



# Multi-channel anomaly detection using graphical models

Bernadin Namooano<sup>1</sup> · Christina Latsou<sup>1</sup> · John Ahmet Erkoyuncu<sup>1</sup>

Received: 25 August 2023 / Accepted: 20 June 2024  
© The Author(s) 2024

## Abstract

Anomaly detection in multivariate time-series data is critical for monitoring asset conditions, enabling prompt fault detection and diagnosis to mitigate damage, reduce downtime and enhance safety. Existing literature predominately emphasises temporal dependencies in single-channel data, often overlooking interrelations between features in multivariate time-series data and across multiple channels. This paper introduces G-BOCPD, a novel graphical model-based annotation method designed to automatically detect anomalies in multi-channel multivariate time-series data. To address internal and external dependencies, G-BOCPD proposes a hybridisation of the graphical lasso and expectation maximisation algorithms. This approach detects anomalies in multi-channel multivariate time-series by identifying segments with diverse behaviours and patterns, which are then annotated to highlight variations. The method alternates between estimating the concentration matrix, which represents dependencies between variables, using the graphical lasso algorithm, and annotating segments through a minimal path clustering method for a comprehensive understanding of variations. To demonstrate its effectiveness, G-BOCPD is applied to multichannel time-series obtained from: (i) Diesel Multiple Unit train engines exhibiting faulty behaviours; and (ii) a group of train doors at various degradation stages. Empirical evidence highlights G-BOCPD's superior performance compared to previous approaches in terms of precision, recall and F1-score.

**Keywords** Time-series · Anomaly detection · Multi-channel · Multivariate · Graphical model

## Introduction

Time-series data analysis, the process of extracting meaningful knowledge from time-ordered data and using this knowledge to diagnose and predict the past and future behaviour of assets, has become a field of interest for both academia and industry in various domains, including transportation, medicine and economics (Choi et al., 2021). With the evolution of Industry 4.0 and the vast amount of available data, time-series analysis has attracted a lot of interest in maintenance systems to detect or forecast asset health status,

prevent potential catastrophic failures, allow for better maintenance planning and, in turn, reduce costs (Manco et al., 2017; Zhao et al., 2020). Time-series data can be described as a collection of observations arranged in a time order (Hamilton, 2020). Classification, clustering, anomaly (or outlier) detection and forecasting are the most common data mining techniques employed in the current literature for time-series data analysis (Tang et al., 2019). In this context, anomaly detection refers to the process of identifying anomalous data with different severity characteristics such as fluctuation (i.e. statistical outliers) that do not conform to the common rules in time-series. An anomaly in time-series refers to a data point or sequence that deviates from the normal range. By learning the normal system operations patterns and identifying the outliers, system failures can be promptly detected.

Anomalies can be effectively detected if the properties of time-series data, including temporality, dimensionality and non-stationarity, and the status of data quality are properly understood. A detailed explanation of the properties can be found in the work conducted by Choi et al. (2021). Traditional anomaly detection approaches, including linear-based,

---

✉ Bernadin Namooano  
bernadin.namooano@cranfield.ac.uk

Christina Latsou  
christina.latsou@cranfield.ac.uk

John Ahmet Erkoyuncu  
j.a.erkoyuncu@cranfield.ac.uk

<sup>1</sup> Centre of Digital Engineering and Manufacturing, Cranfield University, College Rd, Wharley End, Bedford MK43 0AL, UK

distance-based and density-based (Alqahtani et al., 2021) are limited to manipulating univariate or single-channel time-series data sets (i.e. capturing anomalies for a single feature or variable), while they are not sufficiently advanced to handle the absence of labelled training data (Zhao et al., 2020). Additionally, as the complexity of industrial applications increases and multi-channel multivariate time-series data becomes more prevalent, the aforementioned anomaly detection methods have become impractical. In this work, multi-channel,<sup>1</sup> multivariate,<sup>2</sup> nonstationary<sup>3</sup> and noisy<sup>4</sup> time-series data will be studied. A typical example of multi-channel multivariate time-series data can be considered the data generated by sensors attached to railway train engines. Thus, on Diesel Multiple Unit (DMU) train engines, a single channel time-series data is considered the data collected from one engine. The channels should have a similar data architecture. Each channel comprises multivariate sensor measurement data (e.g. volume, pressure, temperature, speed and control signals).

With the increased availability of data, data-driven methods effectively utilise historical data to learn asset behaviour and classify or cluster newly observed data. Recent literature has explored approaches, including support vector machine methods for induction motors (Gangsar & Tiwari, 2019), neural network-based applications for engineering anomaly detection (Lin et al., 2019), deep convolutional neural networks (DCNN) for track damage detection (Chen et al., 2019; James et al. 2019; Yuan et al., 2020) and back-propagation neural networks for track spikes detection (Lu & Shen, 2020), among others. Additionally, applying fuzzy C-means clustering methods, Yin and Huang (2015) addressed faults in railway vehicle suspension. Similarly, clustering methods, such as dynamic time warping-based approaches, have been investigated for fault detection in railway point machine (Du et al., 2019). Other common approaches include shapelet analysis (Beggel et al., 2018) and matrix profile methods (Yeh et al., 2018) for unified motif discovery, anomaly detection and fault diagnosis. While these studies offer valuable insights into potential methods, their focus remains on single-variate time-series, limiting their scope.

In the context of anomaly detection and diagnosis for multivariate time-series data, relevant works are discussed. Xi et al. (2006) proposed numerosity reduction to enhance the

efficiency of Dynamic Time Warping (DTW) in one-nearest-neighbour for time-series classification. In 2013, Senin and Malinchik introduced SAX-VSM, combining Symbolic Aggregate Approximation and Vector Space Model for automated discovery and ranking of interpretable characteristic patterns, achieving accurate classification. Patri et al. (2014) proposed a shapelet forest approach for multivariate times series classification, featuring shapelet extraction, distance calculation, and new feature creation for fault detection. In 2017, Lee et al. presented a CNN model for fault detection and classification in semiconductor manufacturing processes. In more recent literature, Hsu and Liu (2021) introduced a CNN model employing data augmentation and a diagnostic layer for adaptive feature extraction, fault detection and diagnosis in semiconductor manufacturing. Chen et al. (2023) proposed LR-SemiVAE, a long-short term memory (LSTM)-based semi-supervised variational autoencoder (VAE) for anomaly detection, and RT-SemiVAE, a transformer-based VAE with parallel multi-head attention, offering more accurate anomaly detection and precise root cause localisation in a large-scale cloud environment. Additionally, Kim et al. (2023) introduced an unsupervised time-series anomaly detection method using Transformers, outperforming LSTM and CNN with self-attention mechanisms. In 2024, Bignoumba et al. presented ALignment-driven Neural Network, transforming data into pseudo-aligned latent values through a novel matrix representation. Although these methods have been applied with varying degrees of success, their application to real-world operational data, with multiple channels of data and diverse complexities, remains limited.

In the context of anomaly detection and diagnosis for multi-channel time-series data, relevant works are discussed. Zheng et al. (2014) proposed a multi-channel DCNN to handle varying lengths of sensor data. The CNN initially learns features from each sensor's subsequences using one channel, enabling the entire model to understand the features of different sensors. The CNN-extracted features are then combined and utilised with a conventional Multi-Layer Perceptron for time-series classification. An MC-CNN was also developed by Yang et al. (2015) to identify human movements, learning features and classifying datasets. Moreover, Canizo et al. (2019) proposed a deep learning-based method for supervised multi-channel time-series anomaly detection in industrial settings, utilising CNNs and Recurrent Neural Networks using independent CNNs (convolutional heads) for each sensor. Similarly, Huang et al. (2020) adopted multi-domain features into a DCNN model for tool wear prediction. In recent literature, Xu et al. (2023) presented a model for multivariate time-series anomaly detection, employing a dual-channel feature extraction module that focuses on spatial and time features. Although various approaches have been proposed for either multi-channel or multivariate data,

<sup>1</sup> A set of single channels with univariate or multivariate input time-series data with different features or variables. Each channel has a distinct distribution of values.

<sup>2</sup> A set of univariate time-series data with more than one time-dependent features or variables. Each feature depends not only on its past values but also has some dependency on other features. Using this dependency, future values can be predicted.

<sup>3</sup> The statistical properties of time-series change over time.

<sup>4</sup> Unwanted changes to signals during their capture, storage, transmission, processing, or conversion.

there is a lack of methods that effectively handle anomalies in both types of time-series simultaneously.

In multi-channel multivariate time-series anomaly detection, capturing both temporal dependency (i.e. relations between current and previous states) and interrelations between features (i.e. intra-dependency) is crucial. While intra-dependency is not a new research topic, capturing the interrelations among multiple channels has not yet been addressed (Li et al., 2022). Existing literature suggests that interrelations can be extracted from feature sparsity and effectively represented through graphical models (Fang et al. 2023; Xu et al., 2023). Graphs can visually represent time-series data, with nodes denoting process variables at each timestamp and edges representing connections between nodes. They can effectively model dependencies in multi-channel multivariate time-series data, capturing topological structures and causal relationships between features (Chen et al., 2022; Deng & Hooi, 2021; Zhao et al., 2020). In recent years, there has been an increasing interest in the use of graph machine learning for anomaly detection and diagnosis. Notably, Wu et al. (2020) introduced a graph learning approach for multivariate time-series data, able to extract the directed spatial dependency between features. Additionally, Tao and Du (2023) proposed a graph-based sparse learning framework for anomaly detection on unstructured point cloud data, representing smooth free-form surfaces. The literature also highlights relevant research on graph convolutional networks (GCN), graph dynamic autoencoders and fault detection algorithms that integrate graph modelling and minimum spanning tree graphs (Goswami et al., 2023). Motivated by the success of graphical models in diagnosing anomalies in industrial systems, this study introduces a method based on a graphical model for multi-channel multivariate time-series. The proposed model can capture inter-dependency within single-channel data and intra-dependency among multiple channels. By introducing multi-channel feature extraction, the learning process of the graphical model can capture the correlation between the hidden features across different channels, enabling prompt, accurate and automated anomaly detection and diagnosis.

This research work contributes to the literature by proposing a generic graphical model-based annotation method for detecting anomalies in multiple channels of sensor data with a shared structure in an automatic way. By adopting the graphical lasso algorithm method for a dependant graphical-based model, the proposed solution can capture the conditional dependency of channel features and dynamically detect anomalies based on unlabelled multi-channel data establishing relationships between multiple features at every time step. Therefore, the contributions of this work are summarised as follows:

- A novel graphical model-based annotation method, called G-BOCPD, is introduced for detecting anomalies in multi-channel multivariate time-series data automatically.
- A hybridisation of the graphical lasso and expectation maximisation algorithms is proposed for capturing the internal and external dependencies in multi-channel multivariate time-series data. The segments are annotated to highlight the difference in behaviour within the multivariate time-series. G-BOCPD model computes alternatively the concentration matrix and annotation of the segments, employing the graphical lasso and a minimal path clustering method, respectively.
- G-BOCPD method is validated using two real-world engineering datasets: (i) DMU train engines with faulty behaviours; and (ii) a group of train doors data with different degradation stages. Moreover, quantitative experiments are conducted to demonstrate that the proposed method can provide competitive results compared to alternative methods in the field in terms of precision, recall and F1-score.

The remainder of the paper is organised as follows: Section II describes the methodological approach used for developing the graphical model-based annotation method, G-BOCPD, for anomaly detection in multi-channel multivariate time-series. Section III presents the problem motivation for the proposed method, including a real-world multi-channel multivariate time-series data problem. Additionally, the proposed G-BOCPD is validated through a case study in anomaly detection of multiple unit diesel engines. Section IV presents a summary of critical discussion on the proposed method while comparing the results of this work with those obtained in the literature on methods for detecting anomalies in multivariate time-series. Finally, the conclusions and future work are highlighted in Section V.

## Methodology

The methodology employed for the anomaly detection is shown in Fig. 1. The step 1, where raw data are collected and organised into multiple channels. Step 2 involves implementing a sliding window technique in order to extract relevant time indexed features that may be sensitive to anomaly conditions. Step 3 estimate intra-channel dependency estimation, i.e., make use of graphical models to estimate the conditional dependencies within individual channels of the dataset. Step 4 pertains to inter-channel dependency estimation, also using graphical models to ascertain the dependencies and relationships between different channels. In step 5, the focus shifts to the computation of residuals, which detect intra

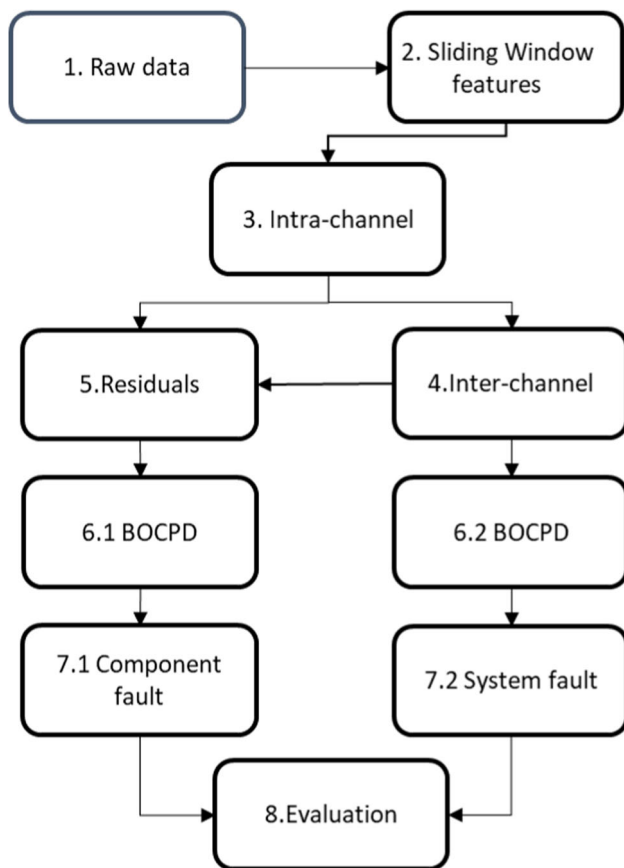


Fig. 1 Methodological approach

channel anomalies. The residuals represent each channel specific variations. Steps 6.1 and 6.2 both involve Bayesian Online Change Point Detection, first applied to each channel to detect change points in 6.1, and to detect change at inter-channel level 6.2 for system-wide fault detection. Steps 7.1 and 7.2 build upon the results of the BOCPD, with 7.1 addressing whether or not change point detected are component faults in steps 6.1, and 7.2 focusing on inter-channel faults, which involve the assessment of identified changes detected in step 6.2. Finally, Step 8 encompasses the evaluation phase, where the results of the previous steps are assessed, using performance metrics to ensure the efficacy and accuracy of the methodology.

## Problem definition

Given a multidimensional time-series,  $T = \{T_1, T_2, \dots, T_k, \dots, T_K\}$ , where each  $T_k \in R^{np}$  a multivariate time-series and  $T_{k,i}$  the observation at the time index  $i$  of the time-series  $T_k$ .  $T_{k,i}$  is a vector of  $p$  variables entries values. To examine the dependency between the multivariate time-series, two latent

variables  $C$  and  $X$  are introduced (Xie et al., 2016) such that:

$$T_k = C + X_k \quad (1)$$

where  $C$  captures shared behaviour among all the multivariate time-series and  $X_k$  the variation within the specific multivariate time-series  $T_k$ . In some conditions,  $C$  represents the behaviour shared across channels whereas  $X_k$  can be seen as a local behaviour of the component. When all the time-series have the same behaviour, the  $X_k$  represents anomalous behaviour, or noise measurement as it captures the deviation of the local component from the common behaviour. The present work assumes that latent variables  $C$  and  $\{X_k\}$ ,  $k = 1, 2, \dots, K$  are mutually independent and follow a Gaussian distribution. It has been shown in such settings (Xie et al., 2016) that:

$$\text{var}(C) = \text{cov}(T_k, T_l) \neq l \quad (2)$$

Considering  $\Theta_s$  the concentration matrix of the window  $w_s$ , the inter-channel layer aims to estimate  $\Theta_s$  for all the  $K$  multivariate time-series and  $\Theta_{0,s}$  such that:

$$\text{var}(C_s) = \Theta_{0,s}^{-1} \quad (3)$$

At the system layer, the estimation of the concentration matrix  $\Theta_{T_k}$  is as follows:

$$\Theta_{T_k,s} = (\Theta_{0,s}^{-1} + \Theta_{k,s}^{-1})^{-1} \quad (4)$$

In these settings, the following likelihood is estimated:

$$\underset{\Theta}{\text{argmin}} \sum_{k=1}^K \left[ \overbrace{\|\lambda * \Theta_{0,s}\|_1}^{\text{sparsity}} + \overbrace{-LL(T_k, \Theta_{k,s})}^{\text{log likelihood}} \right] \quad (5)$$

with  $\lambda$  representing a non-negative tuning parameter guaranteeing the sparsity. A large  $\lambda$  value, for instance, makes the entries of the concentration matrix converge to zeros,  $\|\bullet\|_1$  is the L<sub>1</sub>-norm,  $*$ , the element-wise matrix product is also known as the Hadamard product.

$$\begin{aligned} LL(T, \Theta_{u,s}) \\ = -\frac{1}{2} [(T - \mu_{u,s}) \Theta_{u,s} (T - \mu_{u,s}) + \log \det(\Theta_{u,s}) + \gamma] \end{aligned} \quad (6)$$

$\mu_s$  is the empirical mean of the samples to the window  $w_s$ ,  $\gamma$  a constant value,  $\det(\Theta_{u,s})$  is the determinant of the matrix  $\Theta_{u,s}$ .

A sub-problem of optimisation (5) has been investigated in (Hallac et al., 2018) for only one multivariate time-series ( $K = 1$ ). A similar resolution approach is taken here, as the

optimisation problem shares the same characteristics, i.e., it is a highly non-convex mixed discrete and continuous optimisation problem, with  $\Theta$  being the sets of the precision. The continuous optimisation problem is solved using the ADMM (Alternating Direction Method of Multipliers) algorithm.

**Problem resolution**

A closed form of the continuous optimisation for each multivariate time-series  $T_k$  in (5) can be rewritten as follow:

$$\max_{\Theta} \frac{1}{2} \sum_{s=1}^S (\log \det \Theta_{k,s} - \text{tr}(\mathbf{R}_{k,s} \Theta_{k,s}) - \lambda \|\mathbf{Z}_{k,s}\|_1) \tag{7}$$

Subject to  $\Theta_{k,s}$  positive and  $\mathbf{Z}_{k,s} = \Theta_{k,s}$   
 $\text{tr}(\mathbf{R}\Theta)$  is the trace of the matrix  $\mathbf{R}\Theta$ ,  $\mathbf{R}$  the empirical covariance matrix and  $\Theta = \{\Theta_{k,s}\}$ ,  $s = 1..S$ .  $\mathbf{Z}$  a consensus variable (Danaher et al., 2014). The augmented Lagrangian multipliers scaled by Frobenius norm of the minimisation form of the problem (8) defined by (Boyd, 2010) is given by:

$$L_{\rho}(\Theta_{k,s}, \mathbf{Z}_{k,s}, \mathbf{U}_{k,s}) = \{-\log \det \Theta_{k,s} + \text{tr}(\mathbf{S}_{k,s} \Theta_{k,s}) + \lambda \|\mathbf{Z}_{k,s}\|_1\} + \frac{\rho}{2} \|\Theta_{k,s} - \mathbf{Z}_{k,s} + \mathbf{U}_{k,s}\|_F^2 - \frac{\rho}{2} \|\mathbf{U}_{k,s}\|_F^2 \tag{8}$$

where  $\rho > 0$  is the ADMM penalty parameter,  $\mathbf{U}_{k,s}$  are dual variables.  $\Theta_k = \{\Theta_{k,s}\}$ ,  $\mathbf{Z}_k = \{\mathbf{Z}_{k,s}\}$ ,  $\mathbf{U}_k = \{\mathbf{U}_{k,s}\}$ ,  $s = 1..S$ ,  $\|\bullet\|_F^2$  is the Frobenius norm. The optimal solution to the problem (9) provides an estimation of the concentration matrix  $\Theta_{k,s}$ 's. ADMM for multivariate works in three main steps as follows:

At the  $i$ th iteration, update the parameter  $\Theta_k, \mathbf{Z}_k, \mathbf{U}_k$ :

- (a)  $\Theta$  - Update :  $\Theta_k^{(i)} = \underset{\Theta}{\text{argmin}}\{L_{\rho}(\Theta, \mathbf{Z}_k^{(i-1)}, \mathbf{U}_k^{(i-1)})\}$
- (b)  $\mathbf{Z}$  - Update :  $\mathbf{Z}_k^{(i)} = \underset{\mathbf{Z}}{\text{argmin}}\{L_{\rho}(\Theta_k^{(i)}, \mathbf{Z}_k, \mathbf{U}_k^{(i-1)})\}$
- (c)  $\mathbf{U}$  - Update :

$$\mathbf{U}_k^{(i)} = \mathbf{U}_k^{(i-1)} + (\Theta_k^{(i)} - \mathbf{Z}_k^{(i)})$$

The general algorithm of the multivariate  $\mathbf{Z}_k$  is presented below:

- (a) For each  $T_k k = 1..K$
- (b) Initialise the variables:  $\Theta_{k,s} = \mathbf{I}$ ,  $\mathbf{Z}_k = 0$ , and  $\mathbf{U}_{k,s} = 0$  for each  $s = 1..S$
- (c) Initialise  $\rho > 0$

For each  $i = 1..$  until convergence

- (i)  $\Theta$  - Update :  $\Theta_{k,s}^{(i)} = \underset{\Theta}{\text{argmin}}\{L_{\rho}(\Theta, \mathbf{Z}_{k,s}^{(i-1)}, \mathbf{U}_{k,s}^{(i-1)})\}$  using eigen decomposition method proposed by (Witten and Tibshirani 2009).
- (ii)  $\mathbf{Z}$  - Update :  $\mathbf{Z}_{T_k}^{(i)} = \underset{\mathbf{Z}}{\text{argmin}}\{L_{\rho}(\Theta_{k,s}^{(i)}, \mathbf{Z}, \mathbf{U}_{k,s}^{(i-1)})\}$
- (iii)  $\mathbf{U}$  - Update :  

$$\mathbf{U}_{k,s}^{(i)} = \mathbf{U}_{k,s}^{(i-1)} + (\Theta_{k,s}^{(i)} - \mathbf{Z}_{k,s}^{(i)})$$

In these settings, convergence is reached when  $\Theta_{k,s}^{(i-1)} = \Theta_{k,s}^{(i)}$  to avoid convergence issues, maximum iteration is set as a parameter (set to 1000 for the experiment). The presented graphical lasso method above provides a concentration matrix for each multivariate segment. Concentration matrix of the system at the window  $w_s$  is matrix  $\sum_s$  where the diagonal represents components layer concentration matrices.

$$\sum_s = \begin{pmatrix} \Theta_{0,s} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Theta_{k,s} \end{pmatrix} \tag{9}$$

So far, the estimation of the parameter concern only  $\Theta_k$ 's. A natural way of  $\Theta_{0,s}$  is computing the  $\Theta_{0,s}^{-1}$ 's representing the variance through a regression-based algorithm by using all pairs of variables of the observations. Xie et al. (2016) have shown that its estimation can be done for the observation belonging to the window  $w_s$  using:

$$\Theta_{0,s}^{-1} = \frac{1}{K(K-1)N} \sum_{k=1}^K \sum_{i=1}^n T_{k,i} \bullet T_{k,i}^T \tag{10}$$

Therefore, the  $\Theta_{k,s}$  can be deduced from the Eq. (4):

$$\Theta_{k,s} = (\Theta_{0,s}^{-1} - \Theta_{k,s}^{-1})^{-1} \tag{11}$$


To guarantee the positivity of  $\Theta_{k,s}$ 's, semidefinite positiveness projection (Xu & Shao, 2012) can be used as follows:

$$\Theta_{k,s} = \underset{\Theta}{\text{argmin}} \|\Theta - \Theta_{k,s}\|_{\infty} k = 0..K \tag{12}$$

where  $\|\bullet\|_{\infty}$  is the infinite norm.

**Tracking anomalies through changes**

With the concentration matrixes  $\Theta_{0,s}^{-1}$  and  $\Theta_{k,s}$  computed over the sliding window, a change point detection algorithm can be used for anomaly detection as  $\Theta_{k,s}$  are stationary signals. Hence, BOCPD method (Namoano et al., 2020a, b, c)



**Fig. 2** Train engine multiple units with three engines in the carriages E1, E2, and E3

is used to detect anomalies at both channel and intra-channel levels.

## USE CASES

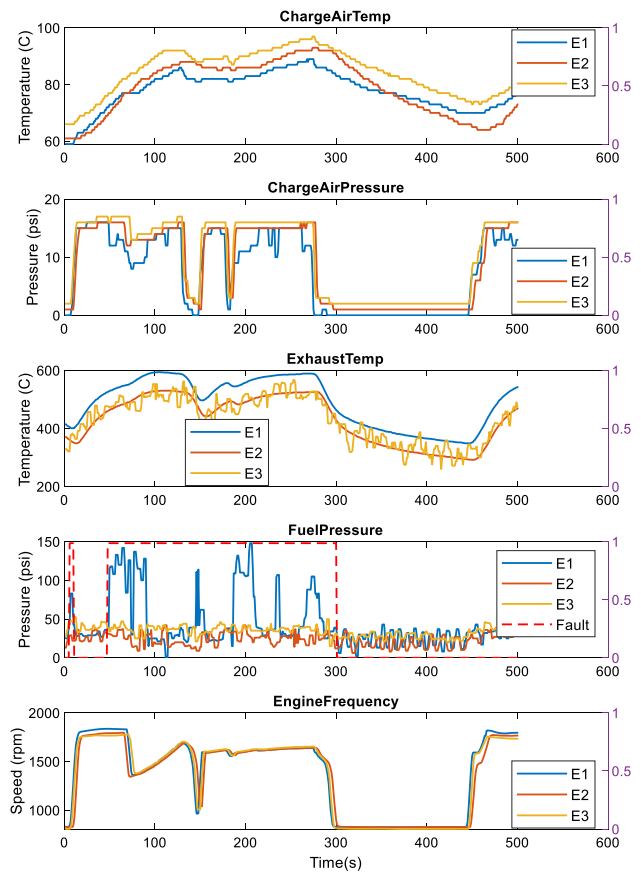
### Datasets

#### Engines

The research reported in this paper is motivated by a real-world problem of anomaly detection on multiple system components. Figure 2 illustrates a diesel multiple-unit train, composed of three carriages. The engines share the load generated from the driver actions. In each engine, sensors, placed in various components (e.g. cooling, exhaust, shaft speed), continuously gather data such as the rotational frequency, charge air pressure, cooling temperature and pressure, as well as measurements of the exhaust gases and ambient temperature.

Figures 3 and 4 show examples of several signals overlaid, collected during the train operation. Patterns comprise idling mode where the engines are on (engine rotational speed is approximately steady and around 800 rpm), acceleration in which there are rise in both pressure and temperature signals, gear changes, and deceleration where both pressures and temperature signals decrease, and the rotational speed changes to the idling rpm. The examples of signals shown include the engine charge air temperature (ChargeAirTemp), the charged air pressure (ChargeAirPressure), exhaust gas temperature (ExhaustTemp), internal fuel pressure (FuelTemp), and engine rotational speed. Each subplot represents the same signal for the three engines.

Figure 4 shows an example of a fuel pressure fault in engine E1. This pressure fault does not affect the E1 engine overall performance as the other signals in comparison to the other engines (the three engines produce roughly the same rotational speed). Conversely, in Fig. 4 where lower power event faults are observed, all the other signals of E1 are affected. Although Fig. 3 shows that anomaly has been detected only in engine E1, different faults have been observed in all the engines. It can be, thus, seen that there is a need to capture the relationship within one channel signal and between channel signals for the detection of faults occurring in the system.



**Fig. 3** Engine normal operation modes with local faults of the fuel pressure in engine E1

Two types of datasets have been used for the analysis. The first category represents a working condition with white noise where due to some unknown cause, the charge air valve has unusual behaviour, leading to a shortage of air in one of the engines. In the second category, one of the engines often does not generate sufficient power to meet the demand, causing low-power events. The analysis aims to annotate the data such that clusters represent the normal conditions and faulty behaviour. The engine datasets are available on the Cranfield cord repository (Namoano et al., 2020b).

#### Doors

The second dataset used is from doors and is collected at a frequency of 20 Hz, where many doors from the same manufacturers are supposed to have the same opening and closing profile. Datasets contain opening and closing speed and the current. As the rate of the use of these doors may not be the same depending on their location as well as the number of users, one may be interested in their variability instead of one component degradation. A way, therefore, to tackle this issue is through multivariate annotations. Figure 5 shows the opening and closing profiles, respectively. In the opening profile,

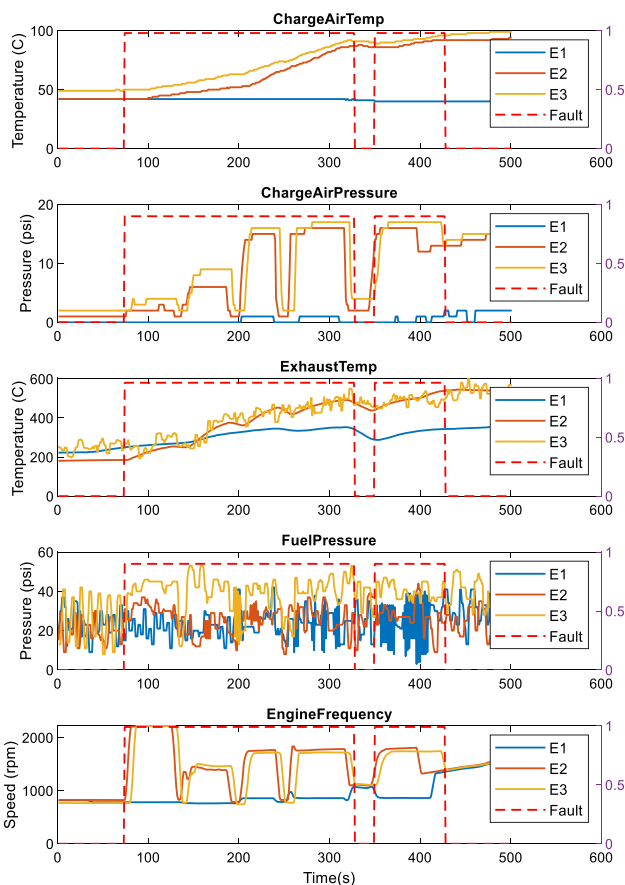


Fig. 4 Engine normal operation modes with low power fault in engine E1

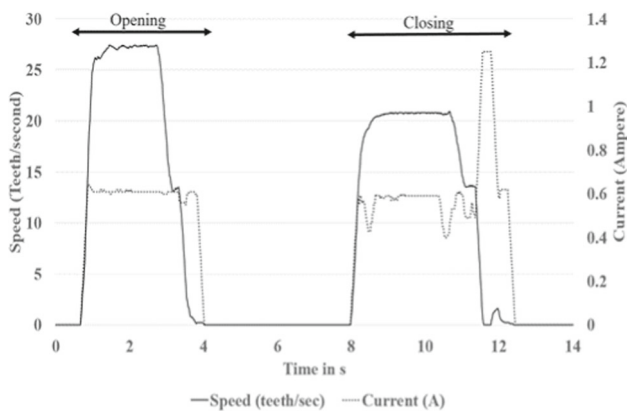


Fig. 5 An example of a door opening and closing profile

the speed, as well as the current, increased steadily up to a maximum. A slight flat constant curve follows it and then decreases to zero. A dip can be seen during the decreasing phase, at roughly half the maximum value, due to a controller avoiding fast decrease, which could lead to severe damage to the door. As it reaches its maximum opening position, the door stays in the “open state” for several seconds before




Fig. 6 Door locations in the system where the data are collected

triggering the closing motion. The closing profile follows a similar pattern to the opening, with three main differences visible in the current motion:

- The peak in the closing profile is lower than the opening one, leading to a larger closing time window.
- Depending on the position of the door with the ground floor (tilted or flat ground), an additional gravitational speed may be added. A controller, therefore, adjusts the door speed by decreasing or increasing the current. Those can be seen in the current speed profile within the time range from 8 to 10 s.
- At the end of the closing profile, an abrupt change in the current can be seen, followed by a slight speed bump. This helps to push the door to its maximal reachable position, where a locking process can be triggered.

The datasets are available on the Cranfield cord repository (Namoano et al., 2020a). Over time, such doors may fail to achieve their goal and show unusual behaviour visible through the profiles. Unusual behaviours include, for example, faster or slower speed increasing or reducing hence the required time to achieve the closing or the opening. These affect the current profile, as the controllers adjust the speed automatically by changing the current. Other key unusual behaviours are the overshoot and the variation of the speed during the steady phase, which is either the manifestation of the defective components of the door or reading sensor issues. The present work seeks to segment and annotate groups of such similar doors, as shown in Fig. 6 where each door represented from the letter A to D may have different utilisation rates. An investigation can be, for instance, looking for unusual deviations between the two groups.

## FEATURES

As already mentioned, feature extraction is critical for the performance of the methods. It enables the capture of system dynamics, which is sensitive to deviations from normalities.

**Table 1** Features extracted from door profiles

Signal	Motion	Feature
Speed	Opening	Overshoot
		Steady-state mean
		Settling time
		Gradients
	Closing	Steady-state standard deviation
		Overshoot
		Steady state mean
		Settling time
		Gradients
		Steady state standard deviation
		Overshoot
		Steady state mean
Current	Opening	Settling time
		Gradients
		Steady state standard deviation
		Overshoot
	Closing	Steady state mean
		Settling time
		Gradients
		Steady state standard deviation

**Table 2** Features extracted from engine datasets

Signal	Motion	Feature
Engine RPM	Acceleration	Overshoot
		Settling time
		Gradients
		Steady state standard deviation
	Gear Change	Change duration
		Change gradient
		Dips
		Gradient
Air Pressure Change	Acceleration	Overshoot
		Settling time
		Gradients
		Steady state standard deviation
	Gear Change	Change duration
		Change gradient
		Dips
		Gradient
Exhaust Temperature	Deceleration	Unexpected peaks
		Steady standard deviation

Tables 1 and 2 present, respectively, the features extracted from engines and door profiles. For the doors, the features include overshoot, the rising gradient, the settling time, and the steady state mean and standard deviation. These features are sensitive to fast or slow doors opening or closing, shaking doors due to spalling defects, and unexpected changes. For engines, features are extracted during acceleration, gear change, deceleration, and steady-state periods. Such features are important for detecting sensors and engine faults.

## Algorithmic parameters setting

For effective detection, different parameters setting is applied to each category of the dataset. Table 3 shows the number of multivariate used in each category ( $K$ ), the number of observations in each multivariate time-series ( $n$ ), the window length  $w_s$ , the penalty ( $\beta$ ), the sparsity penalty ( $\lambda$ ), and the ADAMM regularisation parameter ( $\rho$ ) applied to the model. These optimal parameter values are obtained through a grid search algorithm. Due to computational limitations, the parameters search interval was set [start;end;step] as follows:  $w_s = [2;300;1]$ ,  $\beta = [1-2000,1]$ ,  $\lambda = [0;1;0.01]$ ,  $\rho = [0;1;0.01]$ .

Table 4 summarises the parameter value obtained by empirical tests in BOCPD methods. For BOCPD, a prior Gaussian distribution assumption is made. The hazard parameter, which is used jointly with the Gaussian distribution for sufficient statistics computation, has also been tested empirically. Grid searches were also applied to the

BOCPD methods to select best parameters. The threshold parameter enables the adjustment for the detection sensitivity. Higher thresholds provide higher sensitivity, which can lead to noise detection as anomalies. Hence the thresholds were empirically set to avoid noise present in the data such as the engine E3 Exhaust temperature in Fig. 3, or the subtle signal differences also observed in Fig. 3, across engine charge air temperatures, and pressures. These differences are considered as normal as in normal operation, depending on the leading vehicle, DMU engines have different behaviours, either generating more power or less in comparison to the others engines.

Three performance indicators – precision, recall and F1-score, are involved in the assessment of the approaches. Precision measures the percentage of correctly annotated cases, while recall represents the percentage of actual changes that are annotated correctly. The F1-score serves as the test accuracy metric, representing the correctness of all identified cases.

$$Precision(P) = \frac{Number\ of\ Correct\ classification}{Number\ of\ annotation} \cdot 100 \quad (13)$$

$$Recall(R) = \frac{Number\ of\ Correct\ annotations}{Number\ of\ true\ classification} \cdot 100 \quad (14)$$

$$F1 - score(F1) = 2 \frac{P \cdot R}{P + R} \quad (15)$$

**Table 3** Parameters setting used for the model

Datasets	$K$	$n$	$w_s$	$\beta$	$\lambda$	$\rho$
Engine_1	3	120,000	18	325	0.2	0.98
Engine_2	3	120,000	18	350	0.6	0.60
Doors	2	36,000	20	198	0.11	0.80

**Table 4** Parameters settings used in the compared methods

Dataset	BOCPD	Hazard grid search
Engines_1	Gaussian, hazard = 95, threshold = 0.5	[1;10,000;1]
Engines_2	Gaussian, hazard = 145, threshold = 0.5	[1;10,000;1]
Doors	Gaussian, hazard = 25, threshold = 0.001	[1;10,000;1]

**Table 5** Engine 1 anomaly detection results

Methods	Precision	Recall	F1
1NN-DTW (Xi et al., 2006)	0.30	0.20	0.24
SAX-VSM (Senin & Malinchik, 2013)	0.85	0.39	0.53
Shapelet forests (Patri et al., 2014)	0.90	0.77	0.82
MC-DCNN (Zheng et al., 2014)	0.84	0.90	0.86
FDC-CNN (Lee et al., 2017)	0.91	0.94	0.92
MTS-CNN (Hsu & Liu, 2021)	0.90	0.91	0.90
<b>G-BOCPD</b>	<b>0.97</b>	<b>0.98</b>	<b>0.96</b>

## Results and discussion

To demonstrate the effectiveness of the proposed method in detecting anomalies within multivariate time-series, the performance of the proposed G-BOCPD method has been compared against the performance of six time-series analysis methods, obtained from the literature. These methods include 1NN-DTW (Xi et al., 2006), SAX-VSM (Senin & Malinchik, 2013), Shapelet forests (Patri et al., 2014), MC-DCNN (Zheng et al., 2014), FDC-CNN (Lee et al., 2017) and MTS-CNN (Hsu & Liu, 2021), as discussed in Sect. 1. The evaluation employed precision, recall and F1-score metrics, as discussed in (Hsu & Liu, 2021), to assess the performance of the proposed method through sixfold cross-validation. Tables 5, 6 and 7 summarise the performance comparisons for anomaly detection among the seven aforementioned methods in precision, recall and F1-score, employing data from Engine 1, Engine 2, and Doors, respectively.

**Table 6** Engine 2 anomaly detection results

Methods	Precision	Recall	F1
1NN-DTW (Xi et al., 2006)	0.35	0.25	0.29
SAX-VSM (Senin & Malinchik, 2013)	0.84	0.42	0.56
Shapelet forests (Patri et al., 2014)	0.84	0.80	0.82
MC-DCNN (Zheng et al., 2014)	0.64	0.80	0.71
FDC-CNN (Lee et al., 2017)	0.89	0.78	0.83
MTS-CNN (Hsu & Liu, 2021)	0.91	0.93	0.91
<b>G-BOCPD</b>	<b>0.98</b>	<b>0.97</b>	<b>0.97</b>

Following the anomaly detection results of Engine 1 in Table 5, it can be seen that Shapelet Forests, FDC-CNN, MTS-CNN and G-BOCPD provide high precision with average values equal to or higher than 0.9. Additionally, it is observed that although the average precision results of SAX-VSM and MC-DCNN are very close ( $\sim 0.845$ ), the average recall of SAX-VSM is significantly lower than MC-DCNN by 51%. In terms of the average precision and recall, MC-DCNN, FDC-CNN, MTS-CNN and the proposed G-BOCPD provide better results than 1NN-DTW, SAX-VSM, and Shapelet Forests. Overall, considering the results of Engine 1, the proposed G-BOCPD outperforms the other six methods with high average precision (0.96), recall (0.98) and F1 (0.97). It is also observed that 1NN-DTW would be the least preferred method as the values of average precision, recall and F1-score are the lowest among the other methods. It is worth mentioning that similar observations were made for the results obtained from the analysis of Engine 2 (Table 5) and Doors (Table 6) datasets.

Moreover, in terms of the anomaly detection results of Engine 2 in Table 6, it is seen that only MTS-CNN and the proposed G-BOCPD have an average precision above 0.9. Moreover, in terms of the average precision and recall, Shapelet Forests, FDC-CNN, MTS-CNN and G-BOCPD provide better performance compared to 1NN-DTW, SAX-VSM and MC-DCNN. Overall, the anomaly detection results for Engine 2 demonstrate that the proposed G-BOCPD outperforms the other six methods, achieving high average precision (0.98), recall (0.97) and F1-score (0.97), consistent with the observations from the results of Engine 1.

**Table 7** Doors anomaly detection results

Methods	Precision	Recall	F1
INN-DTW (Xi et al., 2006)	0.61	0.42	0.49
SAX-VSM (Senin & Malinchik, 2013)	0.81	0.60	0.69
Shapelet forests (Patri et al., 2014)	0.84	0.90	0.87
MC-DCNN (Zheng et al., 2014)	0.84	<b>0.96</b>	0.90
FDC-CNN (Lee et al., 2017)	0.86	0.94	0.90
MTS-CNN (Hsu & Liu, 2021)	0.90	0.92	<b>0.91</b>
<b>G-BOCPD</b>	<b>0.94</b>	0.88	<b>0.91</b>

Additionally, similar to Engine 2, the results of Doors in Table 5 show that only MTS-CNN and the proposed G-BOCPD provide highly precise results, equal to or higher than 0.9. It is interestingly observed in Table 5 that the proposed G-BOCPD outperforms the other six methods with high average precision (0.94) and F1-score (0.91). It is also noticed that MTS-CNN and the proposed G-BOCPD have the highest F1-score values, equal to 0.91, compared to the other methods. However, the average recall of the proposed G-BOCPD is 0.88, lower than the recall values of Shapelet Forests, MC-DCNN, FDC-CNN and MTS-CNN. The highest average recall value is obtained from MC-DCNN and equals 0.96. In conclusion, G-BOCPD-based tests reveal that the average precision, recall and F1-score values are the highest for the datasets of Engine 1 and Engine 2, being consistently above 95%, compared to the Doors dataset in which a lower F1-score can be seen (88%). One of the main reasons for this lower recall value for the Doors dataset is that the proposed G-BOCPD poorly segments the data, leading to inefficiency in the clustering. This limitation will be further explored in future research.

## Conclusion

This paper has proposed a novel graphical model-based annotation method, called G-BOCPD, for automatically detecting anomalies in multi-channel multivariate time-series data. The novelty of G-BOCPD lies in the hybridisation of the graphical lasso and expectation maximisation algorithms that enable capturing both internal and external dependencies within the data. This hybrid approach, which addresses the challenge of detecting anomalies by considering the complex relationships between variables, identifies segments in multivariate time-series that exhibit different behaviour or patterns. These segments are then annotated to highlight the variations. The proposed method alternates between estimating the concen-

tration matrix, which represents the dependencies between variables, using the graphical lasso algorithm, and annotating the segments using a minimal path clustering method that allows for a comprehensive understanding of variations within the data.

Two real-world engineering datasets have been used to validate the G-BOCPD method. The first dataset focused on DMU train engines with faulty behavioural patterns while the second dataset involved train doors' data with different utilisation rates that resulted in varying degradation stages. To ensure a comprehensive representation of system dynamics, relevant features such as overshoot, gradients, and steady-state standard deviations were extracted from both datasets. These features provided valuable information about the behaviour and characteristics of the systems under observation. The application of the G-BOCPD method to these case studies demonstrated its practicality and effectiveness in detecting anomalies in diverse scenarios within the engineering domain. By leveraging the combination of graphical modelling techniques and clustering methods, the G-BOCPD method proved capable of automatically capturing complex dependencies and variations within the time-series data.

Quantitative experiments have been conducted to compare the performance of G-BOCPD with alternative time-series methods in terms of precision, recall, and F1-score. The results obtained through sixfold cross-validation on three datasets (i.e. Engine 1, Engine 2, and Doors) have shown that G-BOCPD achieved competitive results, indicating its effectiveness in anomaly detection compared to existing approaches in the field. Across the three datasets, G-BOCPD consistently achieved high average precision (0.96–0.98), recall (0.97), and F1-score (0.96–0.97). It outperformed other methods in terms of average precision and recall, with INN-DTW being the least preferred method. G-BOCPD and MTS-CNN provided highly precise results for the Doors dataset, but G-BOCPD had a lower average recall. Overall, G-BOCPD consistently achieved precision, recall, and F1-score values above 95% for Engine 1 and Engine 2 datasets. The Doors dataset showed a lower F1-score (88%) due to inefficient data segmentation by G-BOCPD, which adversely affected recall. Further research will address this limitation.

Further investigations can be undertaken to enhance the G-BOCPD method. A potential future direction could be to explore methods for estimating the number of clusters. The current approach assumes prior knowledge of the number of clusters, which may not always be available. Therefore, additional research could explore methods for automatically estimating the number of clusters by considering a penalty for each new cluster utilised. This would enable the method to handle cases where the number of clusters is unknown, improving its flexibility and adaptability.

**Acknowledgements** This work was supported by Cranfield's Engineering and Physical Sciences Research Council (EPSRC) Impact Accelerator Account block funding and Instrumentel.

**Author contributions** BN Literature review, methodology, investigation, data collection, experimentation and model building, writing. CL Literature review, methodology, investigation, data collection, experimentation, and model building. JAE Technical advising, review, writing.

**Code availability** The data that support the findings of this study are available from the corresponding author upon request.

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Alqahtani, A., Mohammed, A., & Xie, X. (2021). Deep time-series clustering: A review. *Electronics*, *10*, 3001.
- Beggel, L., Kausler, B. X., Schiegg, M., Pfeiffer, M., & Bischl, B. (2018). Time series anomaly detection based on shapelet learning. *Computational Statistics*, *34*, 945.
- Bignoumba, N., Mellouli, N., & Yahia, S. B. (2024). A new efficient alignment-driven neural network for mortality prediction from irregular multivariate time series data. *Expert Systems with Applications*, *238*, 122148.
- Boyd, S. (2010). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, *3*(1), 1–122.
- Canizo, M., Triguero, I., Conde, A., & Onieva, E. (2019). Multi-head CNN-RNN for multi-time series anomaly detection: An industrial case study. *Neurocomputing*, *363*, 246–260.
- Chen, N., Tu, H., Duan, X., Hu, L., & Guo, C. (2023). Semisupervised anomaly detection of multivariate time series based on a variational autoencoder. *Applied Intelligence*, *53*(5), 1–25.
- Chen, S. X., Ni, Y. Q., Liu, J. C., & Yao, N. (2019). Deep learning-based data anomaly detection in rail track inspection, pp. 3235–3242. In *Structural Health Monitoring 2019: Enabling Intelligent Life-Cycle Health Management for Industry Internet of Things (IIOT) - proceedings of the 12th international workshop on structural health monitoring*, Vol. 2. DEStech Publications, Inc.
- Chen, Z., Chen, D., Zhang, X., Yuan, Z., & Cheng, X. (2022). Learning graph structures with transformer for multivariate time-series anomaly detection in IoT. *IEEE Internet of Things Journal*, *9*(12), 9179–9189.
- Choi, K., Yi, J., Park, C., & Yoon, S. (2021). Deep learning for anomaly detection in time-series data: Review, analysis, and guidelines. *IEEE Access*, *9*, 120043–120065.
- Danaher, P., Wang, P., & Witten, D. M. (2014). The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *76*(2), 373–397.
- Deng, A., & Hooi, B. (2021). Graph neural network-based anomaly detection in multivariate time series. In *AAAI conference on artificial intelligence*.
- Du, H., Li, Z., Chen, R., Yin, Z., Fu, Z., Zhang, Q., Xiao, X., Luo, M., Bao, F. (2019). Dynamic time warping and spectral clustering based fault detection and diagnosis of railway point machines, pp. 595–600. In *2019 IEEE Intelligent Transportation Systems Conference, ITSC 2019*. IEEE.
- Fang, Y., Ren, K., Shan, C., Shen, Y., Li, Y., Zhang, W., Yu, Y., & Li, D. (2023). Learning decomposed spatial relations for multivariate time-series modeling. *Proceedings of the AAAI Conference on Artificial Intelligence*, *37*(6), 7530.
- Gangsar, P., & Tiwari, R. (2019). A support vector machine based fault diagnostics of induction motors for practical situation of multi-sensor limited data case. *Measurement*, *135*, 694–711.
- Goswami, U., Rani, J., Kodamana, H., Kumar, S., & Tamboli, P. K. (2023). Fault detection and isolation of multi-variate time series data using spectral weighted graph auto-encoders. *Journal of the Franklin Institute*, *360*(10), 6783–6803.
- Hallac, D., Vare, S., Boyd, S. & Leskovec, J. (2018). Toeplitz inverse covariance-based clustering of multivariate time series data, pp. 5254–5258. In *Proceedings of the twenty-seventh international joint conference on artificial intelligence*, Vols. 2018-July. International Joint Conferences on Artificial Intelligence Organization.
- Hamilton, J. D. (2020). *Time series analysis*. Princeton University Press.
- Hsu, C.-Y., & Liu, W.-C. (2021). Multiple time - series convolutional neural network for fault detection and diagnosis and empirical study in semiconductor manufacturing. *Journal of Intelligent Manufacturing*, *32*(3), 823–836.
- Huang, Z., Zhu, J., Lei, J., Li, X., & Tian, F. (2020). Tool wear predicting based on multi - domain feature fusion by deep convolutional neural network in milling operations. *Journal of Intelligent Manufacturing*, *31*(4), 953–966.
- James, A., Jie, W., Xulei, Y., Chenghao, Y., Ngan, N. B., Yuxin, L., Yi, S., Chandrasekhar, V., & Zeng, Z. (2019). TrackNet - A deep learning based fault detection for railway track inspection, pp. 1–5. In *2018 International Conference on Intelligent Rail Transportation, ICIRT 2018*. IEEE.
- Kim, J., Kang, H., & Kang, P. (2023). Time-series anomaly detection with stacked transformer representations and 1D convolutional network. *Engineering Applications of Artificial Intelligence*, *120*, 105964.
- Lee, K. B., Cheon, S., & Kim, C. O. (2017). A convolutional neural network for fault classification and diagnosis in semiconductor manufacturing processes. *IEEE Transactions on Semiconductor Manufacturing*, *30*, 135–142.
- Li, W., Wenbo, Hu., Chen, T., Chen, N., & Feng, C. (2022). StackVAE-G: An efficient and interpretable model for time series anomaly detection. *AI Open*, *3*(July), 101–110.
- Lin, C.-J., Chu, W.-L., Wang, C.-C., Chen, C.-K., & Ting Chen, I. (2019). Diagnosis of ball-bearing faults using support vector machine based on the Artificial Fish-Swarm Algorithm. *Journal of Low Frequency Noise, Vibration and Active Control*, *1*, 146134841986182.
- Lu, C., & Shen, Y. (2020). Research on location detection of railway track spike based on neural network. *IOP Conference Series: Materials Science and Engineering*, *711*(1), 012019.

- Manco, G., Ritacco, E., Rullo, P., Gallucci, L., Astill, W., Kimber, D., & Antonelli, M. (2017). Fault detection and explanation through big data analysis on sensor streams. *Expert Systems with Applications*, 87, 141–156.
- Namoano, B., Emmanouilidis, C., Ruiz-Carcel, C., & Starr, A. G. (2020a). change detection in streaming data analytics: A comparison of Bayesian online and martingale approaches. *IFAC-PapersOnLine*, 53(3), 336–341.
- Namoano, B., Starr, A., Emmanouilidis, C., & Carcel, C. R. (2020a). Doors datasets.
- Namoano, B., Starr, A., Emmanouilidis, C., & Carcel, C. R. (2020b). Engines datasets.
- Patri, O. P., Sharma, A. B., Chen, H., Jiang, G., Panangadan, A. V., & Prasanna, V. K. (2014). Extracting discriminative shapelets from heterogeneous sensor data, pp. 1095–1104. In *2014 IEEE international conference on Big Data (Big Data)*.
- Senin, P., & Malinchik, S. (2013). SAX-VSM : Interpretable time series classification using SAX and vector space model. In *2013 IEEE 13th international conference on data mining*, pp. 1175–80.
- Tang, Z., Chen, Z., Bao, Y., & Li, H. (2019). Convolutional neural network-based data anomaly detection method using multiple information for structural health monitoring. *Structural Control and Health Monitoring*, 26, 1.
- Tao, C., & Du, J. (2023). PointSGRADE: Sparse learning with graph representation for anomaly detection by using unstructured 3D point cloud data. *IJSE Transactions*, 1, 1–14.
- Wu, Z., Pan, S., Long, G., Jiang, J., Chang, X., & Zhang, C. (2020). Connecting the dots: Multivariate time series forecasting with graph neural networks, pp. 753–763. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining, KDD '20*. Association for Computing Machinery.
- Xi, X., Keogh, E., Shelton, C., Wei, L., & Ratanamahatana, C. A. (2006). Fast time series classification using numerosity reduction, pp. 1033–1040. In *Proceedings of the 23rd international conference on machine learning, ICML '06*. Association for Computing Machinery.
- Xie, Y., Liu, Y., & Valdar, W. (2016). Joint estimation of multiple dependent Gaussian graphical models with applications to mouse genomics. *Biometrika*, 103(3), 493–511.
- Xu, M. H., & Shao, H. (2012). Solving the matrix nearness problem in the maximum norm by applying a projection and contraction method. *Advances in Operations Research*, 2012, 1–15.
- Xu, Z., Yang, Y., & Gao, X. (2023). DCFF-MTAD: A multivariate time-series anomaly detection model based on dual-channel feature fusion. *Sensors (basel)*, 24(10), 3258.
- Yang, J. B., Nguyen, M. N., San, P. P., Li, X. L., & Krishnaswamy, S. (2015). Deep convolutional neural networks on multichannel time series for human activity recognition. *Ijcai, 2015*, 3995–4001.
- Yeh, C.-C., Zhu, Y., Ulanova, L., Begum, N., Ding, Y., Dau, H. A., Zimmerman, Z., Silva, D. F., Mueen, A., & Keogh, E. (2018). Time series joins, motifs, discords and shapelets: A unifying view that exploits the matrix profile. *Data Mining and Knowledge Discovery*, 32(1), 83–123.
- Yin, S., & Huang, Z. (2015). Performance monitoring for vehicle suspension system via fuzzy positivistic c-means clustering based on accelerometer measurements. *IEEE/ASME Transactions on Mechatronics*, 20(5), 2613–2620.
- Yuan, Z., Zhu, S., & Zhai, W. (2020). Damage detection of rail fastening system through deep learning and vehicle-track coupled dynamics, pp. 148–53. In *Lecture notes in mechanical engineering*.
- Zhao, H., Wang, Y., Duan, J., Huang, C., Cao, D., Tong, Y., Xu, B., Bai, J., Tong, J., & Zhang, Q. (2020). Multivariate time-series anomaly detection via graph attention network. In *Proceedings - IEEE international conference on data mining, ICDM 2020-Nov*, pp. 841–50.
- Zheng, Y., Liu, Q., Chen, E., Ge, Y., & Leon Zhao, J. (2014). Time series classification using multi-channels deep convolutional neural networks Bt - web-age information management, pp. 298–310, edited by F. Li, G. Li, S. Hwang, B. Yao, and Z. Zhang. Springer.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

# Multi-channel anomaly detection using graphical models

Namoano, Bernadin

2025-08-01

Attribution 4.0 International

---

Namoano B, Latsou C, Erkoyuncu JA. (2024) Multi-channel anomaly detection using graphical models. *Journal of Intelligent Manufacturing*, Volume 36, July 2024, pp. 4319-4330

<https://doi.org/10.1007/s10845-024-02447-7>

*Downloaded from CERES Research Repository, Cranfield University*