ADEPT: Anomaly Detection, Explanation and Processing for Time Series with a Focus on Energy Consumption Data

Benedikt Tobias Müller, Marvin Ender, Jan Erik Swiadek(✉), Mengcheng Jin, Simon Winkel, Dominik Niedziela, Bin Li, Jelle Hüntelmann, and Emmanuel Müller

TU Dortmund University, Dortmund, Germany
jan-erik.swiadek@tu-dortmund.de

Abstract. Anomaly detection techniques are applicable for recognizing excessive energy consumption and device failure, thereby contributing to the maintenance of operational and sustainable energy supply systems. In this context, human decision makers can benefit from receiving explanation attempts for detected anomalies as part of a semi-automated software solution. Therefore we introduce the framework ADEPT, which comprises interfaces for processing user-supplied time series data and interactively visualizing explanatory anomaly information. Our framework features several shallow and deep machine learning algorithms for anomaly detection and explanation. We demonstrate ADEPT using energy consumption data collected from our university campus.

Keywords: Anomaly detection · Energy consumption · Explainability

1 Introduction

Conserving and efficiently utilizing energy are essential principles in sustainable development. They also reduce cost and afford greater autonomy in times of scant or uncertain energy supply. However, it is challenging to monitor energy consumption for large-scale systems because of the complexities inherent to visualizing and understanding high-dimensional time series data. Anomalous events can arise in various patterns and across all subsets of the deployed sensors, so they are hard to detect and interpret even for domain experts. On the technical side, recent methods show convincing performance in anomaly detection [1], but fully-automated solutions can't make use of human domain knowledge and lack in transparent explanations. This creates demand for semi-automated solutions offering intuitive and trustworthy anomaly explanations, which can help humans reliably find, e.g., periods of unusual consumption or sensor defects.

For the reasons stated, anomaly detection and explanation are challenging machine learning problems. Some already existing tools and frameworks make pertinent research results available to users, though they primarily focus on different application domains, like eX^2 [2] for cybersecurity, DAART [8] for military purposes, and EXAD [9] for big data tracing. Other options for explainable anomaly detection include MSDA [3], a library targeting data scientists, and VADETIS [4], which does not go beyond comparative explanations. We thus propose ADEPT, a novel interactive framework providing easily accessible anomaly explanations in addition to energy consumption monitoring for homeowners and facility managers maintaining energy supply systems. Furthermore, our framework enables the research community to assess and compare machine learning methods concerning anomaly detection and explainability.

2 Framework Overview

Because our target audience is partly non-technical, a simple and easy-to-use way of comfortably and quickly assessing anomalies is needed. For this purpose we chose a microservice-based web application. That means users can access the application from every device, no calculations need to take place on the device itself, and the framework is extensible. A demonstration video of the ADEPT web application is available here: https://youtu.be/Uk28ipbJGiY

2.1 Flexible Architecture

The framework is composed of a multitude of microservices, one for each core component. This allows researchers and developers to quickly and independently test innovative techniques, simply by adapting the appropriate component of the framework. Due to ADEPT's dynamic and flexible design, data can be read from many different sources, like real-time sensor data or static data files. These design principles extend to all parts of the software including the machine learning pipeline, which consists of normalization, feature engineering and model training as well as the detection, explanation and visualization of anomalies.

Fig. 1. Screenshot of the ADEPT web interface

2.2 Interactive Exploration

ADEPT enables users to easily browse through and analyze existing data for anomalies without prior knowledge in data science. A screenshot of the ADEPT web interface can be found in Fig. 1. The first row contains the raw data and a configuration panel. Here users can select their desired sensors, features, a time period, and an anomaly detection algorithm. Detected anomalies are displayed in the second row. The table on the left side presents the timestamps and types of the anomalies, while the diagram on the right side features calculated anomaly likelihood scores with a threshold. Selecting one of the anomalies from the table fills the bottom row with corresponding explanation results, allowing users to analyze the anomaly in depth. The bottom left tile depicts an example-based explanation of the feature most responsible for the selected anomaly, contrasting the anomaly with normal patterns. This gives users immediate feedback on the shape of the anomaly. The bottom right tile displays the feature attribution of the anomaly, i.e., how much each feature contributes to its anomaly score.

3 Detection Models and Explainability Challenges

Each anomaly detection model comes with different strengths and weaknesses, which is why we need several models to detect all kinds of anomalies reliably. On a high level, models are categorized as shallow or deep models. Shallow models provide high training efficiency and are often sufficient for detecting simple anomalies, such as low-dimensional or extreme-value anomalies. Thanks to their simplicity, explanations can usually be derived directly. Meanwhile, explaining the results of deep models requires more sophisticated techniques, but their enhanced detection performance is essential for our use case because of the typically complex anomaly patterns hidden in high-dimensional energy consumption data. Among others, ADEPT makes use of the shallow algorithms Isolation Forest [5] and One-Class SVM [7] as well as the deep model LSTM Autoencoder [6]. Each model returns anomaly likelihood scores for all measurement timestamps. It is possible to train all models on-demand with user-supplied data. Shallow models can also be fitted using real-time data from the machine learning pipeline.

To facilitate human understanding and trust in the models, ADEPT includes explanation techniques that highlight anomalies by presenting normal patterns for comparison. We identify three key challenges in this regard. Firstly, a uniform normal state might not exist, as the meaning of "normal" not only depends on the respective time series, but can also vary over the course of a single time series, e.g., due to concept drifts. Secondly, the interplay of multiple time series can obfuscate the normal state and result in very complex anomalies. For instance, a slight deviation from a correlation involving a large number of dimensions is hardly noticeable. Thirdly, there are many approaches for extracting normal patterns and it has not yet been fully explored how intuitive their results are under different conditions. By implementing multiple methods for extracting normal patterns, ADEPT makes it possible to compare them. In order to find the dimension with the greatest contrast between a detected anomaly and the corresponding normal state, feature attributions are provided. In the case of LSTM Autoencoder, they are calculated using Integrated Gradients [10].

4 Evaluation and Future Work

For evaluating the detection performance of our models we use data from TU Dortmund University that consists of electricity, heat and water consumption measurements across more than 40 buildings on the university campus. This data is provided to us by the facility managers of our university, who also collaborate with us in interpreting normal states and abnormal events as domain experts. Thus far, ADEPT helped us find many events in the campus data that were confirmed by the facility managers as known anomalies, but also a few previously unknown, more subtle occurrences. These findings might enable them to optimize the energy efficiency of some buildings with abnormal energy consumption.

Considering the challenges we laid out before, there are still limits to the capabilities of ADEPT. At the same time, this creates opportunities for future research, in which our framework could help with assessing and comparing anomaly detection and explanation methods. Moreover, we plan to conduct regular stakeholder meetings for discussing its usage in the decision making process regarding the sustainability goals of our university's energy supply system.

Acknowledgements. This work was supported by the Research Center Trustworthy Data Science and Security, an institution of the University Alliance Ruhr.

References

1. Aggarwal, C.C.: Outlier Analysis. Springer, 2 edn. (2017)
2. Arnaldo, I., Veeramachaneni, K., Lam, M.: eX2: a framework for interactive anomaly detection. In: Joint Proceedings of the ACM IUI 2019 Workshops (2019)
3. Arunachalam, A.: MSDA (2021), https://pypi.org/project/msda/
4. Khelifati, A., Khayati, M., Cudré-Mauroux, P., Hänni, A., Liu, Q., Hauswirth, M.: VADETIS: an explainable evaluator for anomaly detection techniques. In: 37th IEEE International Conference on Data Engineering. pp. 2661–2664 (2021)
5. Liu, F.T., Ting, K.M., Zhou, Z.: Isolation forest. In: Proceedings of the 8th IEEE International Conference on Data Mining. pp. 413–422 (2008)
6. Malhotra, P., Ramakrishnan, A., Anand, G., Vig, L., Agarwal, P., Shroff, G.: LSTM-based encoder-decoder for multi-sensor anomaly detection. CoRR (2016), https://arxiv.org/abs/1607.00148
7. Schölkopf, B., Williamson, R.C., Smola, A.J., Shawe-Taylor, J., Platt, J.C.: Support vector method for novelty detection. In: Advances in Neural Information Processing Systems 12. pp. 582–588 (1999)
8. Smith-Renner, A., Rua, R., Colony, M.: Towards an explainable threat detection tool. In: Joint Proceedings of the ACM IUI 2019 Workshops (2019)
9. Song, F., Diao, Y., Read, J., Stiegler, A., Bifet, A.: EXAD: A system for explainable anomaly detection on big data traces. In: ICDMW 2018. pp. 1435–1440 (2018)
10. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: Proceedings of the 34th International Conference on Machine Learning. pp. 3319–3328 (2017)