

A Comprehensive Guide to Time-Series Anomaly Detection

John Paparrizos
The Ohio State University
Columbus, OH, USA
Aristotle University
Thessaloniki, Greece
paparrizos.1@osu.edu

Paul Boniol
Inria, DI ENS, PSL, CNRS
Paris, France
paul.boniol@inria.fr

Qinghua Liu
The Ohio State University
Columbus, OH, USA
liu.11085@osu.edu

Themis Palpanas
Université Paris Cité; IUF
Paris, France
themis@mi.parisdescartes.fr

Abstract

Anomaly detection is a fundamental data analytics task across scientific fields and industries. In recent years, an increasing interest has been shown in the application of anomaly detection techniques to time series. In this tutorial, we take a holistic view of anomaly detection in time series and comprehensively cover detection algorithms ranging from the 1980s to the most current state-of-the-art techniques. Importantly, the scope of this tutorial extends beyond algorithmic discussion, delving into the latest advancements in benchmarking and evaluation measures for this area. In particular, our interactive systems enable the exploration of methods and benchmarking results, thereby promoting user comprehension. Furthermore, this tutorial extensively explores automated solutions for unsupervised model selection, introduces a new taxonomy, and engages with the challenges and recent findings, particularly the difficulty for these solutions to outperform simple random choice. Driven by the limited generalizability of current detection algorithms, we review recent applications of foundation models for anomaly detection to motivate further research in the area.

CCS Concepts

• **Do Not Use This Code → Generate the Correct Terms for Your Paper;** *Generate the Correct Terms for Your Paper;* Generate the Correct Terms for Your Paper; Generate the Correct Terms for Your Paper.

Keywords

Time-Series Analysis; Anomaly Detection; Outlier Detection

ACM Reference Format:

John Paparrizos, Paul Boniol, Qinghua Liu, and Themis Palpanas. 2018. A Comprehensive Guide to Time-Series Anomaly Detection. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

A wide range of sensing, networking, and processing solutions enable the collection of large amounts of data over time [10, 14,

17, 24, 25, 28, 30, 31]. Analytical tasks over such ordered sequences of real-valued data, known as *time series*, are becoming increasingly important in virtually every domain. Anomaly detection has received ample academic and industrial attention over the past decades. *Anomalies* refer to data points that do not conform to some notion of normality or an expected behavior based on previously observed data. In practice, anomalies can correspond to erroneous data (e.g., broken sensors), or data of interest (e.g., anomalous behavior of the measured system) [1]. Detecting such cases is crucial for many applications [12]. Multiple surveys, benchmarks, and experimental studies summarize and analyze the state-of-the-art methods [3, 20, 29], exploring different aspects of the problem.

Relation to Previous Tutorials. Recent tutorials related to anomaly detection have focused (i) on specific types of methods (such as deep-learning methods [9]) ignoring several of the state-of-the-art methods, (ii) on specific types of temporal data (such as spatiotemporal data [38]), or (iii) on a much more general topic (only briefly discussing anomaly detection as a subpart of time series analysis [36]). Earlier versions of this tutorial, presented at EDBT 2023 [6], ICDE 2024 [7], VLDB2024 [16], KDD2025 [26] primarily focused on state-of-the-art time-series anomaly detection methods, their operational mechanisms, and a limited discussion on automated solutions and emerging foundation models.

In contrast, this tutorial undergoes significant expansion, offering a detailed explanation of recently proposed methodologies with interactive exploration for the audience. Beyond the introduction of anomaly detection techniques, this tutorial adopts a comprehensive perspective on the field, with a focus on evaluation measures and benchmark analysis. A brand new section has been introduced, outlining automated solutions with a taxonomy and evaluation results. Furthermore, the tutorial explores the latest advancements in applying foundation models to time-series analysis and anomaly detection, opening up new research opportunities. To our knowledge, this is the first tutorial that (i) extensively covers the time-series anomaly detection landscape; (ii) incorporates the latest developments in automated solutions and foundation models for TSAD; and (iii) elaborates on concepts through experimental studies and use of interactive demonstration systems.

Audience and Expected Background. This tutorial is for researchers and data analysts and will focus on recent advances in time series analysis and anomaly detection. The tutorial aims to initiate new collaborations between members of the data mining community and machine learning practitioners in various domains and increase the interest in anomaly detection for time series. We will present material that includes the necessary background to follow the entire presentation and technical details of the current

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXX.XXXXXXX>




Figure 1: Process-centric anomaly detection taxonomy.

state-of-the-art solutions for anomaly detection. We will also discuss the drawbacks of existing solutions and the open problems. Both experts and newcomers in the area will be able to follow the material and benefit from the tutorial.

2 Presenters

John Paparrizos is an assistant professor at The Ohio State University, leading The DATUM Lab. His research focuses on adaptive solutions for data-intensive and machine-learning applications. His doctoral work was recognized at the 2019 ACM SIGKDD Doctoral Dissertation Award competition. He has also received the inaugural ACM SIGMOD Research Highlight Award, a NetApp Faculty Award, and the 2023 IEEE TCDE Rising Star Award. His ideas have been widely adopted across scientific areas, Fortune 100-500 companies (e.g., Exelon and Nokia), and organizations such as ESA.

Paul Boniol is a researcher at Inria, member of the VALDA project-team. Previously, he worked at ENS Paris-Saclay (Centre Borelli), Université Paris Cité, EDF Research lab, and Ecole Polytechnique. His research interests lie between data analytics, machine learning, and time-series analysis. His Ph.D. thesis focused on subsequence anomaly detection and time-series classification. His work has been published in the top data management and analytics venues.

Qinghua Liu is a Ph.D. student at The Ohio State University and a member of The DATUM Lab. His research interests include time-series analysis and anomaly detection. He received the B.Eng. degree in Electronic Engineering from Tianjin University in 2022. Previously, he worked at Bosch, SRIID, CASIA, and JD.com Inc.

Themis Palpanas is a Senior Fellow of the French University Institute (IUF), and Distinguished Professor of computer science at Université Paris Cité (France). He has authored 15 patents, received 3 best paper awards and the IBM SUR award, has been Program Chair for VLDB 2025 and IEEE BigData 2023, General Chair for VLDB 2013, and has served Editor in Chief for BDR. He has been working in the fields of Data Series Management and Analytics for more than 15 years, and has developed several of the state of the art techniques. He has delivered 19 tutorials in top conferences.

3 Tutorial Scope

In this **3-hour lecture-style tutorial**, we will go through the problem of anomaly detection in time series, starting from fundamental definitions to open problems and opportunities.

- Introduction, Motivation and Foundations (30min).
- Taxonomy of Anomaly Detection Methods (45min).
- Evaluation Measures (30min).
- Anomaly Detection Benchmarks (25min).
- Automated Solutions for Anomaly Detection (25min).
- Conclusion and Open Problems (25min)

Support Material. Supplementary material, including slides from previous tutorials, are available at: <https://thedatumorg.github.io/TSAD-Tutorial>. For a comprehensive discussion of recent advances in time-series anomaly detection, we also refer readers to our recent surveys [5, 26].

Tutorial Overview. As the interest in Time-Series Anomaly Detection (TSAD) expands across scientific fields and industries, this tutorial provides a comprehensive view of this task. We start with definitions for time series and anomalies. Our goal is four-fold: (i) introduce the motivation related to the anomaly detection task in time series by describing different types of time-series anomaly and the taxonomy of anomaly detection algorithms proposed by different communities; (ii) describe recently proposed benchmarks and experimental evaluations along with the discussion on the challenges and problems inherent in the evaluation process; (iii) explore the latest trends in automated anomaly detection solutions for time series, providing a taxonomy and insights for each methodology alongside experimental investigations; (iv) discuss the challenges and opportunities for time-series anomaly detection including Foundation Model-empowered analysis.

This tutorial features (i) an extensive review of literature spanning several decades, from early methods of the 1980s to the latest state-of-the-art approaches, and developed in different communities, from data management to machine learning; (ii) insights into the recent advancements in automated solutions, enriching the research landscape with novel perspectives; (iii) interactive systems for exploring a method's computation steps and experimental evaluation, making complex concepts engaging. We hope this tutorial will equip the audience with a fundamental comprehension of TSAD and inspire them to work in this dynamic and evolving area.

3.1 Introduction, Motivation and Foundations

We will start by discussing examples of scientific and industrial applications that rely on TSAD.

Type of Time Series. We first introduce the different types of time series. Specifically, we define the time series as an ordered sequence of real values on one (for *univariate* time series) or multiple dimensions (for *multivariate* time series). Moreover, we define *static* and *streaming* time series as sequences with a fixed length or continuously arriving subsequences. Finally, the normal behavior (i.e., subsequences representing the normal and recurrent behavior) might change over time. In this case, we differentiate *single* and *multiple normalities* time series.

Type of Anomalies. We then introduce the different types of anomalies. The first two categories, *point* and *contextual* anomalies, refer to data points deviating remarkably from the rest of the data globally or given a specific context, respectively. The third category, *Collective* anomalies, corresponds to sequences of points that do not repeat a typical (previously observed) pattern. Moreover, their combination also matters. For instance, we need to differentiate time series containing *single* anomalies from time series containing *multiple* anomalies (either *similar* or *different*). For all these definitions and classes, we will provide explicit examples.

3.2 Taxonomy of Anomaly Detection Methods

We will dive into the different anomaly detection methods proposed in the literature. As many papers appear every year proposing new

methods for anomaly detection in time series based on different applications [3, 11], it is beyond our scope to cover all proposed methods extensively here. In this tutorial, we will summarize the popular categories of methods.

Classification by Input. We will first mention the three categories of methods based on the external knowledge provided to them. First, *unsupervised* methods take time series as input and are not provided by any other information. Then, *semi-supervised* methods take as input time series without any anomalies and are trained on normal data only. Finally, *supervised* methods take as input separately both normal and abnormal data. Thus, the model is trained to discriminate the anomalies from the normality.

Classification by Methodology. Then, we will describe the following categories of methods as depicted in Figure 1. First, *distance-based* approaches analyze subsequences by utilizing distances to a given model to detect anomalies. Second, *density-based* methods focus on detecting recurring or isolated behaviors by evaluating the density of the points or subsequences space. Third, *prediction-based* methods consist of two main categories: forecasting-based and reconstruction-based methods. The former, such as recurrent or convolutional neural network-based [22, 23], use the past values as input, predict the following, and use the forecasting error as an anomaly score. The latter, such as the autoencoder-based approach [32], are trained to reconstruct the time series and use the reconstruction error as an anomaly score.

3.3 Evaluating Anomaly Detection Methods

After describing various anomaly detectors, we will focus on how to evaluate them. The choice of benchmarks and accuracy measures may significantly bias the evaluation.

Evaluation Measures. We will start by describing the evaluation measures. Briefly, we will first discuss traditional measures, such as Precision, Recall, and F-score, that assess the methods by assuming each time-series point can be marked as an anomaly or not (e.g., by a threshold on an anomaly score). We will then discuss range-based variants [35] that aim to overcome shortcomings of traditional measures when evaluating time series containing subsequence anomalies. We will discuss Area Under the Curve (AUC) measures that, contrary to the previous measures, eliminate the need to define a threshold. We will finally discuss Volume Under the Surface (VUS) [4, 27] measures, which provide more robustness.

Benchmark Study. Then, we will discuss recent benchmarks proposed for anomaly detection in time series task [13, 15, 20, 29, 33]. Such benchmarks provide an extensive collection of time series from various domains and evaluate multiple methods belonging to the aforementioned categories. In this tutorial, we will discuss the results and conclusions of these recent benchmarks, as well as the criticisms that have been expressed about the characteristics and suitability of some datasets for this task [20, 37].

3.4 Automated Solutions for TSAD

Recent benchmarks and evaluation studies (described in the previous sections) have demonstrated that no overall best anomaly detection methods exist when applied on very heterogeneous time series (i.e., coming from very different domains). The primary question that arises is: How can we *automatically* identify the most

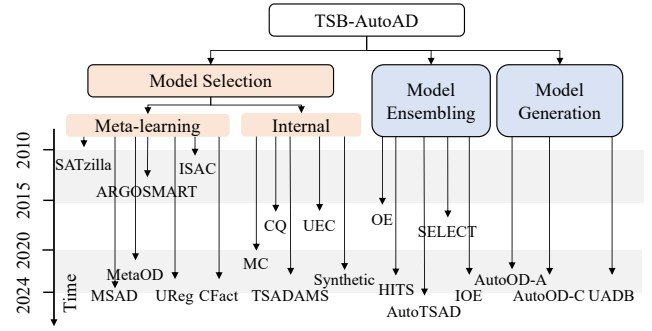


Figure 2: A taxonomy of automated solutions for TSAD.

accurate anomaly detector in the time series dataset? In this tutorial, we will present a taxonomy for the methods proposed in this field and provide recent experimental evaluations on them.

AutoAD Taxonomy. We will begin by outlining the pipeline of automated solutions and then delve into different categories of methods. The works in this field can be classified into two main categories as depicted in Figure 2. *Model Selection* refers to identifying the optimal model and its corresponding hyperparameters from a predefined set. Subsequently, the selected model is utilized for anomaly detection. In this category, internal evaluation methods evaluate the effectiveness of a model by using surrogate metrics for anomaly detection, independent of external data such as ground truth labels for anomalies [21] and meta-learning-based methods leverage the knowledge of the performance of various anomaly detectors on historical labeled datasets to enable the automatic model selection for new datasets [34]. On the other hand, *Model Ensembling* and *Model Generation* entail the construction of a completely new model based on the predefined set. This newly generated model then can operate independently as an anomaly detector. The former methods involve constructing models that combine ensembles from the predefined set [2]. Model generation methods generate pseudo-labels to transform the unsupervised anomaly detection problem into a supervised framework [8].

Experimental Results. We will then introduce a recent study [18, 19] on evaluating automated solutions and provide insights into the design choices. The study reveals that a majority of these solutions fail to outperform the simple random choice, indicating both the necessity for further investigation and the substantial potential within this domain. Furthermore, we will delve into the various scenarios in which they are applicable.

3.5 Challenges and Opportunities

In the last part of the tutorial, we will discuss the challenges and perspectives for next-generation time-series anomaly detection.

Foundation Model Empowered TSAD. Foundation models represent a new paradigm in large-scale self-supervised or unsupervised learning. They have demonstrated adaptability across modalities, extending beyond web data to scientific domains. Scaling the model, dataset size and data diversity have also been shown to result in remarkable transfer capabilities and excellent few-shot learning on novel datasets and tasks. Their capacity to process and analyze a large amount of data with temporal dependencies paves the way for more sophisticated and precise detection of anomalies. The tutorial will survey recent advances in applying foundation models

to TSAD, focusing on two complementary directions: (i) adapting models originally trained on other modalities (e.g., text, vision) for TSAD tasks, and (ii) developing time-series-native foundation models for anomaly detection. Subsequently, we will review the design choices and insights these methodologies offer.

Towards Automated Anomaly Detection. Given the absence of labeled data and one-fits-all anomaly detectors in unsupervised anomaly detection, it is challenging to determine the most accurate detector for various applications. Moreover, achieving optimal performance requires in-depth knowledge of the myriad of methods. This necessity drives data analysts into an exhaustive and time-consuming trial-and-error procedure process, making selecting suitable anomaly detectors cumbersome. In this section, we will further explore automated solutions (Section 3.4) and outline recent efforts. Following this, we will discuss prospective research directions, encompassing domain generalization and the development of incremental automated solutions. This field represents a promising area of research, and we hope our tutorial can serve as a catalyst to steer research efforts towards this direction.

3.6 Conclusions

We will conclude this tutorial by summarizing the main insights obtained from recent benchmarks of various anomaly detectors and automated solutions. Finally, we will summarize the new trends and challenges with the advent of automated solutions and foundation models and discuss the open research directions.

Acknowledgments

Supported by EU Horizon projects AI4Europe (101070000), Twin-ODIS (101160009), ARMADA (101168951), DataGEMS (101188416) and RECITALS (101168490), and by YITAIΘA & NextGenerationEU project HARSH (YIT3TA – 0560901) that is carried out within the framework of the National Recovery and Resilience Plan “Greece 2.0” with funding from the European Union – NextGenerationEU. This work was granted access to the HPC resources of IDRIS under the allocation 2025-A0191012641 made by GENCI.

References

- [1] Charu C Aggarwal. 2017. An introduction to outlier analysis. In *Outlier analysis*. 1–34.
- [2] Charu C Aggarwal and Saket Sathe. 2015. Theoretical foundations and algorithms for outlier ensembles. *Acm sigkdd explorations newsletter* 17, 1 (2015), 24–47.
- [3] Ane Blázquez-García, Angel Conde, Usue Mori, and Jose A Lozano. 2021. A Review on outlier/Anomaly Detection in Time Series Data. *ACM CSUR* 54, 3 (2021).
- [4] Paul Boniol, Ashwin K Krishna, Marine Bruel, Qinghua Liu, Mingyi Huang, Themis Palpanas, Ruey S Tsay, Aaron Elmore, Michael J Franklin, and John Paparrizos. 2025. VUS: effective and efficient accuracy measures for time-series anomaly detection. *The VLDB Journal* 34, 3 (2025), 32.
- [5] Paul Boniol, Qinghua Liu, Mingyi Huang, Themis Palpanas, and John Paparrizos. 2024. Dive into time-series anomaly detection: A decade review. *arXiv preprint arXiv:2412.20512* (2024).
- [6] Paul Boniol, John Paparrizos, and Themis Palpanas. 2023. New Trends in Time-Series Anomaly Detection. In *EDBT*.
- [7] Paul Boniol, John Paparrizos, and Themis Palpanas. 2024. An Interactive Dive into Time-Series Anomaly Detection. In *ICDE*.
- [8] Lei Cao, Yizhou Yan, Yu Wang, Samuel Madden, and Elke A Rundensteiner. 2023. AutoOD: Automatic Outlier Detection. *Proceedings of the ACM on Management of Data* 1, 1 (2023), 1–27.
- [9] Raghavendra Chalapathy, Nguyen Lu Dang Khoa, and Sanjay Chawla. 2020. Robust Deep Learning Methods for Anomaly Detection. In *SIGKDD (KDD '20)*.
- [10] Jens E d'Hondt, Haojun Li, Fan Yang, Odysseas Papapetrou, and John Paparrizos. 2025. A Structured Study of Multivariate Time-Series Distance Measures. *Proceedings of the ACM on Management of Data* 3, 3 (2025), 1–29.
- [11] Manish Gupta, Jing Gao, Charu C Aggarwal, and Jiawei Han. 2013. Outlier detection for temporal data: A survey. *TKDE* 26, 9 (2013), 2250–2267.
- [12] Medina Hadjem, Farid Nait-Abdesselam, and Ashfaq A. Khokhar. 2016. ST-segment and T-wave anomalies prediction in an ECG data using RUSBoost. In *Healthcom*.
- [13] Vincent Jacob, Fei Song, Arnaud Stiegler, Bijan Rad, Yanlei Diao, and Nesime Tatbul. 2021. Exathlon: A Benchmark for Explainable Anomaly Detection over Time Series. *Proc. VLDB Endow.* 14, 11 (oct 2021), 2613–2626.
- [14] Hao Jiang, Chunwei Liu, John Paparrizos, Andrew A Chien, Jihong Ma, and Aaron J Elmore. 2021. Good to the Last Bit: Data-Driven Encoding with CodecDB. In *SIGMOD*. 843–856.
- [15] E. Keogh, T. Dutta Roy, U. Naik, and A Agrawal. 2021. Multi-dataset Time-Series Anomaly Detection Competition 2021. <https://compete.hexagon-ml.com/practice/competition/39/>.
- [16] Qinghua Liu, Paul Boniol, Themis Palpanas, and John Paparrizos. 2024. Time-Series Anomaly Detection: Overview and New Trends. *PVLDB* 17, 12 (2024), 4229–4232.
- [17] Qinghua Liu, Sam Heshmati, Zheda Mai, Zubin Abraham, John Paparrizos, and Liu Ren. 2025. MLLM4TS: Leveraging Vision and Multimodal Language Models for General Time-Series Analysis. *arXiv preprint arXiv:2510.07513* (2025).
- [18] Qinghua Liu, Seunghak Lee, and John Paparrizos. 2025. EasyAD: A Demonstration of Automated Solutions for Time-Series Anomaly Detection. *Proceedings of the VLDB Endowment* 18, 12 (2025), 5431–5434.
- [19] Qinghua Liu, Seunghak Lee, and John Paparrizos. 2025. Tsb-autoad: Towards automated solutions for time-series anomaly detection. *PVLDB* 18, 11 (2025), 4364–4379.
- [20] Qinghua Liu and John Paparrizos. 2024. The Elephant in the Room: Towards A Reliable Time-Series Anomaly Detection Benchmark. In *NeurIPS 2024*.
- [21] Martin Q Ma, Yue Zhao, Xiaorong Zhang, and Leman Akoglu. 2023. The need for unsupervised outlier model selection: A review and evaluation of internal evaluation strategies. *ACM SIGKDD Explorations Newsletter* 25, 1 (2023).
- [22] Pankaj Malhotra, Lovekesh Vig, Gautam Shroff, and Puneet Agarwal. 2015. Long Short Term Memory Networks for Anomaly Detection in Time Series. (2015).
- [23] M. Munir, S. A. Siddiqui, A. Dengel, and S. Ahmed. 2019. DeepAnT: A Deep Learning Approach for Unsupervised Anomaly Detection in Time Series. *IEEE Access* 7 (2019), 1991–2005.
- [24] Themis Palpanas and Volker Beckmann. 2019. Report on the First and Second Interdisciplinary Time Series Analysis Workshop (TISA). *SIGMOD Rec.* 48, 3 (2019), 36–40.
- [25] Ioannis Paparrizos. 2018. *Fast, Scalable, and Accurate Algorithms for Time-Series Analysis*. Ph. D. Dissertation. Columbia University.
- [26] John Paparrizos, Paul Boniol, Qinghua Liu, and Themis Palpanas. 2025. Advances in time-series anomaly detection: Algorithms, benchmarks, and evaluation measures. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*. 6151–6161.
- [27] John Paparrizos, Paul Boniol, Themis Palpanas, Ruey S. Tsay, Aaron Elmore, and Michael J. Franklin. 2022. Volume under the Surface: A New Accuracy Evaluation Measure for Time-Series Anomaly Detection. *Proc. VLDB Endow.* 15, 11 (2022).
- [28] John Paparrizos, Ikradya Edian, Chunwei Liu, Aaron J Elmore, and Michael J Franklin. 2022. Fast adaptive similarity search through variance-aware quantization. In *ICDE*. IEEE, 2969–2983.
- [29] John Paparrizos, Yuhao Kang, Paul Boniol, Ruey S. Tsay, Themis Palpanas, and Michael J. Franklin. 2022. TSB-UAD: An End-to-End Benchmark Suite for Univariate Time-Series Anomaly Detection. *PVLDB* 15, 8 (2022).
- [30] John Paparrizos, Chunwei Liu, Bruno Barbarioli, Johnny Hwang, Ikradya Edian, Aaron J Elmore, Michael J Franklin, and Sanjay Krishnan. 2021. VergeDB: A Database for IoT Analytics on Edge Devices.. In *CIDR*.
- [31] John Paparrizos, Fan Yang, and Haojun Li. 2024. Bridging the gap: A decade review of time-series clustering methods. *arXiv preprint arXiv:2412.20582* (2024).
- [32] Mayu Sakurada and Takehisa Yairi. 2014. Anomaly Detection Using Autoencoders with Nonlinear Dimensionality Reduction. In *MLSDA*.
- [33] Sebastian Schmidl, Phillip Wenig, and Thorsten Papenbrock. 2022. Anomaly Detection in Time Series: A Comprehensive Evaluation. *PVLDB* 15, 9 (2022), 1779–1797.
- [34] Emmanouil Sylligardos, Paul Boniol, John Paparrizos, Panos Trahanias, and Themis Palpanas. 2023. Choose wisely: An extensive evaluation of model selection for anomaly detection in time series. *Proceedings of the VLDB Endowment* 16, 11 (2023), 3418–3432.
- [35] Nesime Tatbul, Tae Jun Lee, Stan Zdonik, Mejbah Alam, and Justin Gottschlich. 2018. Precision and Recall for Time Series. In *NeurIPS*, Vol. 31.
- [36] Qingsong Wen, Linxiao Yang, Tian Zhou, and Liang Sun. 2022. Robust Time Series Analysis and Applications: An Industrial Perspective. In *SIGKDD*.
- [37] Renjie Wu and Eamonn J. Keogh. 2023. Current Time Series Anomaly Detection Benchmarks are Flawed and are Creating the Illusion of Progress. *IEEE Trans. Knowl. Data Eng.* 35, 3 (2023), 2421–2429. doi:10.1109/TKDE.2021.3112126
- [38] Guang Yang, Ninad Kulkarni, Paavani Dua, Dipika Khullar, and Alex Anto Chirayath. 2022. Anomaly Detection for Spatiotemporal Data in Action. In *KDD*.