Unsupervised representation learning for time series anomaly detection and beyond

Ferreira Correia; L.

Chapter 1

Introduction

1.1 Background

As the digitisation of industrial processes progresses, larger and larger datasets result from the data recorded. Ensuring this data is representative of the process is important, since downstream tasks like modelling or optimisation can be negatively impacted by incomplete or contaminated data. For tasks that require system behaviour modelling, data deviating from the norm is hence undesired, and we speak of *anomalous* behaviour. Recorded data manifests itself in many forms depending on the application and domain, one form being time series. Examples of real-world time series applications are diverse, ranging from the medical field [1] and server machines [2] to water distribution [3] and treatment [4]. Beyond that, the automotive industry also presents a real-world setting which can benefit from time series anomaly detection. Cars are complex dynamic systems whose behaviour can be described in many ways, one of which being the longitudinal dynamics, which is largely dictated by the powertrain within the vehicle. Automotive powertrains consist of multiple components, which contribute to a wide range of transient behaviour, including highly dynamic and more static signals.

1.2 Research Questions

This work attempts to answer a number of research questions relevant to the state of the art in the field of time series anomaly detection. Below, each question is posed and subsequently elaborated on, while being directly addressed in the following chapters.

Research Questions

RQ1 Can a non-trivial multivariate time series dataset be generated from physically motivated simulation models?

Measurable progress in the research field of unsupervised time series anomaly detection has been difficult to quantify over the last few years. While significant research volume has flowed into creating ever more complex data-driven models to approach the task, comparatively little work has been dedicated into benchmarking. Within benchmarking, publicly available data plays a large role, as it allows for the independent verification and reproduction of results. This contrasts with data stemming from real-world industrial processes, for instance, which often subject to strict levels of confidentiality and hence cannot not be published. Unfortunately, however, publicly available datasets have shown to be ill-suited due to several flaws and its distance from real-world problems in terms of dataset size and diversity. To answer this research question, the viability of physically inspired models for dataset generation is investigated in order to provide a publicly available dataset with real world-like complexity.

RQ2 Are satisfactory results obtainable when detecting anomalous behaviour in multivariate time series using completely unsupervised methods?

Although many studies in time series anomaly detection claim to propose unsupervised methods, only the model training is performed without labels, while adjacent processes like threshold or window size choice are informed with some amount of labelled data. In the real world, labelled data is either not available or expensive to obtain, and hence truly unsupervised methods are required for deployment. Clearly, proposed methods that use labelled data do not qualify as truly unsupervised and cannot as easily be deployed in the real world. To answer this research question, a novel, completely unsupervised method is benchmarked against the state of the art from literature to better understand its effectiveness in the real world.

RQ3 How well can a time series anomaly detection approach generalise to unseen data stemming from the same dynamic system?

In the real world, training data may not capture all possible scenarios due to very recent start of data collection, for example. Additionally, it may even not be possible to cover all scenarios due to random external effects, say for instance, caused by driving a different route to a destination every time with different levels of traffic and weather conditions. As is evident, for time series anomaly detection

in such settings, approaches that can generalise based on limited training data are desirable, and hence this aspect is further investigated.

RQ4 Can active learning be harnessed to find a suitable threshold for completely unsupervised methods?

Contamination of the training subset with anomalous data can pose major challenges for truly unsupervised time series anomaly detection approaches. Not only does the modelling process become more difficult, but so does the threshold choice. Truly unsupervised approaches are observed to use a far-from-ideal threshold due to the presence of anomalous data and hence the use of active learning, i.e. intelligently querying a domain expert to label data, is investigated to improve the threshold choice.

RQ5 How robust are active learning-based thresholds to mislabelling by the oracle?

In the real world, domain experts acting as oracles may still make mistakes when labelling time series data. Therefore, different mislabelling rates are experimented with to simulate such cases in an effort to study the impact on the anomaly detection performance.

1.3 Structure of this Thesis and Publications

This thesis is structured as follows. First, Chapter 2 provides the reader with relevant background in the field of time series anomaly detection, as well as with the theory behind key concepts discussed in this work. Following that, Chapter 3 is largely based on a survey [5] of the research area. It not only introduces methods proposed for time series anomaly detection but also provides a perspective on the state of benchmarking in the field. Chapter 4 addresses RQ1 by presenting a publicly available dataset [6] aiming to replicate real-world complexity with its physically motivated nature. Furthermore, a real-world dataset is also introduced along with some preliminary analysis revealing key properties within both datasets. Chapter 5 then provides the reader with a range of experiments and the associated results representing the viability of online and unsupervised detection methods and their capability to generalise, therefore addressing RQ2 and RQ3, respectively. This chapter is based on two publications. One paper [7] proposes a novel unsupervised data-driven model, tested on an offline real-world use-case, and another [8] compliments the findings by extending the problem to an online scenario and proposing a novel set of metrics for evaluation. Then, Chapter 6 is mostly based on a publication [9] which further addresses some of the issues

identified in Chapter 5 revolving around threshold choice and attempts to provide an answer to RQ4 and RQ5 by exploring the combination of active learning with unsupervised anomaly detection methods. Furthermore, another survey [10] under review identifies a research gap within predictive maintenance, which is briefly touched upon in Chapter 7. It illustrates how modelling dynamic system behaviour in a data-driven way can be applied beyond anomaly detection to fields like predictive maintenance in the case of a real-world problem. Finally, Chapter 8 summarises the main body of this thesis and concisely answers the research questions posed above, while also outlining potential future work. The full citations for the publications included in this thesis are provided below.

[5] Lucas Correia, Jan-Christoph Goos, Philipp Klein, Thomas Bäck, and Anna V. Kononova. "Online model-based anomaly detection in multivariate time series: Taxonomy, survey, research challenges and future directions". In: *Engineering Applications of Artificial Intelligence* 138 (2024), p. 109323. DOI: 10.1016/j.engappai.2024.109323

[6] Lucas Correia, Jan-Christoph Goos, Thomas Bäck, and Anna V. Kononova. "PATH: A Discrete-sequence Dataset for Evaluating Online Unsupervised Anomaly Detection Approaches for Multivariate Time Series". In: *Big Data Research* (2025), p. 100573. DOI: 10.1016/j.bdr.2025.100573

[7] Lucas Correia, Jan-Christoph Goos, Philipp Klein, Thomas Bäck, and Anna Kononova. "MA-VAE: Multi-Head Attention-Based Variational Autoencoder Approach for Anomaly Detection in Multivariate Time-Series Applied to Automotive Endurance Powertrain Testing". In: *Conference on Neural Computation Theory and Applications*. SciTePress, 2023, pp. 407–418. DOI: 10.5220/0012163100003595

[8] Lucas Correia, Jan-Christoph Goos, Philipp Klein, Thomas Bäck, and Anna V. Kononova. "TeVAE: A Variational Autoencoder Approach for Discrete Online Anomaly Detection in Variable-State Multivariate Time-Series Data". In: *Computational Intelligence*. Vol. 1196. Series Title: Studies in Computational Intelligence. Springer Nature Switzerland, 2025, pp. 106–132. DOI: 10.1007/978-3-031-85252-7_7

[9] Lucas Correia, Jan-Christoph Goos, Thomas Bäck, and Anna V. Kononova. "DQS: A Low-Budget Query Strategy for Enhancing Unsupervised Data-driven

Anomaly Detection Approaches". In: *Journal of Big Data* (2026). DOI: 10.1186/s40537-026-01524-3.

[10] Bernd G Wagner, Qi Huang, Lucas Correia, Thomas Bäck, Anna Kononova, and Niki Van Stein. *REMAINING USEFUL LIFE IN COMPLEX MULTI-COMPONENT SYSTEMS: TAXONOMY, REVIEW, AND RESEARCH DIRECTIONS*. 2025. DOI: 10.36227/techrxiv.176288069.90065228/v1.