INTIMEAD: Interactive Time Series Anomaly Detection

Louis Carpentier[1, 2, 3], **Wannes Meert**[1, 2], **Mathias Verbeke**[1, 2, 3]

[1]Department of Computer Science, KU Leuven, Belgium
[2]Leuven.AI - KU Leuven Institute for AI, Belgium
[3]Flanders Make@KU Leuven, Belgium
louis.carpentier@kuleuven.be, wannes.meert@kuleuven.be, mathias.verbeke@kuleuven.be

Abstract

Time series anomaly detection has received substantial attention over the past two decades, leading to the development of hundreds of algorithms. However, comprehensively understanding this vast landscape remains challenging, particularly for non-experts and novices. In this demonstration paper, we present INTIMEAD, an interactive web application that provides access to more than 30 state-of-the-art time series anomaly detection algorithms. INTIMEAD is intended to explore the performance of existing as well as custom anomaly detection models in an interactive, hands-on manner. By lowering the entry bar, we support practitioners overwhelmed by the large number of existing techniques, while providing a platform for researchers to rapidly analyze their novel anomaly detection algorithms.

INTIMEAD – https://intimead.cs.kuleuven.be
Code – https://github.com/ML-KULeuven/InTimeAD

Introduction

Time series anomaly detection (TSAD) is the automatic detection of unusual or anomalous sequences in time series data (Liu and Paparrizos 2024) and plays a crucial role in many domains, such as industry (Fingerhut et al. 2023), healthcare (Samariya et al. 2023), finance (Ahmed, Choudhury, and Uddin 2017), social networks (Savage et al. 2014), and environmental sciences (Russo et al. 2020). Due to its importance, hundreds of TSAD algorithms have been proposed over the past two decades (Schmidl, Wenig, and Papenbrock 2022). However, comprehensively understanding this vast landscape remains challenging, particularly for non-experts and novices. This difficulty is further amplified by the fact that algorithm performance heavily depends on the time series characteristics (Lemmens et al. 2024).

In this demo paper, we present INTIMEAD, a web application to facilitate the analysis of more than 30 state-of-the-art anomaly detection algorithms in an interactive, hands-on manner. Our goal is to improve the understanding of state-of-the-art anomaly detection methods while minimizing the effort required from users. To achieve this, we (1) build on top of DTAIANOMALY (Carpentier et al. 2025) to provide access to a large set of state-of-the-art models; (2) integrate dynamic visualizations that support an intuitive understanding in the predictions beyond numerical scores; (3) allow flexible configuration of hyperparameters to analyze their impact; and (4) support multiple options for loading time series data, including custom datasets. INTIMEAD is online available at https://intimead.cs.kuleuven.be.

Impact on Practitioners As mentioned before, TSAD is applicable in a wide range of real-world domains. However, the extensive number of available models can be overwhelming for non-experts and novices. INTIMEAD provides access to a variety of state-of-the-art methods through an interactive interface, enabling practitioners to gain hands-on experience with minimal effort. Moreover, INTIMEAD allows practitioners to evaluate models directly on their own data, thereby bridging the gap between the algorithmic design and practical application.

Impact on Researchers TSAD remains an active research area, with new methods continuously being proposed. Accurately assessing the strengths and weaknesses of a novel algorithm is essential for improving its design and performance. INTIMEAD facilitates this process by allowing researchers to integrate their implementations and visualize the predictions across a large collection of time series with minimal effort. This provides an intuitive understanding of algorithmic performance, beyond mere numerical results. INTIMEAD also supports exploration of hyperparameter sensitivity through configurable parameters. Furthermore, new algorithms can be directly compared against all state-of-the-art detectors included in INTIMEAD, enabling rapid benchmarking and analysis.

Positioning to Related Work

Several open-source toolboxes collecting a comprehensive set of TSAD algorithms have been published recently (Carpentier et al. 2025; Middlehurst et al. 2024; Wenig, Schmidl, and Papenbrock 2022; Liu and Paparrizos 2024). While these toolboxes improve access to publicly available implementations, users still need to set up environments and write the code for using these models. INTIMEAD eliminates these steps, enabling immediate analysis of algorithm performance. Furthermore, INTIMEAD focuses on visual exploration, which is essential for building intuition about model behavior, but is often neglected in practice (Wu and

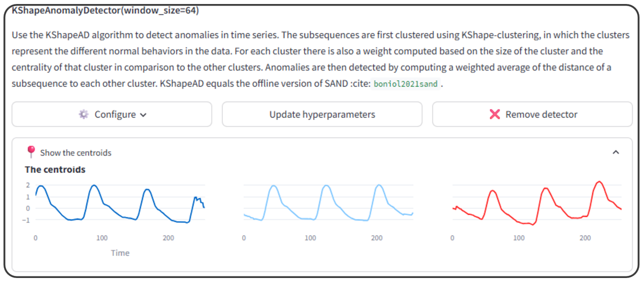

(a) Configuration of KShapeAD

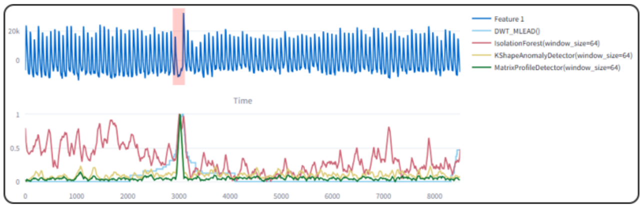

(b) Visual analysis of the the anomaly scores

(c) Numerical analysis of the anomaly detectors

Figure 1: Interface for key features of INTIMEAD

Keogh 2023). Code-based analysis often emphasizes numerical results, thereby limiting visual exploration.

Besides these open-source code bases, several user interfaces have been developed (Boniol et al. 2022, 2024; Liu, Lee, and Paparrizos 2025). These interfaces primarily focus on the global, average performance of a method based on benchmarking results. This approach restricts insights to what is available within the benchmarking results and prevents an in-depth analysis of individual anomaly detectors. INTIMEAD addresses this limitation by analyzing the models on a per-time series basis, enabling more detailed assessment of algorithm performance, including the performance of custom models and on user-provided time series. Moreover, configurable hyperparameters allow researchers and practitioners to explore sensitivity in INTIMEAD and understand how parameter choices affect performance.

INTIMEAD

We first describe the main flow of INTIMEAD, followed by a summary of additional features.

Users begin by loading a time series, which can be inspected immediately. Next, one or more anomaly detectors are selected for analysis (see Figure 1a). Hyperparameters are initialized with default values but can also be manually configured. The initial analysis visualizes the predicted anomaly scores (see Figure 1b) and indicates which methods accurately detect anomalies in the time series. In addition to this visual analysis, INTIMEAD provides a numeric comparison of detectors to support objective selection of the best-performing method (see Figure 1c). For this, users can choose from a variety of evaluation metrics, which are fully configurable (e.g., β in F_β-based metrics or thresholds on anomaly scores).

Code Snippets After performing an initial analysis with INTIMEAD, a more in-depth investigation is often needed. The interface provides code snippets that can be directly copied and executed, facilitating the transition from INTIMEAD to a code-based analysis in DTAIANOMALY.

Model-Specific Visualizations Certain models provide additional insights beyond the raw anomaly scores. For instance, KSHAPEAD (Paparrizos and Gravano 2017) clusters subsequences using centroids and computes anomaly scores based on the distance to the nearest centroid. Visualizing these centroids reveals the expected normal behavior, offering an intuitive explanation of the resulting anomaly scores. INTIMEAD supports such model-specific visualizations, as illustrated in Figure 1a.

Extensibility INTIMEAD can be extended with new anomaly detectors when run locally. To do so, users implement a class `MyDetector` with `.fit(X, y)` and `.decision_function(X)` methods, and simply tell INTIMEAD to include the custom class:

```
import InTimeAD
InTimeAD.run(custom_anomaly_detectors=MyDetector)
```

Custom data loaders and evaluation metrics can also be integrated in a similar manner.

Analyzing Anomaly Detectors in INTIMEAD

INTIMEAD enables simple analysis of anomaly detectors without the need to write any code and provides insights within only a few minutes. Figures 1b and 1c compare four anomaly detectors on a time series from the UCR Anomaly Archive (ID 4): DWT-MLEAD (Thill, Konen, and Bäck 2017), IForest (Liu, Ting, and Zhou 2008), KShapeAD (Paparrizos and Gravano 2017), and Matrix Profile (Zhu et al. 2016). DWT-MLEAD, KShapeAD, and Matrix Profile accurately detect the anomaly, whereas IForest fails to do so. The event-wise F_1-score confirms this. Interestingly, raising the threshold from 0.9 to 0.95 increases IForest's F_1-score from 0.5 to 1, illustrating the sensitivity of the threshold. Threshold-independent metrics AUC-PR and VUS-PR (Paparrizos et al. 2022) further reveal that DWT-MLEAD performs better than the other two methods on this time series.

The ability to perform such rapid analyses is particularly important in real-world applications, facilitating the development of proof-of-concept solutions with minimal effort.

INTIMEAD: TSAD Made Simple

We have shown how state-of-the-art research in time series anomaly detection can be made accessible through an interactive web application. INTIMEAD provides a low-effort platform for hands-on exploration of a comprehensive collection of methods and supports the development and evaluation of new algorithms. In future work, we plan to link algorithm performance to time series characteristics, to explain the algorithms with reference to their invariances.

Acknowledgments

This research is supported by Flanders Innovation & Entrepreneurship (VLAIO) through the AI-ICON project CONSCIOUS (HBC.2020.2795), the Flemish government under the Flanders AI Research Program, and by Internal Funds KU Leuven (STG/21/057 & C3/23/047). We would also like to thank the researchers of the M-Group research group for their valuable comments and feedback.

References

Ahmed, M.; Choudhury, N.; and Uddin, S. 2017. Anomaly detection on big data in financial markets. In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, 998–1001.

Boniol, P.; Paparrizos, J.; Kang, Y.; Palpanas, T.; Tsay, R. S.; Elmore, A. J.; and Franklin, M. J. 2022. Theseus: navigating the labyrinth of time-series anomaly detection. *Proceedings of the VLDB Endowment*, 15(12): 3702–3705.

Boniol, P.; Sylligardos, E.; Paparrizos, J.; Trahanias, P.; and Palpanas, T. 2024. ADecimo: Model Selection for Time Series Anomaly Detection. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 5441–5444.

Carpentier, L.; Seeuws, N.; Meert, W.; and Verbeke, M. 2025. dtaianomaly: A Python library for time series anomaly detection. arXiv:2502.14381.

Fingerhut, F.; Klein, S.; Verbeke, M.; Rajendran, S.; and Tsiporkova, E. 2023. Multi-view contextual performance profiling in rotating machinery. In *2023 IEEE International Conference on Prognostics and Health Management (ICPHM)*, 343–350.

Lemmens, I.; Carpentier, L.; Meert, W.; and Verbeke, M. 2024. Evaluation of anomaly detection methods in time series using Instance Space Analysis. In *BNAIC/BeNeLearn, Date: 2024/11/18-2024/11/20, Location: Utrecht*.

Liu, F. T.; Ting, K. M.; and Zhou, Z.-H. 2008. Isolation Forest. In *2008 Eighth IEEE International Conference on Data Mining*, 413–422.

Liu, Q.; Lee, S.; and Paparrizos, J. 2025. EasyAD: A Demonstration of Automated Solutions for Time-Series Anomaly Detection. *Proceedings of the VLDB Endowment*, 18: 12.

Liu, Q.; and Paparrizos, J. 2024. The Elephant in the Room: Towards A Reliable Time-Series Anomaly Detection Benchmark. In *NeurIPS 2024*.

Middlehurst, M.; Ismail-Fawaz, A.; Guillaume, A.; Holder, C.; Guijo-Rubio, D.; Bulatova, G.; Tsaprounis, L.; Mentel, L.; Walter, M.; Schäfer, P.; and Bagnall, A. 2024. aeon: a Python Toolkit for Learning from Time Series. *Journal of Machine Learning Research*, 25(289): 1–10.

Paparrizos, J.; Boniol, P.; Palpanas, T.; Tsay, R. S.; Elmore, A.; and Franklin, M. J. 2022. Volume under the surface: a new accuracy evaluation measure for time-series anomaly detection. *Proc. VLDB Endow.*, 15(11): 2774–2787.

Paparrizos, J.; and Gravano, L. 2017. Fast and Accurate Time-Series Clustering. *ACM Trans. Database Syst.*, 42(2).

Russo, S.; Lürig, M.; Hao, W.; Matthews, B.; and Villez, K. 2020. Active learning for anomaly detection in environmental data. *Environmental Modelling & Software*, 134: 104869.

Samariya, D.; Ma, J.; Aryal, S.; and Zhao, X. 2023. Detection and explanation of anomalies in healthcare data. *Health Information Science and Systems*, 11(1): 20.

Savage, D.; Zhang, X.; Yu, X.; Chou, P.; and Wang, Q. 2014. Anomaly detection in online social networks. *Social networks*, 39: 62–70.

Schmidl, S.; Wenig, P.; and Papenbrock, T. 2022. Anomaly detection in time series: a comprehensive evaluation. *Proceedings of the VLDB Endowment*, 15(9): 1779–1797.

Thill, M.; Konen, W.; and Bäck, T. 2017. Time series anomaly detection with discrete wavelet transforms and maximum likelihood estimation. In *Intern. Conference on Time Series (ITISE)*, volume 2, 11–23.

Wenig, P.; Schmidl, S.; and Papenbrock, T. 2022. TimeEval: a benchmarking toolkit for time series anomaly detection algorithms. *Proc. VLDB Endow.*, 15(12): 3678–3681.

Wu, R.; and Keogh, E. J. 2023. Current Time Series Anomaly Detection Benchmarks are Flawed and are Creating the Illusion of Progress. *IEEE Transactions on Knowledge and Data Engineering*, 35(3): 2421–2429.

Zhu, Y.; Zimmerman, Z.; Senobari, N. S.; Yeh, C.-C. M.; Funning, G.; Mueen, A.; Brisk, P.; and Keogh, E. 2016. Matrix Profile II: Exploiting a Novel Algorithm and GPUs to Break the One Hundred Million Barrier for Time Series Motifs and Joins. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, 739–748.