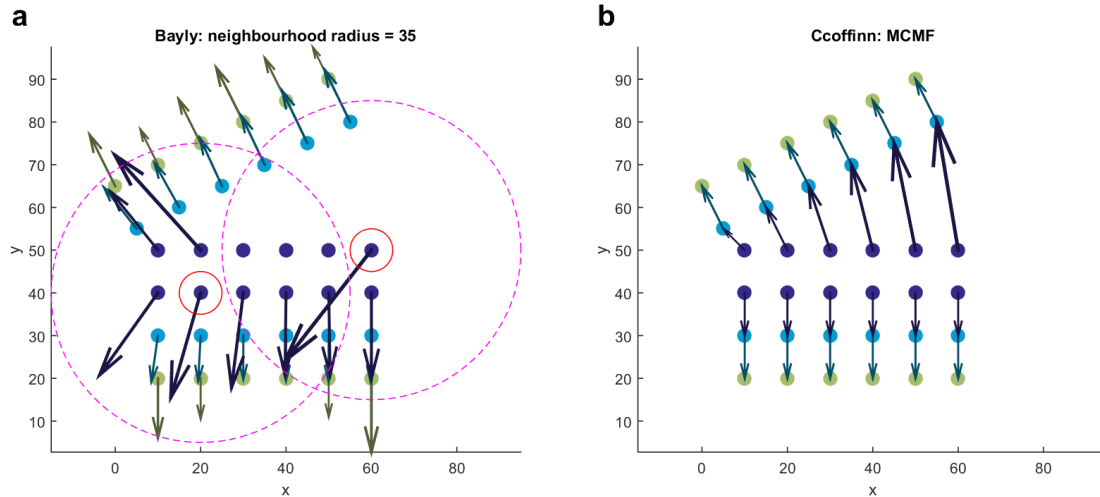


Figure S.9: In a) is shown the activation diagram of a simulated wave spreading simultaneously towards the top and the bottom of the plot; the interpretation of colours and arrows obtained using Bayly's algorithm is identical to the previous figure. In b) is given the tracking information obtained using Ccoffinn's MCMF algorithm.

estimation is correct in the straight wave, when the first and last temporal slices are ignored, both local CVs and propagation directions are almost all incorrect in the arrow-shaped wave, suggesting at a major sensitivity of the Bayly's algorithm to the wave shape, rather than purely to the actual conduction velocity. The panel S.10b gives the surface fitted to the left encircled point's surroundings, showing why CV is mildly overestimated at the point. The panel S.10c then shows the surface fitted to the right encircled point, showing why the direction of propagation is incorrectly estimated. The panel S.10d contains the correct tracking outcome obtained using MCMF algorithm.

#### S7.4 Wavefront jaggedness yields overestimation of CV

In Figure S.11 is shown the tracking output of Bayly's algorithm on a straight wave with jagged wavefront that propagates from bottom to top. Even though the true CV is 10 units per time step, the estimated CV is clearly and systematically estimated to be higher. In certain cases (e.g., the encircled one), the direction is also estimated poorly, due to the asymmetrical positioning of neighbouring active points in the neighbourhood.

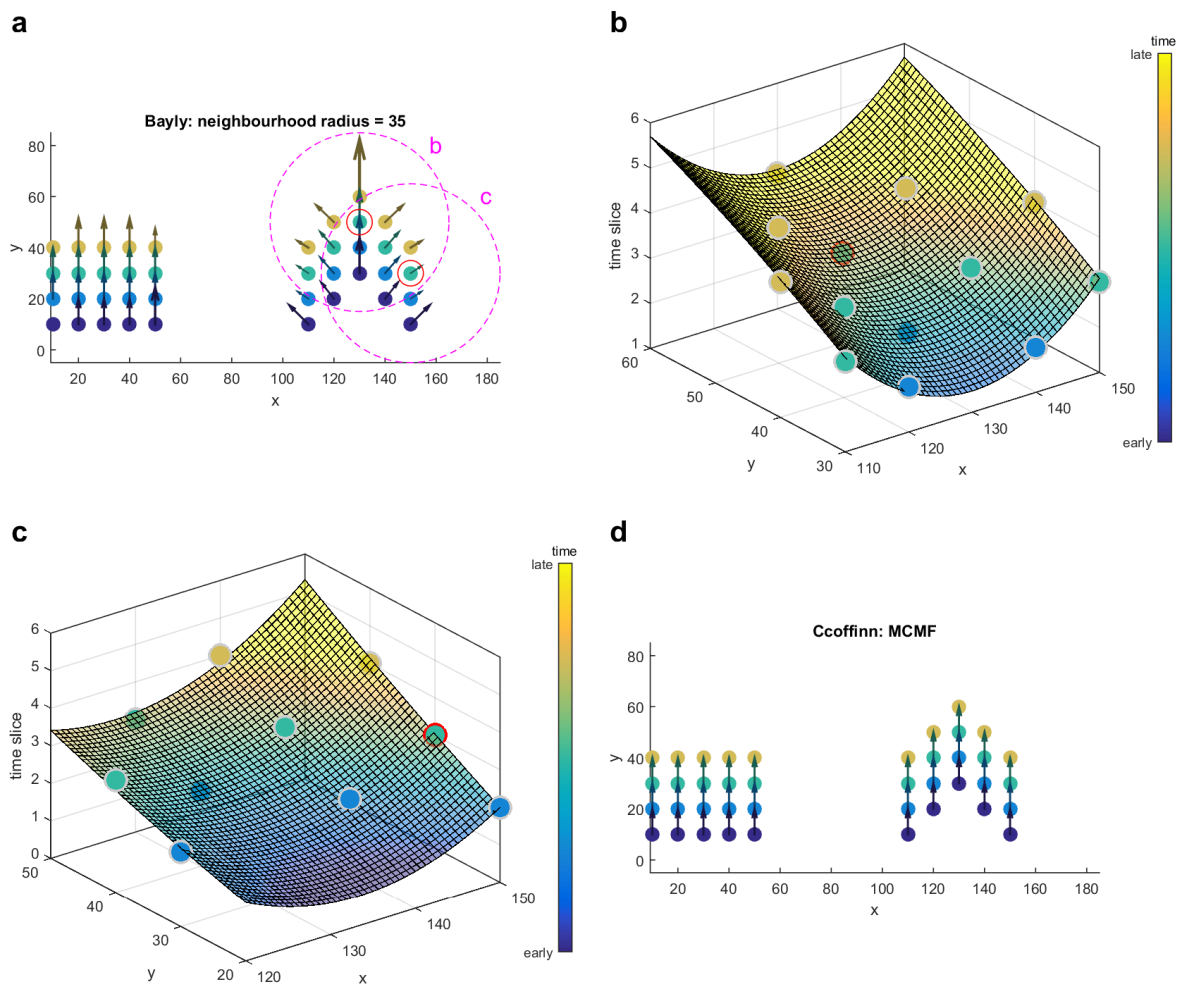


Figure S.10: In a) is shown the activation diagram of two simulated waves of identical conduction velocity: a straight wave to the left, and a arrow-shaped wave to the right; the interpretation of colours and arrows obtained using Bayly's algorithm is identical to the previous figure. In b) and c) are the surfaces fitted to the respective points in a) and their surroundings. In d) is given the tracking information obtained using Ccoffinn's MCMF algorithm.






Figure S.11: In a) is shown the activation diagram of a simulated straight wave with jagged wavefront travelling from bottom to top; the interpretation of colours and arrows obtained using Bayly’s algorithm is identical to the previous figure. In b) is given the tracking information obtained using Ccoffinn’s MCMF algorithm.

## S7.5 Heterogeneously spreading target waves and changing wavefront size

A case where the basic version of MCMF algorithm (allowing at most one arrow entering and leaving a blob) may struggle primarily is the case when wavefront size changes considerably (such as in origination of a target wave), possibly under-tracking the propagation. An example of such under-tracking is shown in Figure S.12a. The issue is the most problematic between frames 1 and 2, where the basic MCMF algorithm generates one arrow where there should be four. The CV is estimated correctly in the frame, but the information on direction of propagation is insufficient. In the last two frames it can be clearly seen how the algorithm chooses the shortest arrows possible, underrepresenting the faster-propagating parts of the wave, thus underestimating CV.

In Figure S.12b is shown the performance of the generalised MCMF algorithm with  $maxArrows = 'auto'$ , showing improved tracking of the previously underrepresented parts of the wave. The perpendicular arrow pairs in certain regions is an artifact caused by the use of synthetic data in this example which rarely observed in real data; weighted spatial averaging of the arrows might be performed to obtain more smooth velocity field. Figure S.12c,d shows the output of Baylys method on the same data (differing in neighbourhood size of 35 and 60 respectively). It is noteworthy that Baylys algorithm also does not work perfectly in this case as arrows leaving blobs in frames 1 and 2 are either not present due to poor conditioning or they are too long and the CV is overestimated in some other regions as well (e.g., the arrow from the blob at the location of  $x=90, y = 100$ ).

Unfortunately, the generalisation of MCMF with  $maxArrows > 1$  does not always perform better than the basic version of the algorithm. The first example is a wave leaving field of view in a corner, as shown in Figure S.13, where the generalised algorithm considers the wave leaving the field of view to be genuinely narrowing, while the basic algorithm and Bayly's method maintain its relative straightness. At the same time, it is important realise that the same wave morphology might happen within field of view due to conduction heterogeneity making a wave narrow down; in such a case, the generalised MCMF with  $maxArrows > 1$  would provide the best tracking information.

There are two cases where MCMF with  $maxArrows > 1$  may generate incorrect tracking information. First, in the case of steep increase in wavefront size, the method becomes more vulnerable to noise causing a non-spiking blob to be randomly detected as spiking in the vicinity of the main wave; in such a case, the larger value  $maxArrows$  is allowed, the more artefactual arrows will enter or leave these "fake-spiking" blobs, polluting the tracking information. A second problematic case happens when a slow wave is imaged at high sampling frequency: even if the wavefront maintains constant size, only a fairly small proportion of the wavefront will be tracked due to the fact that only the became-wavefront blobs are candidates for incoming arrows. At the same time, the set of blobs allowing outgoing arrows (previous-border blobs) will be comparatively large. As a consequence, multiple arrows

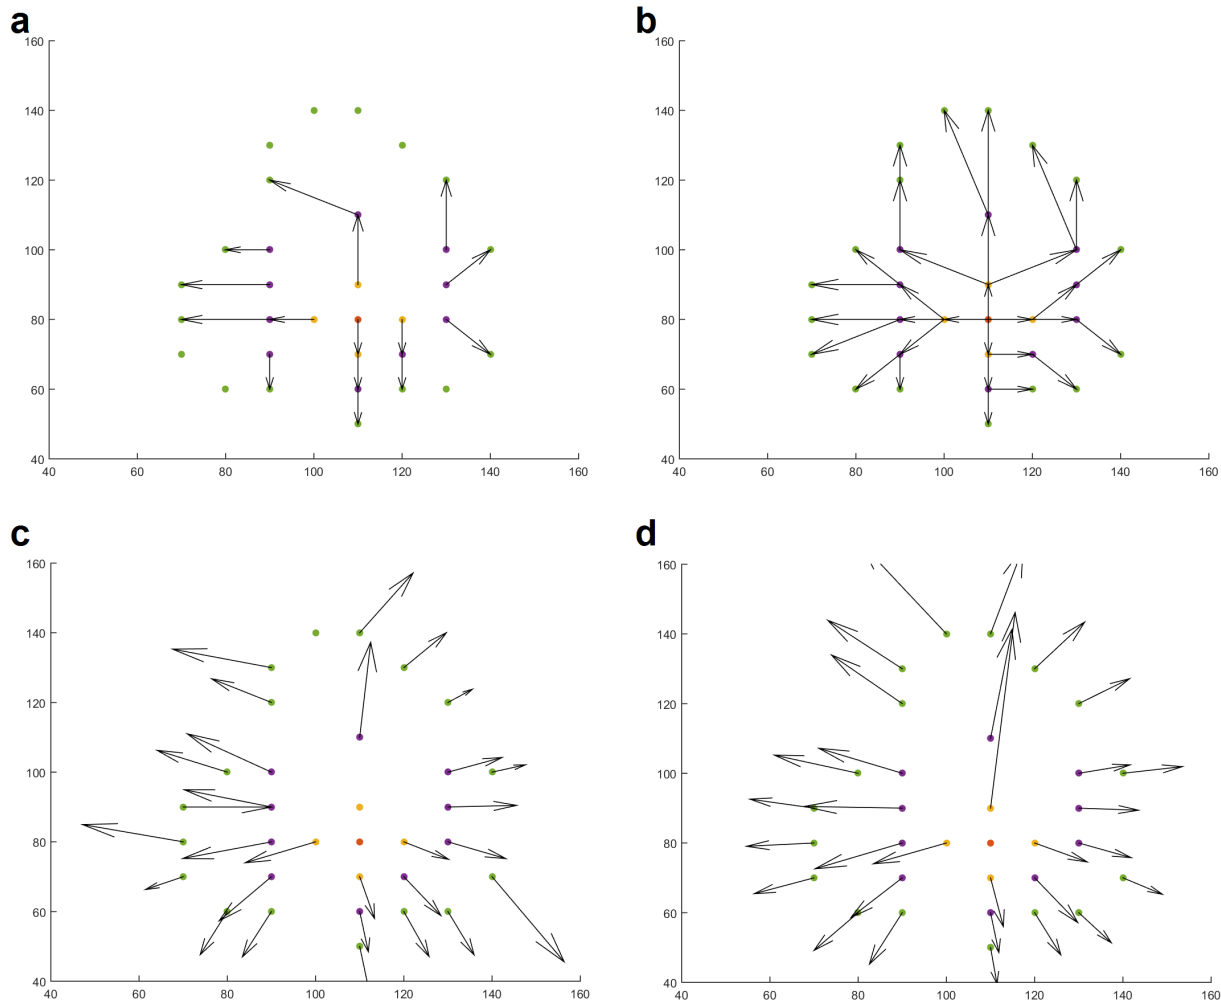


Figure S.12: In a) is shown the tracking of an uneven target wave using the basic MCMF approach. In b) is the same scenario using generalised MCMF tracking with the parameter *maxArrows* set to 'auto'. In c) and d) are shown the outputs of Bayly's algorithm for neighbourhood size 35 and 60 respectively.




Figure S.13: A segment of a spiral wave leaving the field of view, as tracked by MCMF with  $maxArrows = 'auto'$  (a),  $maxArrows = 1$  (b), and Bayly's method (c).




Figure S.14: Four frames of an originating target wave tracked by MCMF with  $maxArrows = 1$ .

will be pointing to most became-wavefront blobs, most of these being necessarily wrong, overestimating conduction velocity and roughness of propagation.

We note that in real use, an obvious underrepresentation of propagation with MCMF and  $maxArrows = 1$  happened only fairly rarely in the data available to us and the cases with the largest potential for step difference (such as the first two frames capturing the origin of a target wave) are implicitly treated by the fact that Ccoffinn requires waves to have a sufficient size/density to be tracked in the first place. In Figure S.14 is shown a sequence of four frames showing origin of a target wave (sample dataset **dataCa2\_target**), as tracked using MCMF with  $maxArrows = 1$ , showing no major systematic flaw of tracking. The problem of under-tracking is somewhat present in the top left part of the wave between frames 3,4, but the bulk of the wavefront is tracked correctly, even the bottom part which is also heterogeneously quickly conducting.

For the reasons stated above, we believe that too high values of  $maxArrows$  are expected to bring more problems than benefit and values of 1 or 2 are likely to work best, unless a particular dataset requires otherwise.

## References

1. Bayly, P. V., B. H. KenKnight, J. M. Rogers, R. E. Hillsley, R. E. Ideker, and W. M. Smith, 1998. Estimation of conduction velocity vector fields from epicardial mapping data. *IEEE Transactions on Biomedical Engineering* 45:563–571.
2. Konstantinova, P., A. Udvarev, and T. Semerdjiev, 2003. A study of a target tracking algorithm using global nearest neighbor approach. *Proceedings of the 4th International Conference on Computer Systems and Technologies E-Learning - CompSysTech '03* 290–295. <http://portal.acm.org/citation.cfm?doid=973620.973668>.
3. Jung, P., J. Wang, R. Wackerbauer, and K. Showalter, 2000. Coherent structure analysis of spatiotemporal chaos. *Physical review. E, Statistical physics, plasmas, fluids, and related interdisciplinary topics* 61:2095–8. <http://www.ncbi.nlm.nih.gov/pubmed/11046503>.
4. Miragoli, M., G. Gaudesius, and S. Rohr, 2006. Electrotonic modulation of cardiac impulse conduction by myofibroblasts. *Circulation Research* 98:801–810.
5. Campbell, K., C. J. Calvo, S. Mironov, T. Herron, O. Berenfeld, and J. Jalife, 2012. Spatial gradients in action potential duration created by regional magnetofection of hERG are a substrate for wavebreak and turbulent propagation in cardiomyocyte monolayers. *The Journal of physiology* 590:6363–79. <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3533198&tool=pmcentrez&rendertype=abstract>.
6. Edwards, J., and L. Blatter, 2014. No TitleCardiac alternans and intracellular calcium cycling. *Clinical and Experimental Pharmacology and Physiology* 41:524–532.
7. Bub, G., A. Shrier, and L. Glass, 2002. Spiral wave generation in heterogeneous excitable media. *Physical review letters* 88:058101.
8. Steinberg, B. E., L. Glass, A. Shrier, and G. Bub, 2006. The role of heterogeneities and intercellular coupling in wave propagation in cardiac tissue. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences* 364:1299–1311.
9. Rohr, S., 2004. Role of gap junctions in the propagation of the cardiac action potential.
10. Fast, V. G., B. J. Darrow, J. E. Saffitz, and a. G. Kléber, 1996. Anisotropic activation spread in heart cell monolayers assessed by high-resolution optical mapping. Role of tissue discontinuities. *Circulation research* 79:115–127.
11. Spach, M. S., and J. M. Kootsey, 1983. The nature of electrical propagation in cardiac muscle. *The American journal of physiology* 244:H3–H22.

12. Schlemmer, A., S. Berg, T. K. Shajahan, S. Luther, and U. Parlitz, 2015. Quantifying spatiotemporal complexity of cardiac dynamics using ordinal patterns. *In* Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS. volume 2015-November, 4049–4052.
13. Schlemmer, A., S. Berg, T. Shajahan, S. Luther, and U. Parlitz, 2015. Entropy Rate Maps of Complex Excitable Dynamics in Cardiac Monolayers. *Entropy* 17:950–967. <http://www.mdpi.com/1099-4300/17/3/950>.
14. Bub, G., A. Shrier, and L. Glass, 2005. Global organization of dynamics in oscillatory heterogeneous excitable media. *Physical Review Letters* 94.
15. Bub, G., K. Tateno, A. Shrier, and L. Glass, 2003. Spontaneous initiation and termination of complex rhythms in cardiac cell culture. *Journal of cardiovascular electrophysiology* 14:S229–S236.
16. Burton, R. A. B., A. Klimas, C. M. Ambrosi, J. Tomek, A. Corbett, E. Entcheva, and G. Bub, 2015. Optical control of excitation waves in cardiac tissue. *Nature Photonics* 9:813–816.
17. Ambrosi, C., and E. Entcheva, 2014. Optogenetic control of cardiomyocytes via viral delivery. *Methods Mol Biol* 215–228.