MOX–Report No. 20/2012

Statistical models for detecting Atrial Fibrillation events

Ieva, F.; Paganoni, A. M.; Zanini, P.

MOX, Dipartimento di Matematica "F. Brioschi"
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox@mate.polimi.it http://mox.polimi.it

Statistical models for detecting Atrial Fibrillation events

Francesca Ieva[1], Anna Maria Paganoni[1], Paolo Zanini[1]

April 20, 2012

[1] MOX– Modellistica e Calcolo Scientifico
Dipartimento di Matematica "F. Brioschi"
Politecnico di Milano
via Bonardi 9, 20133 Milano, Italy
francesca.ieva@mail.polimi.it
anna.paganoni@polimi.it
paolo.zanini@mail.polimi.it

Abstract

Atrial Fibrillation is the most common cardiac arrhythmia that naturally tends to become a chronic condition and chronic Atrial Fibrillation leads to an increase in the risk of death. The study of time series of time intervals between an R peak in the electrocardiogram and the following one is an effective way to investigate the presence of Atrial Fibrillation and to detect when a single event starts and ends. This work presents a new statistical method to deal with identification of Atrial Fibrillation events. Some simulations in order to assess the performances of the proposed method are detailed and the results obtained applying this method to real data concerning patients affected by Atrial Fibrillation are discussed.
Keywords: Atrial Fibrillation, RR intervals, Time Series Analysis, Ljung-Box statistic.
AMS Subject Classification: 62P10, 62M10

1 Introduction

Atrial Fibrillation (AF) is the most common cardiac arrhythmia and involves the two upper chambers (atria) of the heart. During AF, the normal electrical impulses generated by sinoatrial node are overwhelmed by disorganized electrical impulses that originate in the atria and pulmonary veins, leading to conduction of irregular impulses to the ventricles that generate the heartbeat. The result is an irregular heartbeat, which may occur in episodes lasting from minutes to weeks, or it could occur all the time for years. The natural tendency of AF is to become a chronic condition and

chronic AF leads to an increase in the risk of death.
The main device used in order to investigate the heartbeat is the Electrocardiogram (ECG). The ECG is a diagnostic tool that measures and records the electrical activity of the heart in exquisite details. The interpretation of these details allows for diagnosis of a wide range of heart diseases. The stylized shape of an ECG is depicted in Figure 1 (upper panel) where atrial contraction shows up as the P wave; ventricular contraction is identified as a series of three waves, Q,R and S, known as the QRS complex. The third wave in an ECG is the T wave which reflects the electrical activity produced when the ventricles recharge for the next contraction (repolarization); for further inquiry about clinical details on ECG see [7].

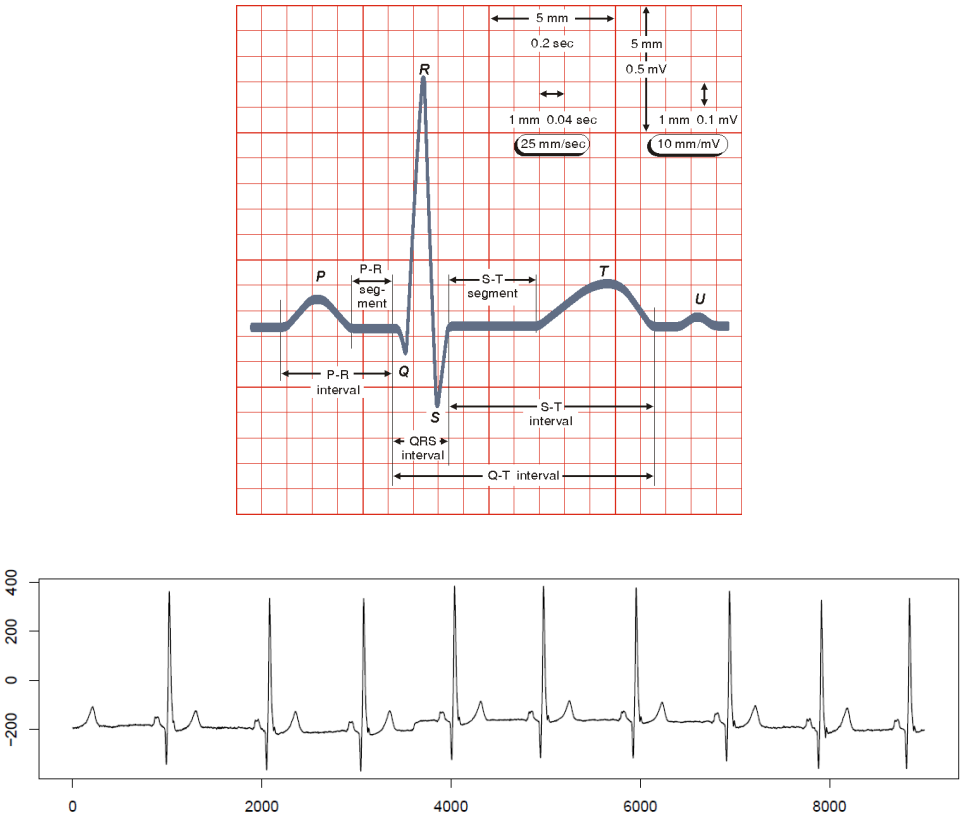

Figure 1: Upper panel: stylized shape of a physiological single beat, recorded on ECG graph paper. Main relevant points, segments and waves are highlighted. Bottom panel: some cardiac cycles recorded in one of the leads of the electrocardiogram with sampling frequency of 1000 Hz.

As we said before Atrial Fibrillation can be diagnosed by the ECG. Characteristic findings are the absence of P waves, with unorganized electrical activity in their place, and irregular RR intervals (the time between an R

peak and the following one, see the bottom panel of Figure 1) due to irregular conduction of impulses to the ventricles.
The study of RR intervals is an effective way to investigate the presence of AF and to detect when a single event starts and ends. Several examples exist in literature (see [1], [6],[9]), which are focused on the peculiar variance of RR intervals during the AF process, which is much greater with respect to the variance during the physiological heartbeat.
Anyway, in many situation, an AF event does not follow a physiological time slot, but comes after other types of arrythmia. On the same time, in many cases, the irregular heartbeat does not disappear when the event finishes. According to these problems, analyzing the RR intervals, it may be possible to look at an irregular heartbeat even when the AF event itself is not already started or already finished. So a method based on detection of changes in the variance of the process can result inaccurate and can fail in those cases where AF episode has the characteristics described above.
Then methods which are not based on the analysis of the process variance are needed, in order to identify suitable quantities to characterize the different phases, 'pre AF', 'AF' and 'post AF'.
The main purpose of this work is to identify such quantities by means of suitable statistical tools. To achieve this goal, we analysed some data of 8 patients affected by AF. In particular we had for each patients the RR intervals sequence from two hours before to two hours after an event of AF. The idea is to consider the process of RR intervals as a time series (see [4], [5]), identifying specific parameters which enable us to detect when an event of AF starts or ends and then pointing out suitable test statistics.
The paper is then organized as follows: in Section 2 we present some elements of time series processes theory, focusing especially on ARIMA models. In Section 3 we illustrate the new method we propose to deal with identification of AF events. In Section 4 we show some simulations in order to assess the performances of the proposed method. To conclude, in Section 5, we present the results obtained applying our method to real data consisting of RR intervals of patients affected by AF.
All the simulations and the analyses of real data are carried out using `R` statistical software [11].

2 AutoRegressive Integrated Moving Average (ARIMA) models

In this section we present some elements of time series processes theory, focusing especially on ARIMA models for linear stationary and non-stationary processes.

2.1 Linear stationary processes

A time series is a set of observations generated sequentially in time. When N successive values of such a series are available, they will be indicated with $z_1, z_2, ..., z_N$. Then we can regard z_t as the observation at time t. In the following we will consider a discrete time series arising by sampling a

stochastic process representing a phenomenon. A statistical phenomenon that evolves in time according to probabilistic laws is called a stochastic process. The time series to be analysed may be seen as a particular realization of the system under study, produced by such underlying process. In other words, in analyzing a time series $z_1, ..., z_N$ we regard it as a realization of a stochastic process $Z_1, ..., Z_N$.
A particular class of stochastic processes, called stationary processes, is based on the assumption that the process is in a particular state of statistical equilibrium. A stochastic process is said to be strictly stationary if its properties are unaffected by a change of time origin; that is, if the joint probability distribution associated with m observations $z_{t_1}, ..., z_{t_m}$, made at any set of times $t_1, ..., t_m$, is the same as that associated with m observations $z_{t_1+k}, ..., z_{t_m+k}$, made at times $t_1 + k, ..., t_m + k$, for any lag k. From now on we denote both the process and the observations as z_t, according to the literature of time series analysis.
In general, analyzing stationary process, there are three kind of linear models who are mainly used. The first is the Autoregressive Process of order p (briefly AR(p)) which may be written as

$$z_t = \phi_1 z_{t-1} + ... + \phi_p z_{t-p} + a_t \tag{1}$$

where the white noise a_t may be regarded as a series of shocks which drive the system and it is such that $E[a_t] = 0$, $Var[a_t] = \sigma_a^2$ and $Cov(a_{t-i}, a_{t-j}) = 0$ for $i \neq j$ and $(\phi_1, ..., \phi_p) \in \mathbb{R}^p$ are the unknown parameters. Introducing the *backward shift operator* B, defined as

$$Bz_t = z_{t-1} \qquad B^j z_t = z_{t-j}$$

the model (1) may be written as

$$(1 - \phi_1 B - ... - \phi_p B^p) z_t = a_t$$

or shortly

$$\phi(B) z_t = a_t \tag{2}$$

where $\phi(B) = 1 - \sum_{j=1}^{p} \phi_j B^j$.
Another kind of model is called Moving Average process of order q (briefly MA(q)). This process may be written as

$$z_t = a_t - \theta_1 a_{t-1} - ... - \theta_q a_{t-q}$$

$$z_t = (1 - \theta_1 B - ... - \theta_q B^q) a_t$$

or shortly

$$z_t = \theta(B) a_t \tag{3}$$

where $\theta(B) = 1 - \sum_{j=1}^{q} \theta_j B^j$ and $(\theta_1, ..., \theta_q) \in \mathbb{R}^q$ are the unknown parameters.
These two models may be combined, obtaining an ARMA(p, q) model

$$z_t = \phi_1 z_{t-1} + ... + \phi_p z_{t-p} + a_t - \theta_1 a_{t-1} - ... - \theta_q a_{t-q}$$

$$\phi(B) z_t = \theta(B) a_t \tag{4}$$

Each of the model previously illustrated has some peculiar features. In particular we focus on those expressed by autocorrelation and partial autocorrelation function, that play a crucial role in the choice of the optimal order (p, q) of the model. It may be shown that, according to the fashion of autocorrelation and partial autocorrelation function, autoregressive and moving average process have different behaviors. In fact, it may be proved that the autocorrelation function of an AR(p) has infinite nonzero elements and consists of a mixture of damped exponentials and damped sine waves. On the other hand the autocorrelation function of a MA(q) process is zero beyond the order q. In other words, the autocorrelation function of a Moving Average process has a cut-off at lag q.
The partial autocorrelation function of an AR(p) process, instead, is nonzero only for the first p elements. In other words, the partial autocorrelation function of an AR(p) process has a cut-off at lag p. Oppositely, the partial autocorrelation function of a MA(q) process is infinite in extent and is dominated by damped exponentials and/or damped sine waves.
For an ARMA(p, q), instead, both autocorrelation function and partial autocorrelation function are infinite in extent. Further details on optimal model order detection using autocorrelation and partial AC functions can be found in [2].

2.2 Linear non-stationary processes

Many empirical time series have no constant mean. Even so, they exhibit a sort of homogeneity in the sense that, apart from local level, or perhaps local level and trend, one part of the series behaves much like any other part. Models which describe such homogeneous non stationary behavior can be obtained by supposing some suitable difference of the process to be stationary. We may consider the properties of the important class of models for which the d-th difference ($\nabla^d z_t = z_t - z_{t-d}$) is a stationary ARMA process.
Then, let us consider the model

$$\phi(B)\nabla^d z_t = \theta(B)a_t \tag{5}$$

We call the process (5) an AutoRegressive Integrated Moving Average (ARIMA) process. If the autoregressive operator $\phi(B)$ in (5) is of order p and the moving average operator $\theta(B)$ is of order q, then we say we have an ARIMA(p, d, q) process.

2.3 Model diagnostic checking

Let consider the following ARIMA model

$$\phi(B)\nabla^d z_t = \theta(B)a_t \tag{6}$$

and suppose to fit it obtaining ML estimates $(\widehat{\boldsymbol{\phi}},\widehat{\boldsymbol{\theta}})$ for the parameters. Then we shall refer to the quantities

$$\widehat{a}_t = \widehat{\theta}^{-1}(B)\widehat{\phi}(B)\nabla^d z_t \tag{7}$$

as the residuals. As the length increases, the $\widehat{a}_t$ become closer to the white noise a_t.
Now suppose the form of the model were correct and that we knew the true parameter values $\boldsymbol{\phi}$ and $\boldsymbol{\theta}$. Then the estimated autocorrelation $r_k(a)$ of the process a would be distributed approximately normally with zero mean (see [3]). Now, in practice, the time value of the parameters is unknown and only the estimates $(\widehat{\boldsymbol{\phi}},\widehat{\boldsymbol{\theta}})$ are available for calculating the $\widehat{a}$. Then, autocorrelation $r_k(\widehat{a})$ of the $\widehat{a}$ can yield valuable evidence concerning the lack of fit. An interesting way to analyze the goodness of fit of the model is then to consider the $r_k(\widehat{a})$ not individually, but taken as a whole. Let suppose that we have the first K autocorrelations $r_k(\widehat{a})$ $(k = 1, 2, ..., K)$ from any ARIMA(p, d, q) process. Then it is possible to show (see [8]) that, if the fitted model is appropriate, the statistic

$$Q = \widetilde{n}(\widetilde{n} + 2) \sum_{k=1}^{K} r_k^2(\widehat{a})/(\widetilde{n} - k) \tag{8}$$

is approximately distributed as $\chi^2(K - p - q)$, where $\widetilde{n} = n - d$, with n equal to the number of observations. Therefore, an approximate test of the hypothesis of model adequacy may be performed. The statistic Q is called Ljung-Box statistic.

3 A method to detect changes in a time series

We now consider a situation where a phenomenon evolves according an ARIMA process. We wish to analyse a time series and to detect when such a phenomenon starts and/or ends. If this specific phenomenon is characterized by an higher (or lower) variability with respect to the current situation, then there is a huge number of methods known to be effective in detecting these changes in variability. Examples are control charts (see [10]) and methods based on graphical analysis among others (see [9]). However, there are a lot of situations where a phenomenon is not characterized by a modification in the variability, but by a change in the process that generates the observations. In these cases methods such those mentioned above are useless. We wish to present here an ad hoc method for dealing with such situations.
In particular the main goal is to identify the beginning and the end of a specific phenomenon generated by an ARIMA process. This means to identify the model parameters, i.e., values of d, p and q. As we have previously presented, in the case of a stationary model the autocorrelation and partial autocorrelation function will quickly die out. Knowing that the estimated autocorrelation function tends to follow the behavior of the theoretical autocorrelation function, failure of this estimated function to die out rapidly might logically suggest that we should treat the underlying stochastic process as non-stationary in z_t, but possibly as stationary in $\nabla^d z_t$. Once identified one or more possible values for d, we move to the choice of p and q. This may be done considering the specific behaviors of the autocorrelation and partial autocorrelation functions and corresponding cut-off.

Then to identify the starting and ending times of the phenomenon of interest we propose the following procedure. Consider the first N observations (with N much smaller then the number n of observations) and fit the identified model over this sub-sample. Then the p-value of the Ljung-Box test (choosing a value for K) is recorded. These operations have to be repeated over the sub-sample from the second to the $N+1$-th observation. Once reached the last observation, the procedure ends producing a "time series" of the p-values which may be used to detect the beginning and the end of the phenomenon of interest.
The purpose now is to construct a test for checking the null hypothesis that the phenomenon is present against the alternative hypothesis that the phenomenon is absent. This may be written in a more rigorous way as follow:

$$H_0 : p = \overline{p} \wedge d = \overline{d} \wedge q = \overline{q} \text{ vs. } H_1 : p \neq \overline{p} \vee d \neq \overline{d} \vee q \neq \overline{q} \tag{9}$$

In order to build the critical region for the test (9) the first M p-values can be considered and the rejection region can be constructed according to the following decisional criterion: the null hypothesis is rejected if at least 1 of the M p-values considered are less than a certain number, say y. In this way we are able to evaluate an approximate level α thanks to a Bonferroni inequality. More specifically:

$$\alpha = P_{H_0}(\bigcup_{i=1}^{M}(p_i < y)) \leq \sum_{i=1}^{M} P_{H_0}(p_i < y) \leq \sum_{i=1}^{M} y = My \tag{10}$$

where p_i is the i-th p-value and the property of valid p-value holds. If we call α^* the value of α established, we set $y = \alpha^*/M$.
The method to detect start and/or end of a specific phenomenon evolves according the following steps:

1. implement the test in (9)-(10) over the first M p-values. At the $N+M$-1-th observation, the output is setted at 0 if there is statistical evidence to reject the null hypothesis, while is setted at 1 otherwise;
2. repeat the step 1 after a shift of one observation until the last one is reached.

Once the procedure ends, an output of 0's and 1's is available. 1 indicates the presence of the phenomenon, 0 the absence. Starting and end points can be then detected through this last 0/1 time series.

4 Simulations

In order to validate the proposed method, different situations have been tested and analysed, with the following aims:

- to point out settings where our method performs at best
- to assess the robustness of the method varying α^* and N
- to make a sensitivity analysis over the parameter K of the Ljung-Box statistics

The method presented in this paper is a technique to detect modification in the process underlying the observed phenomenon. We chose an ARIMA (0,1,1) as reference process (RP), considering a sequence of 7000 realizations from a process, say P_{pre}, then 40000 realizations from the reference model and finally 7000 realizations from another different process, say P_{post}. For all the simulations the value of M was fixed equal to 100.
In order to investigate the performances of the proposed method, we tested it in different situations: in the first, second and third simulation, P_{pre} and P_{post} are very different from RP, whereas in the fourth are not. In particular we set $P_{pre} \equiv P_{post}$, and we considered an ARIMA (4,1,2), ARIMA (5,1,3) ARIMA (2,2,0) and ARIMA (1,1,1) respectively. For all these settings, we assumed $K = 5$, $N = 600$ and $\alpha^* = 0.01$.
Figures 2a, 2b and 2c show that our method works very well in the first 3 settings, where it is appreciable the correspondence among the real starting end end points (red lines) and the 0/1 sequence.
In the fourth simulation, instead, the method is not able to catch the phenomenon under study, as it is shown in Figure 2d.
This behavior may be explained by the fact that this method works very well when the modification of the process underlying the generation of observations is considerable. On the other hand, when the modification is not so significant, the method loses efficiency.

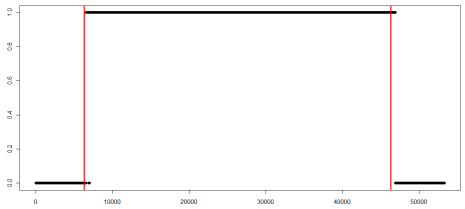

(a) Output of the method: before and after the phenomenon under study the process is an ARIMA(4,1,2).

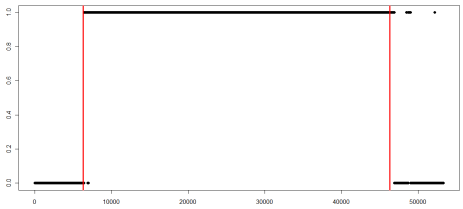

(b) Output of the method: before and after the phenomenon under study the process is an ARIMA(5,1,3).

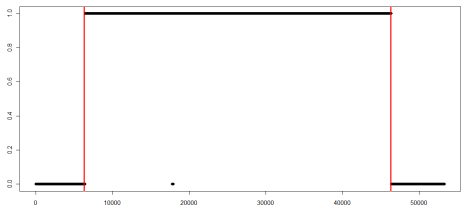

(c) Output of the method: before and after the phenomenon under study the process is an ARIMA(2,2,0).

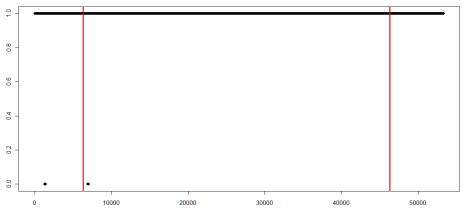

(d) Output of the method: before and after the phenomenon under study the process is an ARIMA(1,1,1).

Figure 2: Analysis of the output of the method changing the process underlying the observations before and after the phenomenon. Red lines represent the start and the end of the phenomenon.

In the following we focus on the first situation, where the generating process is an ARIMA (0,1,1), anticipated and followed by a process of observations generated from an ARIMA (5,1,3), because cases (2a) and (2c) give similar results. We analysed how the power of the test in (9)-(10) is affected by α^* and N. For this analysis we considered $K = 5$.
If α^* was the real probability of the I type error, the power would increase as α^* grows. We do not have the real probability of the I type error, but only an upper estimate. In spite of this fact, we would observe the power growing up as long as α^* increases. Another parameter that affects the power of the test is N. Again, the bigger is N, the greater the power of Ljung-Box test. So also the power of the global test should raise.
In Figure 3 the output of the method varying α^* (along the rows) and N (along the columns) is shown. It can be inferred that the behavior of the method is consistent, since the number of errors before and after the phenomenon decreases as long as α^* and N increase, as we expected.

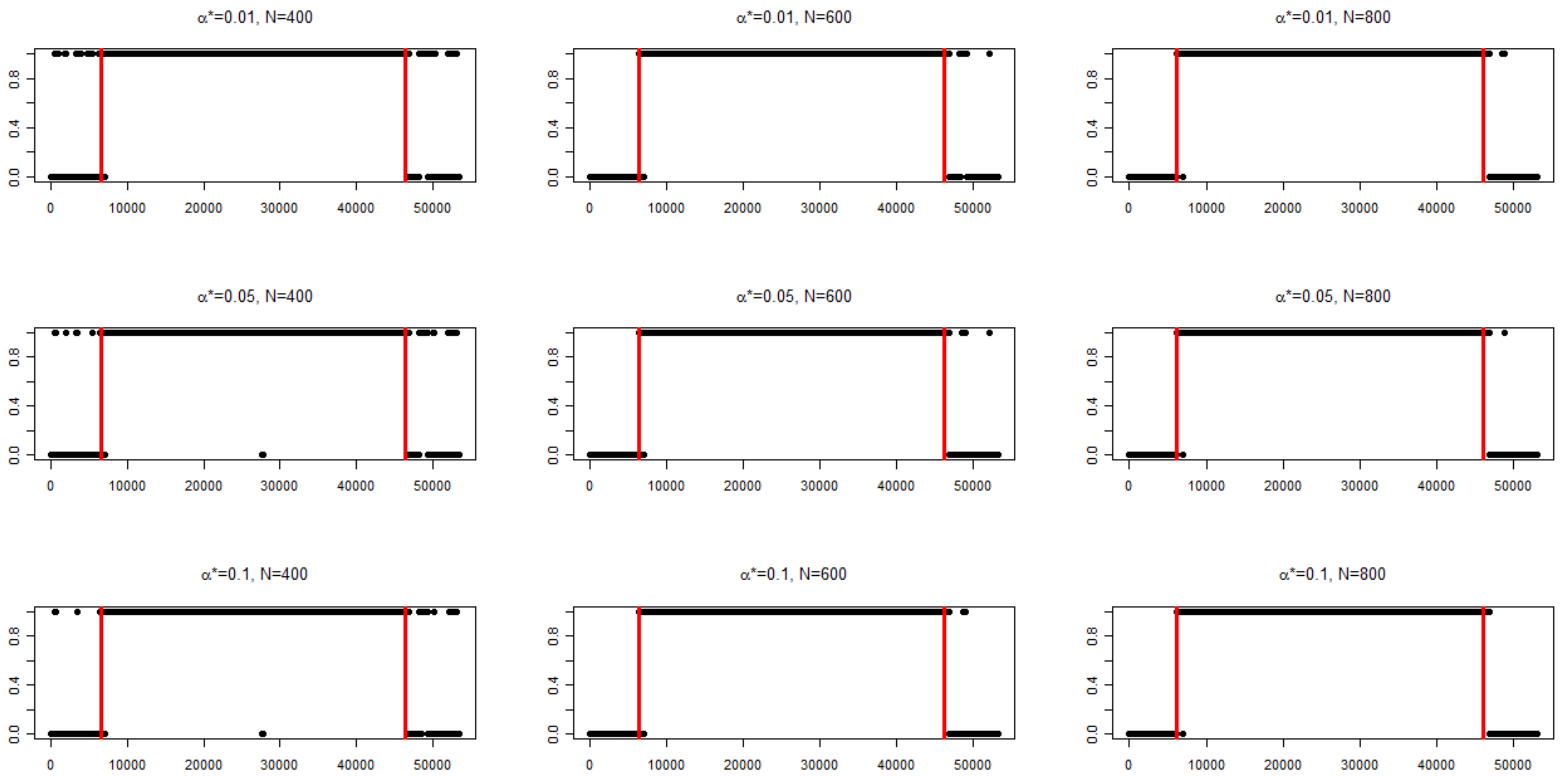

Figure 3: Output of the method varying α^* (along the rows) and N (along the columns).

In Table 1 we show the average number of observations of II type error for 40 simulations (we consider $\alpha^* = 0.05$). The obtained results suggest that it is possible to increase the power of the test tuning N in a suitable way. Then, one could think to set a very high value of N in order to obtain a satisfactory power. However, this is not costless. In fact, increasing values of N reflects on delay in starting and end points detection.

N	*Average number of observations of II type error*	*Standard deviation of observations of II type error*
400	4999	1157
600	2283	1030
800	883	673

Table 1: Average number of observations of II type error varying N.

Hence the choice of the parameter N is regulated by a trade-off between the desired power of the test and the delay in the detection of the phenomenon.

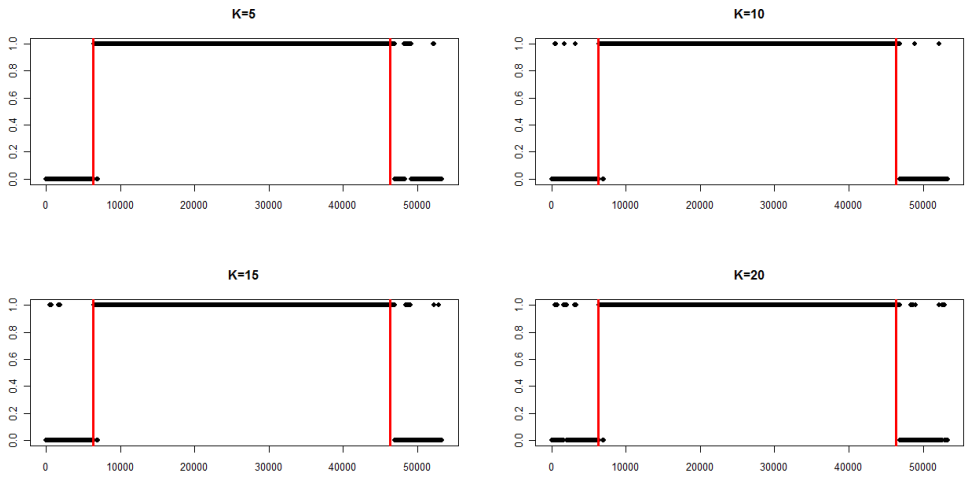

Figure 4: Output of the method varying K.

To conclude the simulations's analysis we would like to infer about the parameter K of the Ljung-Box statistics in order to understand if the method is affected by a modification of its value. Let consider the situation where observations before and after the phenomenon were generated by an ARIMA(5,1,3), and fix $\alpha^* = 0.01$ and $N = 600$. In Figure 4 the output of the method for different values of K (5,10,15 and 20) is shown. Although the outputs are different, it does not appear any pattern of dependence on K. Therefore we can observe that the dependency of the method from K is feeble.

5 An application to Atrial Fibrillation

Let consider now an application of the method proposed in this paper to real data. Specifically we analysed RR intervals of 8 patients during Atrial Fibrillation (AF).
The data available are the RR intervals of such patients from two hours before to two hours after an event of AF. The duration of the phenomenon is different between patients and it is displayed in Table 2.

Pat. num.	**Duration AF (min.)**	**Number of observations**
1	521	41085
2	613	43178
3	433	52937
4	13	1066
5	56	4326
6	442	52661
7	319	28229
8	229	17989

Table 2: Duration and number of observations of the event of AF.

The main goal is to detect the event of AF starting from the time series of RR intervals.
In some cases, the variability of RR intervals during AF is very high with respect to the physiological heartbeat. However this remarkable change in the variability of the phenomenon could be absent, so the traditional methods based on the variability are inefficient in detecting AF starting point.
The first step consists in the identification of a model for the RR intervals during AF. According to the rule presented in Section 3, we used the autocorrelation and partial autocorrelation functions to determine a suitable model. As it is shown in Figure 5, the autocorrelation function of ∇z_t is truncated after the lag number one, while that of $\nabla^2 z_t$ is zero after the lag two. This behaviour, as it has been presented in Section 2, is typical of an ARIMA (0,1,1) and (0,2,2). We then set RP$\equiv$ARIMA (0,1,1). Once identified a model to describe the RR intervals during AF, we would like to analyse the performances of the method in detecting start and end of such a phenomenon.

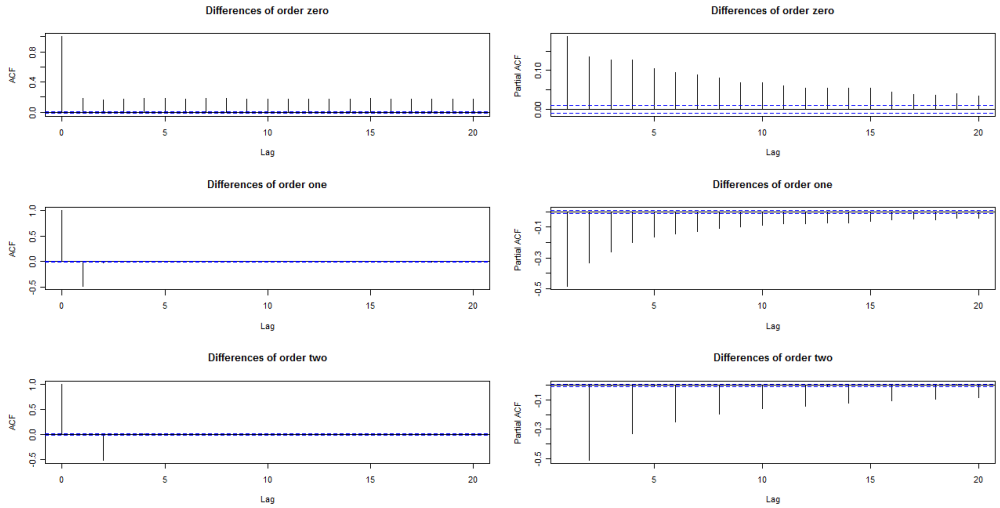

Figure 5: Patient 1: autocorrelation (left panels) and partial autocorrelation (right panels) functions for the time series of RR intervals, of the differences of order one and of the differences of order two.

In order to achieve this goal, let fix the parameters of $K = 5$, $\alpha^* = 0.01$ and $M = 100$, and analyse the output of the proposed method varying N, that in Section 4 has been noted to be the most important parameter that affects the performance of our method. Figure 6 shows the output of the method applied to patients 1 and 5, varying N. We present here only the output for two patients, because the results for the other patients are quite similar.

Some practical considerations can be extrapolated observing the Figure 6. First of all we may point out to the behavior of the method as long as N increases and how this behavior relates to the corresponding simulation case. Then we may analyse the delay in the detection of start and end of AF and the number of errors.

Dealing with the delay, since each observation is the time between an R peak and the following one, we can evaluate the time of the delay in the detection of the event of AF and not only the number of observations. As it is shown in Table 3, the delay in detecting the phenomenon is negligible if compared with the duration of AF. Moreover, in some cases the method is able to detect in advance the event of AF.

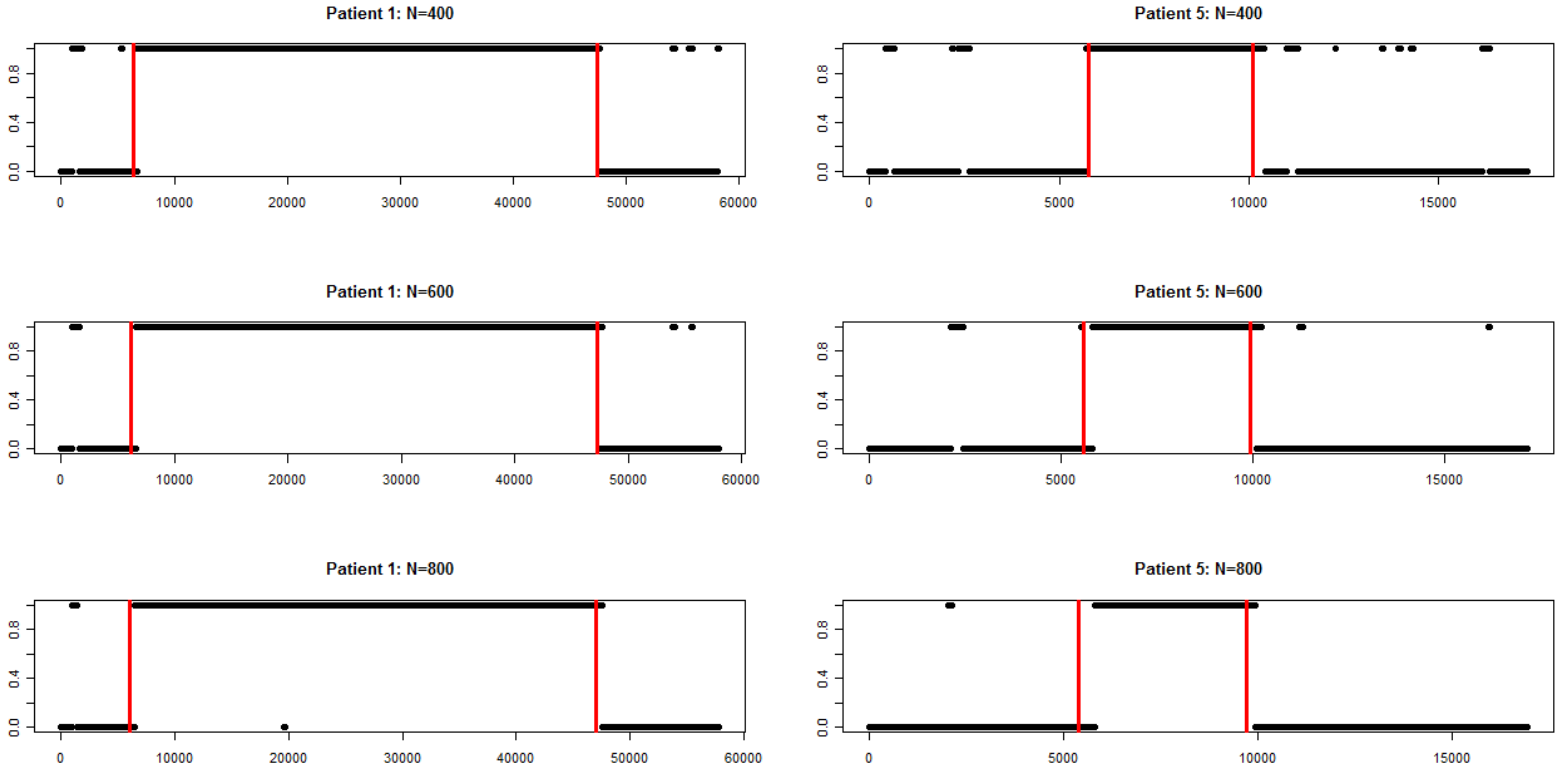

Figure 6: Output of the method for the patients 1 (left panels) and 5 (right panels) varying N.

Another important point which we should focus on is the number of errors made by the method. From a first qualitative analysis, from Figure 6, we can observe that the most part of errors seems to involve a few number of consecutive observations.

Pat. num.	**N=400** (min.)	**N=600** (min.)	**N=800** (min.)
1	4.4	5	5.4
2	4.2	6	6.6
3	-2.4	0.3	-4.6
4	4	5.9	8.4
5	-1.5	2.7	5.5
6	-2.2	1.1	1.1
7	16.8	16.9	29.2
8	4.8	5.6	7.7

(a) Delays detecting the start of AF.

Pat. num.	**N=400** (min.)	**N=600** (min.)	**N=800** (min.)
1	2.6	4.5	5.1
2	5.5	8.3	11.9
3	6.7	7.2	7.7
4	7.3	9.7	9.6
5	4.9	4.9	3.5
6	-1.8	-6.4	-5.7
7	3.2	4.8	7
8	19.1	5.3	7.1

(b) Delays detecting the end of AF.

Table 3: Delays of the method's output.

Then a correction can be implemented in order to reduce the number of errors (in this case, the whole time interval during the 'pre AF' and 'post AF' phases detected in a wrong way is considered as an error). We introduced an artificial time delay: if the output is indicating the absence of the phenomenon under study (then is set to zero), the first time the method signals a one we wait a prefixed time to set the output to one; if after this time the method is still indicating the presence of the phenomenon, we set the output to one, else we don't change the output. The introduction of this correction and its duration are problem driven. Since AF is not a dead risk pathology, the problem concerning the number of errors is more important than the detection delay, then we chose to insert an artificial time delay of 3 minutes. Doing that, we decreased noticeably the number of errors, as shown in Table 4.

Pat. num.	**I type errors (pre)**	**I type errors (post)**	**II type errors (pre)**	**II type errors (post)**	**Duration AF (min.)**
1	1	0	3	2	521
2	0	0	0	0	613
3	17	6	1	1	433
4	0	0	5	3	13
5	1	0	4	1	56
6	23	6	4	1	442
7	7	0	10	3	319
8	0	0	10	5	229
Total	49	12	37	16	

Table 4: Number of errors before and after the introduction of the artificial time delay (we fixed $N = 600$).

6 Conclusions

In this paper we proposed a statistical tool to identify, analysing the RR intervals series, starting and ending point of an event of AF, a common cardiac arrhythmia characterized by an irregular heartbeat. We presented a method based on time series analysis and we performed a statistical test to automatically recognize the phases 'pre AF', 'AF' and 'post AF', especially in those situations where the AF event does not follow a physiological time slot and/or the irregular heartbeat does not disappear when the event finishes.
Then we tested the proposed method on different simulated data, taking a reference ARIMA model for the AF phase, and varying the model of 'pre AF' and 'post AF' phases. When the reference model was quite different from the others, we obtained good results.
Then we applied the method to real RR intervals data. The results we obtained confirmed the goodness of the proposed method, that seems to be able to identify starting and ending of an event of AF even when AF follows or comes before irregular heartbeat time slots. This fact provide us that this methodology may become an helpful tool for the detection of AF.

Acknowledgments

The authors wish to thank prof. Luca Mainardi responsible of the Biomedical Signal Processing Laboratory of the Department of Bioengineering, Politecnico di Milano for supplying the data and dr. Valeria Vitelli for technical support in the statistical analysis.

References

[1] Bootsma, B. K., Hoelen, A. J., Strackee, J., Meijler, F. L.: Analysis of R-R Intervals in Patients with Atrial Fibrillation at Rest and During Exercise. Circulation **41**, 783-794 (1970)

[2] Box, G. E. P., Jenkins, G. M.: Time series analysis - forecasting and control. Holden-Day, San Francisco (1970)

[3] Box, G. E. P., Pierce, D. A.: Distribution of Residual Autocorrelations in Autoregressive-Integrated Moving Average Time Series Models. Journal of the American Statistical Association, **65**, 1509-1526 (1970)

[4] Cerutti, S., Mainardi, L. T., Porta, A., Bianchi, A. M.: Analysis of the Dynamics of RR Interval Series for the Detection of Atrial Fibrillation Episodes. Computers in Cardiology, **24**, 77-80 (1997)

[5] Cammarota, C., Rogora, E.: Independence and symbolic independence of non stationary heartbeat series during atrial fibrillation. Physica A **353**, 323-335 (2005)

[6] Couceiro, R., Carvalho, P., Henriques, J., Antunes, M., Harris, M., Habetha, J.: Detection of Atrial Fibrillation using model-based ECG analysis. Pattern Recognition, 2008. ICPR 2008, 1-5 (2008)

[7] Lindsay, A.E.: ECG learning centre, (2006) [online] http://library.med.utah.edu/kw/ecg/index.html

[8] Ljung, G. M., Box, G. E. P.: On a Measure of Lack of Fit in Time Series Models. Biometrika, **65**, 297-303 (1978)

[9] Mainardi, L., Sornmo, L., Cerutti, S.: Understanding Atrial Fibrillation: the signal processing contribution. Morgan and Claypoll Publishers, (2008)

[10] Montgomery, D. C.: Introduction to Statistical Quality Control. McGraw-Hill, New York (2005)

[11] R Development Core Team: R: A language and environment for statistical computing, R Foundation for Statistical Computing, Vienna, Austria (2008)

MOX Technical Reports, last issues

Dipartimento di Matematica "F. Brioschi",
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

20/2012 Ieva, F.; Paganoni, A. M.; Zanini, P.
Statistical models for detecting Atrial Fibrillation events

19/2012 Faggiano, E.; Antiga, L.; Puppini, G.; Quarteroni, A.; Luciani G.B.; Vergara, C.
Helical Flows and Asymmetry of Blood Jet in Dilated Ascending Aorta with Normally Functioning Bicuspid Valve

18/2012 Formaggia, L.; Vergara, C.
Prescription of general defective boundary conditions in fluid-dynamics

17/2012 Manzoni, A.; Quarteroni, A.; Gianluigi Rozza, G.
Computational reduction for parametrized PDEs: strategies and applications

16/2012 Cutri', E.; Zunino, P.; Morlacchi, S.; Chiastra, C.; Migliavacca, F.
Drug delivery patterns for different stenting techniques in coronary bifurcations: a comparative computational study

15/2012 Mengaldo, G.; Tricerri, P.; Crosetto,P.; Deparis, S.; Nobile, F.; Formaggia, L.
A comparative study of different nonlinear hyperelastic isotropic arterial wall models in patient-specific vascular flow simulations in the aortic arch

14/2012 Fumagalli, A.;Scotti, A.
An unfitted method for two-phase flow in fractured porous media.

13/2012 Formaggia, L.; Guadagnini, A.; Imperiali, I.; Lever, V.; Porta, G.; Riva, M.; Scotti, A.; Tamellini, L.
Global Sensitivity Analysis through Polynomial Chaos Expansion of a basin-scale geochemical compaction model

12/2012 Guglielmi, A.; Ieva, F.; Paganoni, A.M.; Ruggeri, F.
Hospital clustering in the treatment of acute myocardial infarction patients via a Bayesian semiparametric approach

11/2012 Bonnemain, J.; Faggiano, E.; Quarteroni A.; Deparis S.
A Patient-Specific Framework for the Analysis of the Haemodynamics in Patients with Ventricular Assist Device