

Early signs of critical slowing down in heart surface electrograms of ventricular fibrillation victims.

Berend Nannes¹, Rick Quax^{1,2}[0000-0002-0299-0074], Hiroshi Ashikaga³[0000-0003-3689-6892], Mèlèze Hocini^{4,5,6}[0000-0002-7957-8766], Remi Dubois^{5,6}[0000-0002-3320-6086], Olivier Bernus^{5,6}[0000-0003-3917-5791], and Michel Haïssaguerre^{4,5,6}[0000-0001-7644-6146]

¹ Computational Science Lab, University of Amsterdam Amsterdam, The Netherlands

² Institute of Advanced Studies, University of Amsterdam, Amsterdam, The Netherlands

³ Cardiac Arrhythmia Service, Johns Hopkins University School of Medicine, Maryland, USA

⁴ Electrophysiology and Cardiac Stimulation, Bordeaux University Hospital, Bordeaux, France

⁵ IHU LIRYC, Electrophysiology and Heart Modeling Institute, Bordeaux, France

⁶ University of Bordeaux, Bordeaux, France

Abstract. Ventricular fibrillation (VF) is a dangerous type of cardiac arrhythmia which, without intervention, almost always results in sudden death. Implantable automatic defibrillators are among the most successful devices to prevent sudden death by automatically applying a shock to the heart when fibrillation occurs. However, the electric shock is very painful and could lead to dangerous situations when a patient is, for example, driving or biking. An early warning signal for VF could reduce the risk in such situations or, in the future, reduce the need for defibrillation altogether. Here, we test for the presence of critical slowing down (CSD), which has proven to be an early warning indicator for critical transitions in a range of different systems. CSD is characterized by a buildup of autocorrelation; we therefore study the residuals of heart surface electrocardiograms (ECGs) of patients that suffered VF to investigate if we can measure positive trends in autocorrelation. We consider several methods to extract these residuals from the original signals. For three out of four VF victims, we find a significant amount of positive autocorrelation trends in the residuals, which might be explained by CSD. We show that these positive trends may not be measurable from the original body surface ECGs, but only from certain areas around the heart surface. We argue that additional experimental studies involving heart surface ECG data of subjects that did not suffer VF are required to quantify the prediction accuracy of the promising results we get from the data of VF victims.

Keywords: Critical slowing down · Ventricular fibrillation · Critical transition · Early warning signal.

1 Introduction

Ventricular Fibrillation (VF) is a dangerous type of cardiac arrhythmia where the ventricles, instead of beating normally, tremble uncontrollably. As a result, the heart is unable to regulate the blood circulation around the body, resulting almost always in sudden death. The most important contributions to the treatment of VF to date are automated defibrillators. For example, an implantable cardioverter-defibrillator (ICD) can be used for patients at an increased risk of suffering dangerous arrhythmia. The ICD detects cardiac arrhythmia and automatically intervenes by applying an electric shock. For these devices it is crucial to correctly detect the onset of VF in real time to be able to prevent permanent damage or death of the patient. In earlier research, several algorithms have been proposed to this end [1]. The ICD does not prevent the arrhythmia itself, but provides treatment by shock after the occurrence. Intervention with ICD is not without drawbacks: the electric shock is experienced by patients as very painful [2] and can in some cases cause psychological problems [3]. Moreover, the shock could lead to dangerous situations if the patient is, for example, driving or biking. In the latter case, an early warning signal seconds before the shock could greatly reduce risk of accidents. Furthermore, early warning signals in combination with future methods to prevent VF from happening would reduce the need for defibrillation. Many possible indicators have been proposed in earlier work [4] from which Heart Rate Variability (HRV) is the most promising. However, contrary evidence has also been found, as HRV is influenced by characteristics - for example, medical conditions - of the individual patients [5]. Evidently, a generic early warning signal for VF that is independent of patient characteristics is highly desirable.

If we look at VF from a general perspective, we can argue that the shift from a state where the ventricles are pumping normally to a state in which they quiver is a sudden transition between two different dynamical regimes for which the manifestation of VF is the tipping point. Such abrupt changes in dynamical behavior are seen in many real complex systems in nature and are generally called "critical transitions". In critical transitions, once the tipping point is exceeded, it is not easy to return to the previous state. A real-life example of this is desertification: once a patch of land reaches a barren state, it is hard for vegetation to reappear. This "irreversible" character has made possible predictors of these critical transitions much sought-after. Earlier research [6] has shown that there exists a domain-free early-warning signal for critical transitions in complex dynamical systems in different fields of research: critical slowing down (CSD). The theory behind CSD is based on the fact that, mathematically, some critical transitions in real systems can be interpreted as catastrophic bifurcations. It has been shown [7] that systems that approach such bifurcations experience a decrease in resilience; the system needs more time to recover from perturbations when a critical event is close. This decrease in resilience can be measured by increasing autocorrelation and standard deviation in the corresponding time series data (as we will explain in the Method section). These symptoms have indeed been identified in a wide range of real complex systems. For example: the ending

of multiple different glacial periods by abrupt climate changes are preceded by the building up of autocorrelation in deuterium measurements [8], and brain activity shows increasing amounts of variance when close to an epileptic seizure [9]. Here, the same underlying principle applies, independent of the scientific field. If we assume that the onset of VF can be viewed as such a critical transition, we expect that we can detect the same early warning signals that are found in a variety of other real systems.

Historically, there has been controversy about the mechanism that drives VF [10], with some research suggesting that fibrillation represents a chaotic system [11], and others stating that it is rather similar to a nonchaotic random signal [12]. Nowadays it is commonly believed that the onset of VF is a transition to spatiotemporal chaos [13,16] and thus a shift to a different, chaotic attractor. This transition is initiated by a wavebreak that arises when a wavefront and waveback of the cardiac excitation meet [15]. Under normal circumstances this never occurs, but in certain conditions the propagating impulse does not die out but returns to reexcite the heart (called reentry [14]). When reentry triggers a wavebreak, it can in turn produce daughter waves, causing new wavebreaks etc., quickly degenerating into spatiotemporal chaos and VF. The transition from normal heart rhythm to abnormal, chaotic, heart rhythm (the process from reentry to VF) takes place through a series of bifurcations [16]. There exist different theories about how we can exactly describe this route to chaos [17]. In the context of nonlinear dynamics, however, we can consider the shift from a normal heart rhythm to VF as a change in topology of electrical wave dynamics or a transition between two states with different basins of attraction [18]. This drastic change in dynamical regime may therefore bear resemblance to critical transitions in other complex systems and may possibly be signaled by CSD.

In this thesis, we investigate if CSD can be observed in the residuals of heart surface electrocardiogram (ECG) recordings from patients that suffered VF. We analyze data sets from four patients provided by IHU LIRYC (Electrophysiology and Heart Modeling Institute) in Bordeaux, France. All patients suffered sudden cardiac arrest due to documented VF resulting from ischemic heart disease (n=2), early repolarization syndrome (n=1), or idiopathic VF (n=1). All patients were male and the age range was between 15 and 74 years old. Each set contains around 1400 ECG signals (leads) over the heart surface, which are estimated using body surface ECG measurements by solving an ill-posed, inverse problem [19]. The original body surface potential maps consist of around 250 leads. Each lead in the set is a 20-second ECG recording; 10 seconds of normal heart rhythm followed by 10 seconds of arrhythmia, with a sampling rate of 1 kHz. We examine the 10 seconds of the signal that precede the tipping point looking for CSD. Specifically, we look for a significant increase in autocorrelation in the residual.

The main challenge of our research is the actual extraction of the residuals; the short-term fluctuations relative to the main ECG waves. To expose the residual, we have to filter out the wave components that are typical to the ECG. The difficulty resides in the high frequency characteristic of some of these typical

waves, which impedes the use of a simple frequency filter. In this thesis, we go over several alternatives to correctly obtain the residuals, which is essential for credible autocorrelation measurements.

Electrocardiography is traditionally prone to different types of noise like power line interference, electrode contact noise and motion artifacts. We have to keep in mind that some of this noise (in particular, high frequency noise) we can not distinguish from fluctuations in the real signal and will end up in the residual. The ill-posed nature of the inverse solution may also induce errors in the residual. Therefore, to quantify the prediction accuracy of autocorrelation measurements for ECG data from VF victims, we have to conduct the same analysis on heart surface ECG estimations of subjects that did not suffer VF. The main goal of this thesis to show the necessity of such data.

2 Method

Critical slowing down (CSD) is the increase of recovery time needed by the system when it is perturbed from the stable state. In our case, this stable state is described by the well-known features of an ECG signal. Therefore, to be able to measure a possible CSD effect, we aim to find the fluctuations relative to the typical ECG curve: the residual. In this section we discuss several methods we considered to extract the residuals from the original signals and argue which method is the most preferable. To look for CSD, we have to analyze the time evolution of the autocorrelation of the extracted residuals. Because autocorrelation measures the likeness of a signal with a lagged version of itself, a slowly varying signal has a higher autocorrelation than a rapidly fluctuating signal. For that reason it is expected that, if residuals show a slowing down effect, we measure a significant positive trend in autocorrelation.

2.1 Measuring the autocorrelation

The most straightforward autocorrelation measure for equispaced data is the lag-1 autocorrelation, where the state of the signal at time t is directly compared to its state in the previous time unit $t - 1$. The lag-1 autocorrelation can be estimated by treating the signal as a first-order autoregressive (AR(1)) process and calculating the corresponding lag-1 autoregression coefficient. We estimate this parameter using Yule-Walker equations with Python library *statsmodels.tsa.stattools*. To capture the time evolution of the autocorrelation, the autoregression coefficient is calculated over a moving window. The window length is chosen to be exactly half the signal length, so that there is a reasonable trade-off between a sufficiently long window to compute the autocorrelation, and a long enough sequence of autocorrelation values to be able to study its time-evolution. (The influence of different window lengths on our result is shown in Appendix A.10)

2.2 Extracting fluctuations

The typical ECG signal is composed of a set of main wave components (PQRST) formed by electrical currents produced by the depolarization and re-polarization of different heart chambers. Depolarization is responsible for the contraction of cardiac muscle, while with re-polarization, cardiac muscle relaxes. A schematic representation is provided in Fig. 1. The P-wave is produced by the depolarization and contraction of both atria. The QRS-complex is composed of the electrical signals from both the depolarization of the ventricles and the re-polarization of the atria. Finally, the re-polarization of the ventricles produces the T-wave. The orientation of the waves is dependent on the polarity (positive or negative) of the electrode.

Normally, an effective technique to extract short term fluctuations from a signal would be to filter out the lower frequency components using by applying a high-pass filter. However, the QRS-complex has a significantly higher frequency range than the P- and T-waves. A high-pass filter that removes all characteristic waves requires a high cutoff frequency, and thus imposes risk of filtering out short-term fluctuations that possibly show a CSD effect. The same problem arises for methods using wavelet decomposition: by removing high frequency subbands from the signal to remove the QRS complex we might accidentally remove fluctuations we want to analyze. For this reason, we have to consider other methods to extract the residuals: by utilizing the knowledge of the recurring wave components we create a model for the signal containing all characteristic waves and subtract it from the original signal.

In the following subsections we go over techniques we considered to extract the residues of the ECG leads.




Fig. 1. The main wave components of an ECG; the P-wave, QRS-complex and T-wave.

Pre-processing We aim to remove characteristic components (Fig. 1) from the signal to extract the fluctuations. Baseline wandering in the signal makes it hard to distinguish wave components from the zero-volt level. Moreover, the height of characteristic waves may vary between heartbeats. In all methods discussed in this section, the baseline trend is first removed from the signal. Some methods require the signal to be cut into segments of one heartbeat cycle. These pre-processing steps are described in detail in Appendix A.1.

Fitting Gaussian curves To isolate the fluctuations in the ECG signal, we can directly subtract the characteristic ECG waves for every cardiac cycle. By fitting a mixture of Gaussian curves to the original ECG that approximate the PQRST-waves we can construct an characteristic version (as shown in Fig. 1) of the signal, which can in turn be subtracted to reveal the residual. The implementation of this method is explained in more detail in Appendix A.2.




Fig. 2. Segments of detrended ECG signals, their corresponding model fit and the residuals after subtracting the model. **A:** The model seemingly fits the signal. However, we observe high frequency errors in the residual (red arrows) resulting in peaks and dips in the AR(1) measurements (right) occurring at the same frequency as the QRS-complexes. **B:** The signal, while we can identify the PQRST-waves, slightly deviates from the characteristic shape portrayed in Fig. 1. The fitting method fails to reproduce the characteristic waves resulting in periodically recurring errors in the residual (green arrows).

Techniques modeling ECG waveforms using Gaussian curves have been around for some time [21] and are generally used to extract clinical features like the location, height and width of the characteristic waves. When we apply our method to the ECG data we observe that indeed, for signals that resemble the PQRST-composition portrayed in Fig. 1 like the example given in Fig. 2A, the Gaussian fitting method seemingly provides a good model. However, when we subtract the model fit, the parts of the residual at the location of the QRS-complex clearly

have higher frequency compared to other parts of the residual, which indicates that they are error artifacts induced by the fitting method. The effect of these errors is clearly visible in the AR(1) measurements: rapidly fluctuating parts of the signal have relatively low autocorrelation; combined with the moving window this results in peaks and dips in AR(1) values. These peaks and dips recur with the same frequency as the QRS-complexes. Clearly, while this modeling technique may be suitable for clinical feature extraction, it is not ideal to extract correct residuals of the signals. Moreover, to obtain a correct model, the signal has to resemble the characteristic shape from Fig. 1. When the signal deviates from this shape the fitting technique becomes infeasible, resulting in errors in the residual. An example of this is given in Fig. 2B. We conclude that this method is ineffective for the extraction of fluctuations from the signal since the majority of the signals in the provided data set do not have the required characteristic PQRST-composition.

Computing an average beat The method described above to extract fluctuations from the signal turns out to be ineffective as it requires each signal to have a certain characteristic shape. In reality, characteristics may be different for each signal. To capture characteristic features of a signal, one can also construct an average beat using the individual beat cycles. Characteristic features recur periodically and are thus automatically present in the average curve. Everything relative to the average beat can be classified as fluctuations. We cut the signal into segments containing one beat cycle. The fluctuations can be extracted beat by beat by subtracting the average beat from every cycle. However, the height as well as the overall shape of the characteristic waves is not constant over the whole signal. Therefore, the average curve must be adjusted for each beat to provide a good model for the signal. A detailed description of the average beat fit and these adjustments is given in Appendix A.3. It turns out that in almost all signals one or more QRS-complexes deviate too much from the average curve, even after the adjustments are made. As a result, the average curve is unable to fit the signal and errors occur in the residual (Fig. 3). For this reason, we conclude that the average beat method is unusable for credible autocorrelation measurements.

Excluding QRS-complexes In the methods discussed above, we encounter the problem that we can not (correctly) extract the fluctuations of parts of the signal around the QRS-complex. This inability is mainly caused by the fact that during the QRS-complex, a lot of change in y-axis value takes place in a limited amount of time points. Here, small fitting errors on the time axis can lead to big errors in y-axis value in the residual, and this heavily influences autocorrelation calculations. Therefore, to get more reliable results, it might be preferable to remove the QRS-complexes, or, generally, all parts with high first-order differences, from the signal altogether. This has a clear disadvantage; we discard information by cutting parts of the signal. However, if the autocorrelation in the residual is gradually building up, this effect should still be observable,




Fig. 3. Segment of an ECG signal and the average beat fit. The average beat is not able to fit all QRS-complexes, leading to jumps in the AR(1) measurements similar to Fig. 2A. This is typical for almost all signals in the data sets.

even if the signal is not complete. We implement a sequence of steps to remove unwanted parts of the ECG. This procedure accounts for the fact that, when cutting out parts of the signal, the difference in y-value at the edges might induce sudden “jumps” in the signal which affect any autocorrelation measurements. The cutting process is illustrated in Appendix A.4. Besides the fact that this procedure solves the problems we encountered in the methods mentioned above, it also opens up the possibility to use a frequency filter to extract the residual. As mentioned in the introduction of section 2.2, this filtering technique was not possible before due to the high-frequency characteristic of the QRS-complexes. Now that we are excluding these parts of the signal, we are able to extract the fluctuations using a simple low-pass filter. This is shown in Fig. 4: after removing the QRS-complexes the resulting signal is filtered using a 10 Hz cutoff frequency, which is sufficient to filter out any recurring ECG features. With this method, we are able to extract the residues of almost all signals without the major errors we encountered using other methods. We keep in mind that while the cutting procedure avoids major jumps in the resulting residual for most leads, it might still cause unwanted disruptions for some (for example, very noisy) signals. However, we have no reason to assume that this would result in more positive than negative trends in autocorrelation.

3 Results

We measure the trend of the lag-1 autocorrelation in the residuals calculated using the filtering method in 2.2. The parameter setting can be found in Appendix A.6. The trend is obtained by fitting a linear function to the calculated AR(1)-coefficients using a least-squares method and taking the slope. Evidently, applying this method to signals that exhibit CSD should result in significant positive slopes.

To determine the statistical significance of the slope of the AR(1) coefficients, we generate a distribution of AR(1)-slopes from 1000 surrogate time series. The surrogate time series are created by taking the Fourier transform of the residual,




Fig. 4. ECG signal from Fig. 3 before (top) and after (middle) applying the cutting procedure. The residual (bottom) is calculated by subtracting a filtered version of the signal (dotted line). A low-pass filter with a cutoff frequency of 10 Hz is used. The relative change in width of the plots represents the portion of the signal that is cut. The resulting AR(1) measurements form a smooth curve compared to the results from methods discussed above.

multiplying the computed coefficients by random phases and transforming back. In the transformation, linear properties (amplitudes) are preserved and nonlinear properties (phase angles) are randomized. This way, the power spectrum is preserved and the surrogate time series have the same overall autocorrelation as the original residual but are random otherwise [22]. The AR(1)-slopes of the surrogates are normally distributed. A slope is considered significant if it does not fall within the two-sided 95% confidence interval of the obtained distribution. An example of this significance test is given in Appendix A.5.

Our null hypothesis is a situation where the heart is not close to VF, and no CSD is found in the corresponding ECG data. In that case, given the significance testing method described above, we should find an equal amount of positive and negative significant trends. We represent this by letting significant positive and negative slopes be drawn from a binomial distribution with success probability 0.5. $H_0 : p = 0.5$. We reject this null hypothesis if $p > 0.5$ with a significance level of 5% (i.e. if the probability of $p = 0.5$ is lower than 5%). For cases where the null hypothesis is rejected the alternative hypothesis $H_a : p > 0.5$ is accepted and are considered to have a substantial amount of significant positive trends that may possibly be explained by CSD.

We evaluate the trend of the autocorrelation in the ECG data of four patients that suffered VF. Each data set contains around 1400 leads. In Fig. 5 the slope of the AR(1) coefficients is plotted against the power (root mean square) of the residual. Each scatter plot represents results from a different patient. Significant positive or negative AR(1)-slopes are highlighted in red. The figure shows that

in three of the four cases there is a substantial majority of leads that show a significant positive trend in the lag-1 autocorrelation, compared to the number of significant negative trends. For these cases, H_0 is rejected and CSD might be at play. The corresponding scatter plots seem to have a similar shape, with the number of positive trends increasing as the power of the residual decreases. A possible explanation for this could be that residuals with high power contain a bigger proportion of measurement noise of the ECG recording, distorting the component of the residual that could contain CSD.




Fig. 5. Trend of the lag-1 autocorrelation in ECG residuals of four different VF events, plotted against the power of the residual. The trend is measured by the slope. Each dot represents one lead. Significant slopes are colored red. The table on the right shows the amount of significant positive/negative trends for the corresponding plot. For three out of the four patients H_0 is rejected (red cells) (Note that some points may be out of bounds for the sake of better visualization)

Before we draw the conclusion that the number of significant positive slopes is in fact the result of a CSD effect we examine sets of test data that consist of ECG signals from hearts that are not close to a VF event. Because we do not have access to similar heart surface electrograms for this category we use open source ECG data from PhysioNet [23] for this purpose. We analyze 9 sets of test data consisting of 10-second samples from 24-hour, 250 Hz Holter ECG recordings. As expected, we find no substantial majority of significant positive trends in any of these test sets. The results of the analysis of the test data are shown in Appendix A.7.

We also perform our analysis on the data set containing the original body surface ECGs that are used to compute the inverse solution. Strikingly, we do not observe the same substantial majority of positive trends we see for three of

the four data sets of estimated heart surface ECGs. If our measurements are in fact the result of CSD, it seems this is not measurable with the original data and that the transformation to the inverse solution does provide extra information necessary to observe this effect.

An overview of all AR(1)-trend measurements is given in Appendix A.11. Given our significance testing method, under the null hypothesis (no CSD) we expect that around 5% of the autocorrelation trends we observe will be significant. Remarkably, for most data sets where H_0 is valid, we find that more than 5% of the trends are significant. It is likely that a portion of these significant trends is the result of residuals that are corrupted by noise, either directly (the noise directly influences the autocorrelation), or indirectly (the noise forces errors in the extraction of the residual, which in turn influences the autocorrelation). We can see that significant trends are relatively common for residuals with high power. These results indicate that, while artifacts can cause more significant trends than expected, they occur both in positive or negative form (see Appendix A.7), and are therefore not likely to lead to a rejection of H_0 .

The sets of heart surface ECGs include triangulation coordinates, mapping each ECG signal to a point in 3D-space that corresponds to the location around the heart for which the inverse solution is calculated. We use this information to reproduce the plots from Fig. 5 where every point is color-coded based on its coordinates. These plots are shown in Appendix A.9. The results show that the points are clustered, which means that the significant positive trends in the heart surface data (red in Fig. 5) can be measured from specific angles, rather than all around the heart surface. We do not have enough information to couple the 3D-coordinates to a physical location of the heart surface; this would not be difficult to realize in future data acquisition. If the substantial amounts of significant trends are the result of CSD this can be valuable information, since one would know exactly where to look for possible early warning signals.

If, in further research, we can prove the presence of CSD in heart surface electrograms of VF victims, it may be possible for implantable devices such as ICDs to detect this effect and provide an early warning signal for an oncoming arrhythmia. If, additionally, we can isolate from which area around the heart surface it can be measured, this may even be done by using just a single lead rather than the full potential map of the heart surface we used for this experiment.

4 Conclusion and Discussion

CSD has been used as a generic early warning signal for critical transition in a wide range of systems ranging from finance to climate. We reason that the heart as a complex system may bear similarities to such systems, since, in the context of dynamical systems theory, the transition from a normal heart rhythm to VF can be understood as a shift between two states with a different attractor. We hypothesized that when the heart is close to VF (i.e., close to the basin of attraction of the chaotic attractor) it may show decreasing resilience to pertur-

bations, which can be measured as CSD. To test this hypothesis we investigate heart surface ECG signals right before the onset of VF.

In our results we indeed find signs of CSD: for three out of four VF victims, we find a substantial majority of significant positive autocorrelation trends in the residuals of heart surface ECG signals compared to the amount of significant negative trends. The heart surface ECGs are estimated by solving an inverse problem using body surface ECGs. If we perform the same analysis on the original body surface data we do not find such a majority, suggesting that, if CSD is in fact present, we have to compute the inverse solution to observe this effect. Furthermore, triangulation coordinates of the heart surface ECGs suggest that the possible CSD effect can only be measured from specific, yet unspecified angles surrounding the heart. We compare the results of the heart surface ECGs of VF victims to results from Holter recordings of subjects that are not close to a VF event. For the latter (no VF) data, we find that none of the nine recordings we analyze contain a substantial amount of significant positive autocorrelation trends and thus exhibit no CSD.

It has become clear, however, that this test data does not serve as an appropriate comparison the VF data for a number of reasons. Firstly, the test signals have a lower sampling rate of 250 Hz compared to the 1000 Hz of the VF data. It is possible that fluctuations in the residuals that show CSD can be captured by a sampling rate of 1000 Hz but not by a sampling rate of 250 Hz. If this is in fact the case, using a sampling rate 250 Hz can not prove the absence of CSD in the test data. Secondly, with Holter recordings, the electrodes are placed on the chest of the patient. If we assume that the substantial amount of significant positive trends in the heart surface data is caused by CSD, our analysis of the data used to compute the inverse solution already seems to indicate that this effect is not as easily measured on the body surface. It would therefore only be logical if the test also does not show CSD. Lastly, the triangulation coordinates of the heart surface ECGs suggest that, if the data shows CSD, it can probably only be measured from specific angles surrounding the heart surface. Since the test data consists Holter recordings with only one lead, and thus only one angle of incidence, it is already unlikely to measure CSD for these signals.

The test data we use serves more as a validation of our method, rather than a validation of our result. While we do not expect CSD, the test signals are prone to the same types of noise as the VF data, which could influence autocorrelation measurements. Our analysis has shown, however, that this does not lead to a substantial majority of significant positive trends in any of the test data sets. This could indicate that the signs of CSD we find for the VF data are no artifacts of measurement noise. On the other hand, we can argue that some types of measurement noise (for example, noise caused by the movement of a patient) could affect multiple leads at once for the heart surface data since each lead covers that same ten-second time span. For this reason, the test data is still not sufficient to rule out the possibility that measurement noise influences our results.

To prove that the predominant amount of significant positive autocorrelation trends we find in the residuals of heart surface ECGs of patients that suffered VF is in fact the result of CSD, we have to directly compare it to data of subjects that did not suffer VF, recorded in a similar manner. We therefore advocate further data acquisition. An important first step would be to obtain the 1400-lead heart surface ECG of subjects that did not suffer VF by solving the inverse problem using the body surface potentials. If the analysis of such data does not result in a majority of significant positive trends, this would be a strong indication that our measurements are showing CSD. Additional data from patients that suffered VF would also enable better statistical grounding. For future ECG recordings it is important to be able to map each signal to an exact physical location so that, if we can indeed prove CSD, we can also pinpoint angles from which this effect is measurable. So far we have checked for CSD by looking for a buildup in autocorrelation in the 10 seconds prior to the VF event. However, it is possible that this buildup initiates at an earlier time. If we can indeed use CSD as an early warning signal in this setting, it would be valuable to detect this as early as possible. For future data sets, it might therefore be useful to record ECG even longer before the onset of VF, where possible.

We concluded that for the extraction of the residuals from the signals it is best practice to remove parts of the signal with high first-order differences, which are otherwise hard to filter out. Evidently, for future research it is desirable to develop a method that can correctly extract residuals for the full signal, for example by improving the average beat method proposed in 2.2 or developing more advanced modeling techniques than the Gaussian fitting method in 2.2.

Acknowledgements This work was partly supported by the Fondation Leducq Transatlantic Network of Excellence 16CVD02. RQ thanks EU Horizon 2020 project TO_AITON (848146) for support.

A Appendix

The full version of the paper including the appendix is available online at: <https://tinyurl.com/ukjfwuc>

References

1. Amann, A., Tratnig, K., Unterkofler, K.: Reliability of old and new ventricular fibrillation detection algorithms for automated external defibrillators. *Biomedizinische Technik/Biomedical Engineering*. 48, 216-217 (2003).
2. Ahmad M., Bloomstein L.Z., Roelke M., Bernstein A.D., Parsonnet V.: Patients' attitudes toward implanted defibrillator shocks. *Pacing and clinical electrophysiology* **23**(6), 934-938 (2000)
3. Sears S.F., Conti J.B.: Quality of life and psychological functioning of ICD patients. *Heart* **87**(5), 488-493 (2002)

4. Lane, R.: Prediction and prevention of sudden cardiac death in heart failure. *Heart*. **91**(5), 674-680 (2005).
5. Reed, M., Robertson, C., Addison, P.: Heart rate variability measurements and the prediction of ventricular arrhythmias. *QJM*. **98**(2), 87-95 (2005).
6. Scheffer, M., Bascompte, J., Brock, W., Brovkin, V., Carpenter, S., Dakos, V., Held, H., van Nes, E., Rietkerk, M., Sugihara, G.: Early-warning signals for critical transitions. *Nature*. **461**(7260), 53-59 (2009).
7. Wissel, C.: A universal law of the characteristic return time near thresholds. *Oecologia*. **65**(1), 101-107 (1984).
8. Scheffer, M., Dakos, V., Nes, E., Brovkin, V., Petoukhov, V., Held, H.: Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences*. **105**(38), 14308-14312 (2008).
9. McSharry, P., Smith, L., Tarassenko, L.: Prediction of epileptic seizures: are non-linear methods relevant?. *Nature Medicine*. **9**(3), 241-242 (2003).
10. Goldberger, A., Bhargava, V., West, B., Mandell, A.: Some observations on the question: Is ventricular fibrillation "chaos"?. *Physica D: Nonlinear Phenomena*. **19**(2), 282-289 (1986).
11. Surawicz, B.: Ventricular fibrillation. *The American Journal of Cardiology*. **28**(3), 268-287 (1971).
12. Kaplan, D., Cohen, R.: Is fibrillation chaos?. *Circulation Research*. **67**(4), 886-892 (1990).
13. Garfinkel, A., Chen, P., Walter, D., Karagueuzian, H., Kogan, B., Evans, S., Karpoukhin, M., Hwang, C., Uchida, T., Gotoh, M., Nwasokwa, O., Sager, P., Weiss, J.: Quasiperiodicity and chaos in cardiac fibrillation. *Journal of Clinical Investigation*. **99**(2), 305-314 (1997).
14. Wit, A., Cranefield, P.: Reentrant excitation as a cause of cardiac arrhythmias. *American Journal of Physiology-Heart and Circulatory Physiology*. **235**(1), H1-H17 (1978).
15. Weiss, J., Qu, Z., Chen, P., Lin, S., Karagueuzian, H., Hayashi, H., Garfinkel, A., Karma, A.: The Dynamics of Cardiac Fibrillation. *Circulation*. **112**(8), 1232-1240 (2005).
16. Weiss, J., Garfinkel, A., Karagueuzian, H., Qu, Z., Chen, P.: Chaos and the Transition to Ventricular Fibrillation. *Circulation*. **99**(21), 2819-2826 (1999).
17. Qu, Z.: Chaos in the genesis and maintenance of cardiac arrhythmias. *Progress in Biophysics and Molecular Biology*. **105**(3), 247-257 (2011).
18. Qu, Z., Hu, G., Garfinkel, A., Weiss, J.: Nonlinear and stochastic dynamics in the heart. *Physics Reports*. **543**(2), 61-162 (2014).
19. Gulrajani, R.: The forward and inverse problems of electrocardiography. *IEEE Engineering in Medicine and Biology Magazine*. **17**(5), 84-101, 122 (1998).
20. Clifford, G., McSharry, P.: Method to filter ecgs and evaluate clinical parameter distortion using realistic ECG model parameter fitting. *Computers in Cardiology*. pp. 715-718. IEEE, Lyon (2005).
21. Suppappola, S., Sun, Y., Chiaramida, S.: Gaussian pulse decomposition: An intuitive model of electrocardiogram waveforms. *Annals of Biomedical Engineering*. **25**(2), 252-260 (1997).
22. Schreiber, T., Schmitz, A.: Surrogate time series. *Physica D: Nonlinear Phenomena*. **142**(3-4), 346-382 (2000).
23. Goldberger, A., Amaral, L., Glass, L., Hausdorff, J., Ivanov, P., Mark, R., Mietus, J., Moody, G., Peng, C., Stanley, H.: PhysioBank, PhysioToolkit, and PhysioNet. *Circulation*. **101**(23) (2000).