

**Automatic real-time fault detection for industrial
assets using metasensors**

T D Butters, S Güttel, J L Shapiro and T J Sharpe

September 2015

MIMS EPrint: **2015.86**

Manchester Institute for Mathematical Sciences
School of Mathematics

The University of Manchester

Reports available from: <http://www.manchester.ac.uk/mims/eprints>

And by contacting: The MIMS Secretary
School of Mathematics
The University of Manchester
Manchester, M13 9PL, UK

ISSN 1749-9097

Automatic real-time fault detection for industrial assets using metasensors

T D Butters^{†}, S Güttel[†], J L Shapiro[§], T J Sharpe^{*}*

^{}Sabisu, Arden Hall, Brooklands Road, Sale, M33 3SJ, UK, tim.butters@sabisu.co*

[†]School of Mathematics, The University of Manchester, Manchester, M13 9PL, UK

[§]School of Computer Science, The University of Manchester, Manchester, M13 9PL, UK

Keywords: Metasensor, Condition Monitoring, Preventative Action, Asset Management.

Abstract

Large-scale industrial plants require physical sensors to continuously measure quantities such as temperatures or pressures. A large number of sensors is required to accurately describe the operating state of the plant, which unfortunately makes it very difficult for them to be effectively monitored by human operators. In this work we present a method to construct so-called metasensors, virtual sensors that compress the information from several sensors in an optimal manner. These metasensors are used as inputs to a novel anomaly detection system that automatically alerts operators to abnormal operation behaviour.

1 Introduction

Industrial plants are composed of a large number of individual assets, each monitored by sensors recording physical quantities such as temperatures and pressures. The number of sensors across an industrial plant is typically large to accurately describe its operating state. This results in a vast number of real-time signals, far too many to be actively monitored by human operators. It is therefore common to identify a small number of key sensors for continual monitoring based on past plant behaviour. Although there is some merit to this approach, it discards the majority of available information and is devoid of any mechanism to identify previously unseen anomalous behaviour. It is also possible to identify too many “important” sensors, leading to the need of further prioritisation that potentially eliminates meaningful sensor data. It is clear that this manual selection approach has many problems.

To improve on real-time sensor monitoring systems the authors have devised a method of constructing so-called *metasensors*. Metasensors are virtual sensors that encapsulate the state of an individual piece of equipment in a single meaningful time series. They are composites of real sensor data, combined in a way that guarantees minimal loss of information. Metasensors are trained on a small amount of data and then continually

updated in real-time at the same frequency as the underlying physical sensors, or using averaged data over a fixed time period if data smoothing is required.

Although the number of metasensors required to describe the current plant state is typically much smaller than the number of physical sensors there are still opportunities for faults to be missed if manual monitoring is relied upon. Automatic anomaly detection provides a method (or methods) of alerting operators to incipient behaviour. Time series analysis techniques are used to identify changes in behaviour that are likely to be due to the development of a fault.

In the system presented here anomaly detection is performed on the constructed metasensors to highlight points of concern and alert operators to potential issues. This allows action to be taken before the asset suffers serious damage. Several anomaly detection algorithms are employed to cover a wide range of anomaly types. We found it beneficial to combine two or more anomaly detection algorithms using a quorum voting system, thereby minimising the false positive detection rate whilst retaining sensitivity.

In this work we provide an introduction to metasensors (section 2) and automated time series anomaly detection (section 3). We then present a case study showing the application of these methods to a Sabis UK ethylene production plant (section 4). This study focuses on furnace components, for which the system successfully identifies tube leaks approximately up to 22 days before they were identified by plant engineers. These results, discussed in section 5, demonstrate the effectiveness of our metasensor anomaly detection system in identifying faults at an early stage whilst adapting to continually changing operational conditions. We conclude the paper with section 6.

2 Metasensors

Metasensors are constructed using sensor data from a relatively short time range. This training set is ideally taken during a period of normal plant operation, and is analysed to determine the optimal combination of these sensors that corresponds to the projection of the time series data from N dimensions — where N is the number of sensors being combined — to a one-




Fig. 1: Example metasensor construction from 3 signals (panels A – C). Panel D shows the resulting metasensor, in which features from all 3 original signals are clearly visible.

dimensional space that maintains the largest possible variance of the data. In this sense a metasensor is the best possible one-dimensional representation of the combined sensor data.

Although it is possible to build metasensors from a large number of sensors spread across different plant areas, we found that the most effective metasensors are built from sensors attached to a single asset. This also allows detected anomalies to be quickly associated with the faulty piece of equipment.

An example of metasensor construction is shown in figure 1. Panels A – C show the original signals to be combined, and the resultant metasensor is shown in panel D. Notice how the metasensor shows features from each original signal, with the underlying slow oscillation inherited from signal 1, the noise from signal 2, and the wild oscillatory behaviour from signal 3.

When combining multiple sensors we use principal component analysis [1] to effectively weight their contribution so that the metasensor variance is maximized. This results in highly variant sensor signals contributing a large portion to the metasensor, and signals with very small relative changes having a small contribution. As a consequence, this approach automatically ignores data from deactivated sensors as the variance of their signals will be zero.

3 Anomaly Detection

3.1 Algorithms

Several anomaly detection algorithms are used to provide flexibility in the types of abnormal behaviour that can be detected. This allows the system to be uniformly deployed over different assets with minimal configuration. Below is a list of the algorithms we have implemented and tested. All of these algorithms have in common that they build a statistical model of the time series, triggering anomalies when the data fails to fit

this model. All of these methods are adaptive, meaning that when the normal behaviour of the asset changes the algorithms will adapt to this new state automatically.

AdWin is a change-point detection algorithm developed by Bifet and Gavaldà [2]. It searches for points in the data that separate periods of behaviour based on the mean value within an expanding window. It is therefore very effective at locating changes in average value, but it is insensitive to changes in the variance of the signal.

StatWin To also capture changes in variance we developed another anomaly detection algorithm called StatWin. StatWin detects when the most recent data point provided by a metasensor lies outside of a statistical envelope defined by previous metasensor behaviour within a moving time window.

KMeans A k -means clustering algorithm is used to model the data, and then the most recent value of the metasensor is compared to this clustering to determine if it is sufficiently “close” to previously seen values within a fixed window. If it is not, the value is considered anomalous. This algorithm is based on an approach described by Hill and Minsker [3]. It effectively detects changes in both average value and variance, but rather than relying on a continuous window of allowable data it uses the cluster centroids to define regions of acceptable metasensor values. This means that this algorithm may recognise anomalies that lie within the range of previously seen data. For example, given a signal comprising entirely of the numbers 0 and 100, this algorithm would indicate a value of 50 as anomalous, as it is not sufficiently “close” to the previous values. This algorithm will be referred to as KMeans within this paper.

BCPD On-line Bayesian change-point detection (BCPD) builds a statistical model of the data and marks the points at which the underlying data distribution changes significantly [4, 5]. This approach typically provides a good general purpose algorithm that is directly applicable to any time series without any changes in configuration.

3.2 Data Windowing

Data windowing provides an efficient method of dealing with continuous data in real-time. All but one of the anomaly detection algorithms, AdWin, use a fixed size moving window. AdWin is designed around the concept of a dynamically expanding window that shrinks when changes are detected to exclude the old “different” data. This means that the window expands gradually as new points are added, and then shrinks instantaneously to remove the oldest data which does not fit the current behaviour.

KMeans and StatWin require a fixed size window to be populated before anomalies can be detected. This is because accurate statistical measurements (in StatWin’s case) and k -means

clustering (in KMeans' case) cannot be performed without a sufficient number of data points. Once “full” new data points are added to the end of the window, the oldest point is removed from the start to maintain a constant window size.

To ensure numerical stability the on-line Bayesian change point detection algorithm also uses a fixed size window, although it does not have the limitation of requiring this window to be fully populated before anomalies can be detected (as was the case with StatWin and KMeans). This change detection algorithm calculates the probability that each data point being analysed is a contiguous block of similar behaviour. Smaller and smaller weights are given to points more distant from the latest value.

3.3 Quorum Voting

Although it is possible to use individual algorithms for specific assets, better results can be seen when several algorithms are combined in a “democratic” way. Quorum voting (or majority voting) provides this framework.

The simplest form of quorum voting collects the anomaly results from each selected algorithm and only reports an anomaly if the number of algorithms detecting an anomaly is greater than a pre-set threshold (often $\geq 50\%$). It is also possible to include a time window to also count anomalies detected by several algorithms at slightly different times. The window width can range from a few minutes to hours depending on the frequency of the underlying sensor signals and the behaviour of the asset. This accounts for the different ways in which the above algorithms detect anomalies, ensuring that the sensitivity is maintained but the false positive rate is minimised.

An example of anomaly detection is shown in figure 2. The three algorithms KMeans, AdWin, and StatWin are combined to produce the quorum voting results shown as red dots in panel A. The detection results of the individual algorithms are shown in panel B.

4 Case Study

Our anomaly detection system has been recently trialled on historical data from a Sabic UK petrochemicals plant, specifically targeting its heat exchangers. Inner tube leaks are relatively common in heat exchangers and exhibit a specific sensor signature. A small hole develops in the inner tube allowing gas to escape, leading to a gradually falling temperature of the gas that remains in the pipe. The temperature of the steam, which is used as the exchange medium, remains constant, or in some cases slightly increases. The size of the hole can vary in different cases, and can increase over time. The size of the hole effects the rate at which the gas temperature decreases.

It was found that inner tube skin leaks could be positively identified soon after they develop, and often several days before they were identified by the plant engineers. Examples of these are shown in figure 3 and figure 4.




Fig. 2: Example anomaly detection on a test signal using the k -means, AdWin and StatWin algorithms. The results from the individual algorithms are shown on panel B, with the signal and quorum voting combined result shown in panel A.




Fig. 3: Anomaly detection results from heat exchanger data taken from a Sabic UK petrochemicals plant. The gas temperature is shown in blue, the steam in green and the detected anomalies in red. Anomalies are detected from the 9th of October, 20 days before the fault was identified by the engineers and the furnace taken of-line.




Fig. 4: Anomaly detection results for heat exchanger data taken from a Sabc UK petrochemicals plant. The gas temperature is shown in blue, the steam in green and the detected anomalies in red. The first anomaly is detected on the 11th of March, the leak was detected by plant engineers on the 2nd of April.

These two examples show the variation in signals (gas and steam temperature) that can be seen on a plant, with significantly more noise being present in figure 4 compared to figure 3. However, in both cases a clear downward trend in one of the signals is seen.

The first example, shown in figure 3, is a low-noise sensor pair in which a consistent but slow downward trend in gas temperature (shown by the blue line) is observed. The initial leak appears to begin on the 4th of October with the first anomaly detected on the 9th. Another anomaly possibly caused by the leak is identified on the 15th, and then a relatively high anomaly arrival rate is seen between the 25th and the 29th of October when the furnace is brought offline.

The sensors in figure 4 show a much higher noise level. The leak begins to develop around the 6th of March just before a feed change occurs, which is the cause of the spikes in temperature on the 7th. The first anomaly is reported on the 11th of March and the second on the 16th. After this point no more anomalies are reported as the signal variance is so high that the behaviour remains statistically consistent.

5 Discussion

5.1 Industrial Case Study

As can be seen from figures 3 – 4, the presented system has the potential to successfully identify asset failures. The inner tube leaks are detected soon after the leak has developed, and before the plant engineers had identified the fault. As well as the large number of assets to be monitored this is also partly due to the slow nature of this particular type of failure. If only a short section of data is assessed (*e.g.* one day), the downward

gradient of the blue temperature sensor is not identifiable. In the case shown in figure 4 this is particularly true because the level of noise is so high.

These two examples are representative of the tube leaks processed, with faults being identified between 22 and 10 days before they are identified using standard techniques. This gives sufficient notice for the plant operators to implement a strategy to either take the exchanger offline so that it can be repaired, or reduce its workload so that it is more likely to remain operational until its next scheduled shutdown. This information has a positive impact on the plant availability as it saves equipment from large scale failures that require significant amounts of time to rectify.

The amount of anomalies detected by any fault identification system is an important aspect in its “trustworthiness.” It is of unrealistic to expect that a versatile system such as this, designed to detect anomalous behaviour across different asset types with minimal configuration, will produce perfect results with no false positives and no missed faults. It seems plausible that the preferred state would be a small number of false positives with no missed true faults rather than a more conservative state in which some faults are missed but there are no false positives. On the other hand, the rate of false positives must be low enough to not lead to a loss in confidence in the results of the system.

The tube leaks shown here are a good example of this, as the detection parameters used are completely identical between the two. Although there is a high level of erratic behaviour in the gas temperature in the case shown in figure 4 only two anomalies are identified. This is due to the adaptive nature of the anomaly detection algorithms, ensuring that the anomalies are considered within the context of current behaviour, and the quorum voting approach.

5.2 Metasensor Anomaly Detection

Although it would be possible to perform anomaly detection on the sensors individually this has several drawbacks. Firstly the number of sensors can be large, which introduces performance issues if the system is intended to be used in real-time. This is especially true if several algorithms are being utilised and then combined through quorum voting. Secondly, if the sensors are treated independently synchronised bulk movement (*e.g.* all temperatures increase at the same time) will not be detected unless the deflection in each sensor is significant. When combined through metasensor construction relatively small simultaneous changes in each sensor would result in a large deflection in the metasensor. This allows the system to detect faults quickly.

5.3 Limitations

A main limitation of this study is that only one type of fault has been considered. However, a large amount of process data was made available for analysis which allowed us to calibrate the system to account for the varying signal qualities

seen on the plant. It was also possible to use the large number of known process events such as feed changes and furnace de-cokes as features that should be identified. The short time-scale changes caused by these events can be considered anomalous as they are not part of the steady-state operational regime. Other events that could be used to identify such rapid changes may come from the alarm system. Analysis of triggered alarm patterns can reveal underlying incipient behaviour [6, 7], and this information could be used alongside the system presented here to identify the root cause of anomalies.

It is possible that, for a large number of signals, a single metasensor could be insufficient to accurately represent the variance of the original sensors. In this case it could be beneficial to use more than one principal component for the metasensor construction. However, as it was found that this system is most effectively used at the asset level, it is unlikely that such a large number of sensors would be compressed into a single metasensor.

6 Conclusions

An effective anomaly detection system has been presented, capable of detecting various types of anomalous behaviour across a large range of industrial assets. The system requires very little configuration, and quickly adapts to changing behaviour to allow its continued use for long periods of time without user intervention. The utilisation of quorum voting minimises the false positive rate of the system whilst retaining the sensitivity of the individual algorithms.

It has been demonstrated that the system successfully identifies asset faults far in advance of serious failure, allowing action to be taken to minimise overall damage and increase overall asset availability.

Acknowledgements

This work was partially funded by the Innovate UK Knowledge Transfer Partnership — KTP 9315.

References

- [1] I. Jolliffe. *Principal Component Analysis*, (Springer, New York, 1986).
- [2] A. Bifet, R. Gavaldà. “Learning from time-changing data with adaptive windowing”, *In SIAM International Conference on Data Mining*, (2007).
- [3] D. J. Hill, B. S. Minsker. “Anomaly detection in streaming environmental sensor data: A data driven approach”, *Environmental Modelling & Software*, **25(9)**, pp. 1014–1022, (2010).
- [4] R. Adams, D. MacKay. “Bayesian online changepoint detection”, *University of Cambridge Technical Report*, (2007).
- [5] P. Fearnhead, Z. Liu. “On-line inference for multiple changepoint problems”, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **69**, pp. 589–605, (2007).
- [6] T. D. Butters, S. Güttel, J. L. Shapiro. “Detecting and reducing redundancy in alarm networks”, *IEEE CASE 2015*, (in press).
- [7] T. D. Butters, S. Güttel, J. L. Shapiro, T. J. Sharpe. “Statistical cluster analysis and visualisation for alarm management configuration”, *Asset Management Conference 2014*, pp. 1–6, (2014).