Proceedings of the 17th World Congress
The International Federation of Automatic Control
Seoul, Korea, July 6-11, 2008

A Robust Auto-tuning on-line Trend Extraction Method

S. Charbonnier, C. Damour

Laboratoire d'Automatique de Grenoble (UMR 5528 CNRS, INPG, UJF)
BP 46, 38402 St Martin d'Hères France
tel : (33) 476-82-64-15 - fax : (33) 476-82-63-88
email : *Sylvie.Charbonnier@inpg.fr*

Abstract: On-line trend extraction is the first step to be achieved by a pattern-matching diagnostic system. Indeed, most pattern-matching diagnostic methods are based on the classification of qualitative or semi-qualitative trends extracted from one or several signals. The relevance of the trend extracted is a key point for the diagnostic system accuracy. This paper presents a trend extraction method which is robust to the presence of artefacts and step-like variations and does not require a priori tuning of the parameters of the method. The parameters are tuned on line by the algorithm itself (auto-tuning method), using a robust estimate of the signal variability. Results obtained on both simulated data and real data show the efficiency of the method.

1. INTRODUCTION

The dynamical behaviour of complex systems may be difficult to translate into a mathematical model. The development of diagnostic systems for theses processes requires the use of adapted methods, which are not based on a dynamical model. Among the methods available, pattern matching consists in detecting on line the occurrence of a specific pattern (or temporal shape) on one or several signals monitored on the system. This method can be used any time a fault signature can be described by a mono or multivariable temporal shape. Pattern matching has been applied to several processes, such as chemical or bio-chemical processes (Bakshi and Stephanopoulos, 1994, Rengaswamy and Ventaakasubramanian, 1995, Colomer et al, 2002, Dash et al., 2003), to determine unstable conditions in a blast furnace for instance (Gamero et al, 2006). Trends extracted are linear or quadratic. Another field of application is Intensive Care Unit patient's monitoring, where the aim of a diagnostic system is to recognize on line a change in the patient's health state and to be able to filter false alarms by detecting un-physiological changes (Hunter and McIntosh, 99, Calvelo et al. 2000). Trends extracted in this application field are usually linear.

Diagnosis by pattern matching is achieved in two steps :

- in a first step, the signals monitored on the system are converted on-line into temporal trends. The signal trend consists of a succession of contiguous temporal episodes describing the temporal evolution of the signal.

- in a second step, the trends are used as inputs to a decision system whose outputs correspond to the faults to be diagnosed on line. The decision system can be either a classifier tuned using a training data-base or a knowledge-based system built using expert's rules.

Both steps – trend extraction and decision- are equally important to obtain a reliable diagnostic system. Indeed, errors on the trend are propagated into the decision module and may result in diagnostic errors.

In previous papers (Charbonnier et al. 2005), we presented a methodology to extract on line trends from an univariate time series, which was more specially developed for physiological parameters recorded on line on patients hospitalised in ICU. The high level of noise corrupting these signals as well as the presence of step variations that must be preserved on the trend makes it difficult to apply trend extraction methods using the first and second derivative of the signal, such as the ones used on signals recorded on chemical systems.

The trend is a succession of contiguous temporal episodes. In our method, an episode is defined by equation (1):

$$\text{Episode} = \{\text{primitive}, k_0, y_0, k_f, y_f\} \qquad (1)$$

with: k_0 the time when the episode begins, y_0 the signal value at time k_0, k_f the time when the episode ends, y_f the signal value at time k_f. Three primitives are used to describe the signal trend: *Steady, Increasing, Decreasing*, corresponding to the physicians' vocabulary.

The method consists in splitting the signal into line segments using least square linear approximation and then classifying the segments into episodes. Though good results were obtained on a large variety of physiological signals (Charbonnier and Gentil, 2007), the method can still be improved. Indeed, we observed on simulated data that the method was sensitive to the presence of numerous or large artefacts (ie sudden variations due to measurement errors) and that shifts could be observed between the beginning of a step variation and its corresponding episode. Moreover, the method requires the tuning of three parameters whose values must be found manually for each signal. This makes the method difficult to use when the number of signals monitored is large.

The aim of this paper is thus to present a modified version of the methodology which is able to overcome these drawbacks, ie an auto-tuning trend extraction method robust to the presence of artefacts and step variations.

The outline of the paper is as follows. The previous version of the method is briefly described and reasons for its inadequacy to handle artefacts and step detection are given in section two. The modifications of the method are presented in section 3. Results obtained by the previous method and the new one on a set of simulated signals are compared and analysed in section 4 and an example of the results obtained by the new method on a real signal is presented.

2. ON LINE TREND EXTRACTION

2.1 Trend extraction

The methodology developed to extract on-line the trend from an univariate time series is briefly described in this section. It was published in Charbonnier et al, 2005. The trend is a succession of contiguous temporal episodes. An episode is defined by equation (1).

The methodology consists of three successive steps, completed on-line at each sampling time.

First step : Segmentation of the data and residual calculation. A segmentation algorithm splits the data on line into successive line segments expressed by equation (2):

$$y(k)=p_i(k-k_{oi})+y_{oi} \qquad (2)$$

with k_{oi} the time when the segment begins, p_i its slope and y_{oi} the ordinate at time k_{oi}. Two successive segments may be discontinuous. The segmentation algorithm uses the cumulative sum (CUSUM) to determine on line the moment when the linear approximation is no longer acceptable and when the new linear function (eq. (2)) that now best fits the data must be calculated. The CUSUM is the cumulative sum of the differences between the signal $y_m(k)$ and the linear model extrapolation $y(k)$, computed at each sampling time. At time $k_{i,1}$, when the CUSUM value crosses a first threshold (th_1, first tuning parameter), the corresponding signal is stored in a block (named block of abnormal values). It is stored while the value of the CUSUM remains below a second threshold. At time $k_{i,2}$, named the segmentation time, when the CUSUM crosses the second threshold (th_2, second tuning parameter), the algorithm calculates the new linear approximation on the data stored in the block of abnormal values, ie data recorded between $k_{i,1}$ and $k_{i,2}$, using the least squares criterion. The CUSUM is reset and the block of abnormal values is emptied (**Figure 1**).

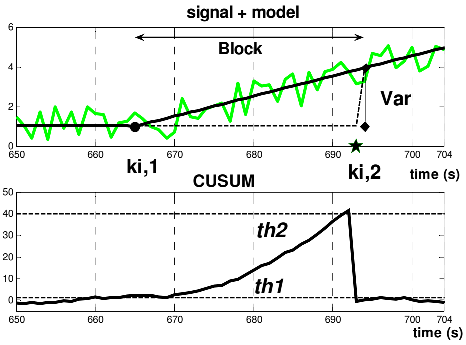

Figure 1 : Segmentation algorithm

The difference between the linear approximation calculated by the segmentation algorithm, $y(k)$, and the measured signal, $y_m(k)$, is calculated. This variable is called the residual $Res(k)$ (eq. 3). It corresponds to the part of the signal that is filtered by the segmentation algorithm.

$$Res(k) = y_m(k) - y(k) \qquad (3)$$

The signal decomposition into line segments is mainly tuned by parameter th_2, which fixes the filtering effect. Any transient variation which change integral is lesser than th_2 is filtered. As instance, a step transient of amplitude A and duration Δ is filtered if the product A.Δ is less than th_2.

Second step : Transformation of segments into episodes. At each segmentation time $k_{i,2}$, the new segment is transformed into one or two episodes, depending on the value of the discontinuity between the previous and current segment. If the absolute value of the difference between the value at the beginning and the end of the segment (see Var on **Figure 1**) is greater than th_c, the primitive is either increasing or decreasing, depending on the sign. Else, it is steady. The meaning of th_c is thus that any variation whose amplitude is greater than th_c is significant from an operator's point of view. Any variation whose amplitude is less than th_c is filtered.

Third step: Aggregation. The current episode is then aggregated, if possible, with the previous one to form the longest possible episode. For a detailed description of the on-line implementation, see (Charbonnier et al., 2005).

2.2. Drawbacks

The method is sensitive to large artefacts and step variations. An artefact is a sudden and large variation of the signal lasting a few sampling periods, mainly due to measurements errors. The method used to calculate the linear approximation in the segmentation algorithm is the least squares approximation. This method gives correct results if the noise corrupting the data can be assumed to be gaussian white noise. However, when artefacts occur on the signal, the assumption is no longer valid and the results may be biased. Moreover, if a step variation occur on the signal when the value of the CUSUM is between th_1 and th_2 (a step variation is present in the part of the signal used to calculate the new linear approximation), the least square approximation provides inaccurate parameters and the step variation may be represented by a sharp linear variation.

Three parameters (th_1, th_2, th_c), which were defined in section 2.1., are to be tuned a priori for each signal. This can be a very tedious task if several signals are recorded on a process. th_1 and th_2 tune the segmentation, yet th_1 does not significantly change the decomposition into segments as long as its value is not significantly greater than the noise level. th_2 determines the filtering effect. If it is too small, a new segment is calculated frequently, and the filtering effect is poor. Else, if it is too large, important transient variations may be missed.

If the value of th_c is too low, small variations are classified as increasing or decreasing episodes. This can result in successive erroneous increasing/decreasing episodes due to noise. Inversely, if the value of th_c is too high, significant variations are missed.

2. MODIFIED VERSION

3.1. Robustification of the method

The method is made less sensitive to artefacts and step changes in two ways :

- The linear approximation method was changed. The Siegel's repeated medians filter was used instead of the least squares method (Fried et al., 2006). This method calculates the parameters of a linear function that minimize the median of the absolute difference between the signal and the linear approximation. It is very powerful when artefacts are present thanks to the median insensitivity to large values. However, when a step variation is present in the data to be approximated, the output of the segmentation algorithm is a step, but its time of occurence can be shifted in time.

- A method to detect abrupt changes (artefacts or step changes) in the signal was implemented. An abrupt change in the level of the signal lasting less than D (tuning parameter) sampling periods is an artefact. It is automatically removed from the data set used to calculate the linear approximation. An abrupt change lasting more than D sampling periods is a step change. The calculation of the new linear approximation is achieved twice : at first using the data stored in the block of abnormal values recorded before the occurrence of the abrupt change, then using the data recorded during the D sampling periods after the abrupt changes.

The detection of abrupt changes was performed in the following way :

The signal $y_m(k)$ is derived numerically to remove the continuous component of the signal (eq.4).

$$y_m^{'}(k) = y_m(k) - y_m(k-1) \tag{4}$$

The median of the derived signal is calculated on a moving time window of size N (eq. 5).

$$med(k) = median_{i=k-N:k}(y_m^{'}(i)) \tag{5}$$

σ, (eq. 6), the median calculated on the time window N of the absolute difference of the derived signal and its median is calculated to provide an estimation of the variability of the derived signal which is more robust to the presence of artefacts than the classical standard deviation.

$$\sigma(k) = median_{i=k-N:k}(\left|y_m^{'}(i) - med(k)\right|) \tag{6}$$

An abrupt change is detected at time k+1 if the absolute value of the difference of the derivative of the signal at this time and the median of the derived signal calculated on the time window N is greater than α times σ (eq. 7).

$$\text{If } \left|y_m^{'}(k+1) - med(k)\right| \geq \alpha\sigma(k) \\ \rightarrow \text{abrupt change detected at k + 1} \tag{7}$$

The time window, N, was chosen equal to 60 sampling periods, which is a tread off between signal stationarity assumption and accurate estimation of the variability. α was chosen equal to 5. Indeed, if the derivative of the signal is a constant value C corrupted with Gaussian white noise of standard deviation σ, 99% of the derived signal is between [C-3.16σ, C+3.16σ]. Choosing α equal to 5 makes it sure to detect changes which are much larger than the level of noise. These values of N and α do not depend on the signal monitored and can be used to process any signal. The detection threshold is a function of the signal variability, which is estimated on-line.

3.2. Auto-tuning of the parameters

In this version, the three tuning parameters (th_1, th_2, th_c) are not fixed but are modified on line, at each segmentation time $k_{i,2}$, according to the signal variability. The signal variability is estimated using the signal residual (eq. (3)), which is the part of the signal filtered by the segmentation algorithm. The residual is a non stationary signal, of zero mean and time varying variance.

The residual variability is estimated in a similar way as section 3.1.: The median of the residual is calculated on a time window of size Nr preceding the segmentation time $k_{i,2}$ (eq. 8).

$$med_{res}(k_{i,2}) = median_{j=k_{i,2}-N_r:k_{i,2}}(res(j)) \tag{8}$$

σ_{res} is the median of the absolute difference of the residual and its median calculated on the time window Nr (eq. 9).

$$\sigma_{res}(k_{i,2}) = \\ median_{j=k_{i,2}-N_r:k_{i,2}}(\left|res(j) - med_{res}(k_{i,2})\right|) \tag{9}$$

The median is used instead of the mean and standard deviation to reduce the influence of artefacts on the signal. Indeed, artefacts have little influence on the value of σ_{res} as long as their cumulated duration is strictly less than half the size of the time window N_r. N_r was fixed to 60 for the same reasons as exposed in section 3.1.

th_1, th_2, th_c are adapted at $k_{i,2}$ from eq. (10), using two tuning parameters, β and Δ.

$$\begin{aligned} th_c &= \beta\sigma_{res}(k_{i,2}) \\ th_2 &= th_c\Delta = \beta\sigma_{res}(k_{i,2})\Delta \\ th_1 &= \sigma_{res}(k_{i,2})/2 \end{aligned} \tag{10}$$

In this version, the meaning of th_c has changed. An increasing (or decreasing) episode is no longer an episode whose amplitude is greater than a fixed value, corresponding to the minimal value above which a variation is considered significant from an operator's point of view. An episode is increasing if its amplitude is significantly greater than the level of noise, from a statistical point of view. Thus, the value of β can be fixed for any signal at 3 or 4, depending on the sensitivity to changes required for the diagnosis of the process. Indeed, under the assumption that the signal is a constant C corrupted with additional Gaussian white noise of standard deviation σ (the trend is *steady*), 99% of the signal remains between the bounds [C-3.16σ, C+3.16σ]. th_2 is tuned by the parameter Δ, which expresses the delay the algorithm takes to detect a step change of amplitude th_c (see section 2.1). Δ can be given the same value to any variable recorded on the process. th_1 is tuned so its value remains below the noise level.

4. RESULTS

4.1. Results on simulated data

The performances of the previous trend extraction algorithm (PTE) and the new one (NTE) presented in this paper were compared on simulated data.

Robustness to abrupt changes

To analyse the robustness of NTE to artefacts and step changes, a set of simulated data was generated (Figure 2, upper part). It is composed of a deterministic reference signal *r* with additional Gaussian white noise with standard deviation σ. The reference signal is composed of constant parts and abrupt changes (amplitude A). Two artefacts occurring less than N (N=60) sampling periods (eq. 5) before the step change, of amplitude A_r and duration 5 sampling periods, are added to the signal. For each simulation, $A = A_r = 6\sigma$. D, the maximal length of an artefact is fixed to 10 sampling periods.

A set of simulated data was carried out with σ varying from 0.2 to 3. 20 simulations were performed for each σ. NTE and PTE were applied to each simulation. To compare the results obtained by the two methods, criterion C was calculated on each simulation. C is the maximal value of the absolute difference between the reference signal r and the output from the segmentation algorithm y, calculated in the vicinity of the 3 abrupt changes (10 sampling periods around the abrupt change).

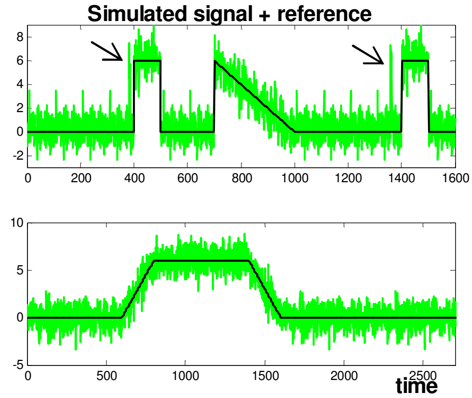

Figure 2 : Simulated data ; Upper part : abrupt changes, arrows show the two additional artefacts Lower part : steady, increasing, decreasing

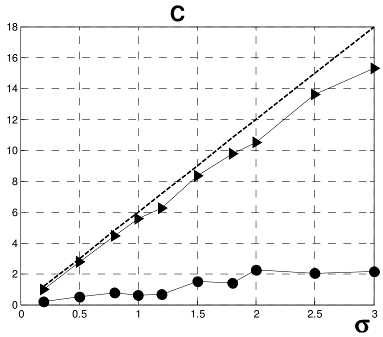

Figure 3: Detection of abrupt changes. Criterion C for different values of σ (▲: PTE, ●: NTE)

Figure 3 presents the median value of C on the 20 simulations carried out for each σ, with σ varying from 0.2 to 3. It can be seen than C is much smaller when NTE is used. The value of C is about the value of σ, meaning that the abrupt changes are detected at the exact time of occurrence. On the contrary, C is just slightly smaller than the amplitude of the abrupt changes (the dotted line is showing the amplitude of the abrupt changes for the corresponding σ), showing that abrupt changes detected with PTE do not match abrupt changes on the reference signal.

Auto-tuning parameters

To evaluate the ability of NTE to correctly extract the trend from any signal without a-priori tuning the three parameters th_1, th_2, th_c, another set of simulated data composed of a reference signal r with additional Gaussian white noise with standard deviation σ was generated (Figure 2, lower part). The reference signal is composed of a constant part (*steady*), a linearly increasing part (amplitude Ai ; *increasing*), a constant part (*steady*), a linearly decreasing part (amplitude -Ai ; *decreasing*) and a constant part (*steady*). Simulations were carried out with σ varying from 0.2 to 3. For each simulation, $\frac{A_i}{\sigma} = 6$. Thus, the amplitude of the increase varies from 1 to 18 while the signal to noise ratio is kept constant. A set of 20 simulations was achieved for each σ.

PTE was tuned to obtain correct results on the simulations carried out with σ=1. For NTE, β was chosen equal to 4 and Δ to 60 sampling periods.

Three criteria were used to compare the results :

- C1 : The mean of the absolute difference between the reference signal and the output from the segmentation algorithm calculated during the whole simulation.

- C2 : The number of sampling periods detected as *steady* by the trend extraction method and actually corresponding to constant parts (in percentage of the number of constant sampling periods on the reference signal r)

- C3 : The number of sampling periods detected as *increasing/decreasing* by the trend extraction method and actually corresponding to increasing or decreasing parts (in percentage of the number of increasing/decreasing sampling periods on the reference signal r)

Figure 4 presents the median value of C1 calculated on the 20 simulations, for σ varying from 0.2 to 3. About the same results are obtained with both methods for σ=0.5 to 1.2. Then, when σ increases, PTE cannot cope with the increase in the noise level. The value of th_2 is not large enough for the PTE segmentation algorithm to correctly filter the noise on the signal. When σ =0.2, C1 increases too because the PTE algorithm filtering effect is too high, and transient parts are not correctly segmented. On the contrary, results obtained with NTE are correct for the whole range of σ.

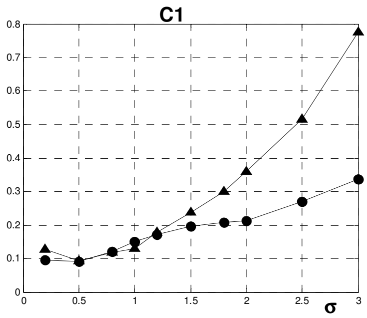

Figure 4 : Filtering effect of the segmentation algorithm; Criterion C1 for different values of σ (▲: PTE, ●: NTE)

Figure 5 presents C2 versus C3, obtained with PTE and with NTE. It shows that NTE obtained correct results for the whole range of σ. The points are in the vicinity of the optimal point (1,1). C2 is nearly 100%, meaning that all the steady parts are correctly detected. C3 is between 70% and 90%. 100% is not reached because of the quite high noise to signal ratio (18 %), which makes the beginning and end of the increasing (decreasing) periods difficult to detect with accuracy. This fact is emphasized for σ=0.2, where the performance decreases to (0.85, 0.55). PTE obtains correct results for σ=0.5 to 1.5. For σ=0.2, the value of th_2 and th_c are too large for the algorithm to detect any increasing or decreasing episode (point (1,0)). When σ is greater than 1.8, th_2 and th_c are too small, which results in the segmentation algorithm filtering effect to be too weak. Thus, many increasing/decreasing episodes due to the noise are extracted during the steady parts and the criteria are shifted towards (0.5, 0.5).

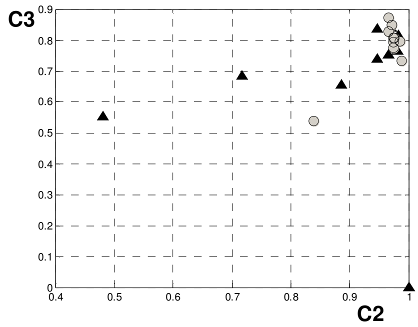

Figure 5 : Detection of trends : Criterion C3 versus criterion C2, for different values of σ (▲: PTE, ❍: NTE)

4.2. Results on real data

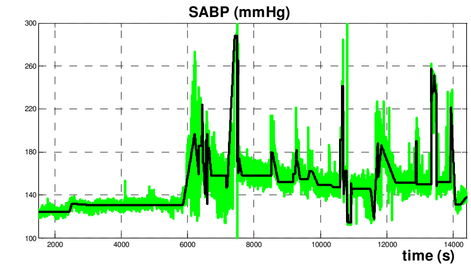

Figure 6 : Results of NTE on Systolic Arterial Blood Pressure data

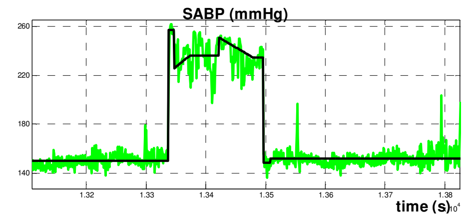

Figure 7 : Results of NTE on Systolic Arterial Blood Pressure data – Zoom in time

NTE (tuned with β=4 and Δ=60) was applied on the systolic arterial blood pressure (SABP) recorded during 4 hours on a patient hospitalised in Intensive Care Unit. Results are displayed **Figure 6**. It shows that NTE is able to correctly extract the trend of this signal without a priori tuning. During the first part of the recording, the patient was asleep. He woke up during the recording which results in a significant change in the variance of the signal around 6000 sampling periods. A small amplitude change in SABP, occurring slightly after 2000 seconds, is detected in the first part of the recording whereas changes of the same amplitude are filtered in the second part. This example shows the interest to apply NTE on physiological signals, since their variance can vary greatly, depending on the patient's clinical state. Step-like variations are correctly detected by NTE (Figure 7). These un-physiological increases, due to medical care, can be easily recognized by a pattern matching system, which can avoid triggering an alarm, as proposed in Charbonnier and Gentil, 2007.

5. CONCLUSION

A new trend extraction method is presented in this paper, which is robust to the presence of artefacts and step-like variations and does not require a priori tuning of the three parameters of the method initially proposed. Detection of abrupt changes in the signal and artefact filtering are achieved using a robust linear approximation method and a robust estimate of the variability of the signal, based on the median. The tuning of the method is achieved by statistical reasoning and adapted on-line. Instead of three fixed parameters (th_1, th_2, th_c) that had to be tuned in the previous method for each signal, only one parameter, Δ, is still to be tuned. It fixes the delay accepted to detect the smallest change on the process variables. The same value of Δ provided correct results for simulated signals with variation range from 1 to 15 as well as on real physiological data.

REFERENCES

Bakshi, B.R., Stephanopoulos G. (1994), « Representation of process trends IV. Induction of real-time patterns form operating data for diagnosis and supervisory control », *Comp. Chem. Engng.* Vol. 18, p. 303-332.

Calvelo D, Chambrin MC, Pomorski D, Ravaux P. Towards symbolisation using data-driven extraction of local trends for ICU monitoring. Artificial Intelligence in Medicine 2000; 1-2: 203-223.

Charbonnier S., Garcia-Beltan C., Cadet C., Gentil S. (2005), "Trends extraction and analysis for complex system monitoring and decision support" *Engineering Applications of Artificial Intelligence* Vol 18, n°1, pp 21-36.

Charbonnier S., Gentil S. (2007) "A trend-based alarm system to improve patient monitoring in Intensive Care Units", *Control Engineering Practice*, volume 15, n° 9, pp. 1039-1050

Dash S., Rengaswamy R., Venkatasubramanian V. (2003) "Fuzzy logic based trend classification for fault diagnosis of chemical processes" *Computers and Chemical Engineering*, vol. 27, pp. 347-362.

Fried R., Bernholt T., Gather U. (2006), "Repeated median and hybrid filter", *Computational statistics & data analysis*, vol.50, pp. 2313-2338.

Gamero F.I., Colomer J., Melendez J., Warren P. (2006) "Predicting aerodynamic instabilities in a blast furnace" *Engng. Applic. Artif. Intell.*, vol. 19, pp. 103-11.

Hunter J, McIntosh N. Knowledge based event detection in complex time series data. *In Proc. AIMDM'99, Lecture Notes in Artificial Intelligence* 1999; 1620: 271-280

Rengaswamy R., Ventaakasubramanian V. (1995), « A syntactic pattern recognition approach for process monitoring and fault diagnosis », *Engng. Applic. Artif. Intell.,* Vol. 8, p. 35-51.