



An Interactive Dive into Time Series Anomaly Detection

Paul Boniol

Inria, ENS, PSL University
Paul.boniol@inria.fr

John Paparrizos

Ohio State University
paparrizos.1@osu.edu

Themis Palpanas

Université Paris Cité; IUF
themis@mi.parisdescartes.fr



THE OHIO STATE UNIVERSITY

Introduction: *Time series are Everywhere*

Energy Production



Edf.fr: tinyurl.com/yc7x5xje

Astrophysics



Virgo: <https://www.virgo-gw.eu/>

Medicine



tinyurl.com/39dx2us4

Volcanology

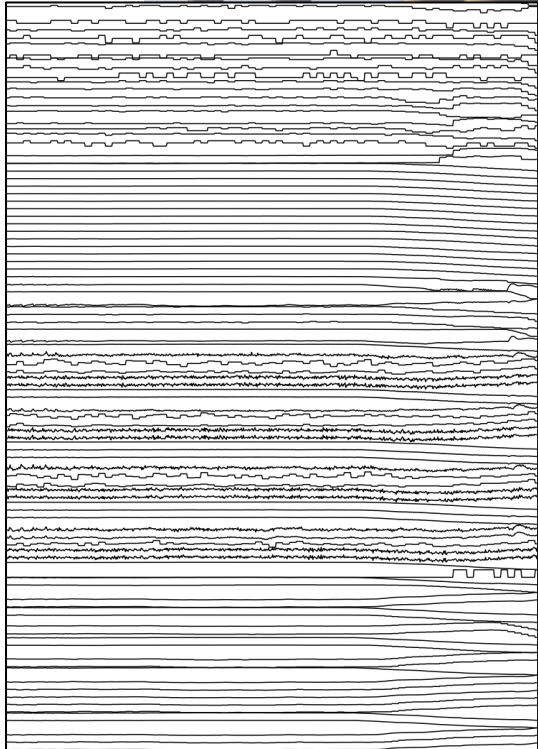


tinyurl.com/ybcttmfz

Introduction: *Time series are Everywhere*

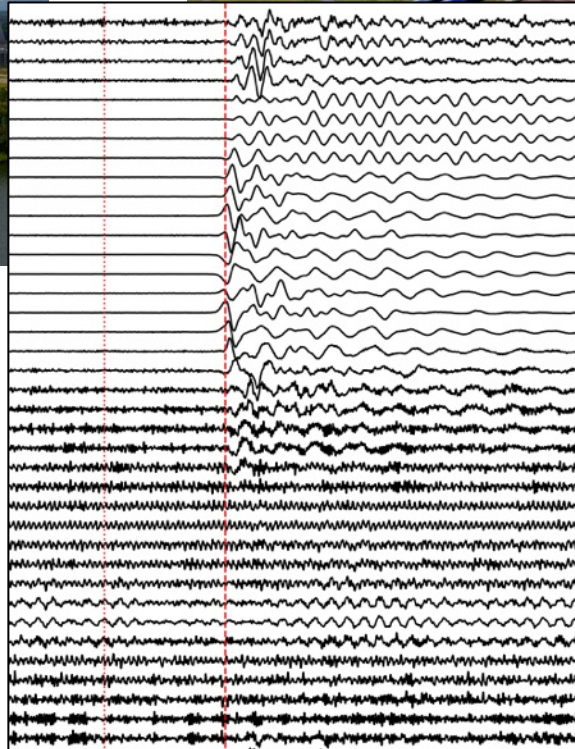
Energy Production

Secondary circuit sensor measurements



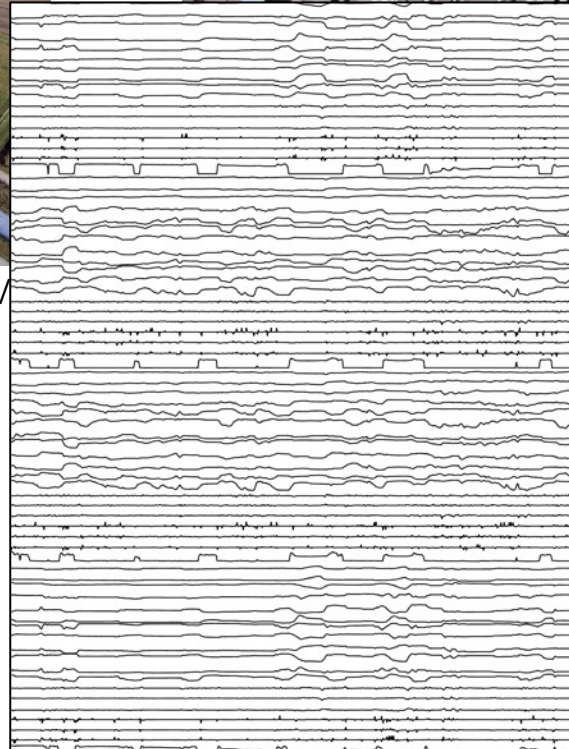
Astrophysics

Fiber-acoustic sensors in the VIRGO north building



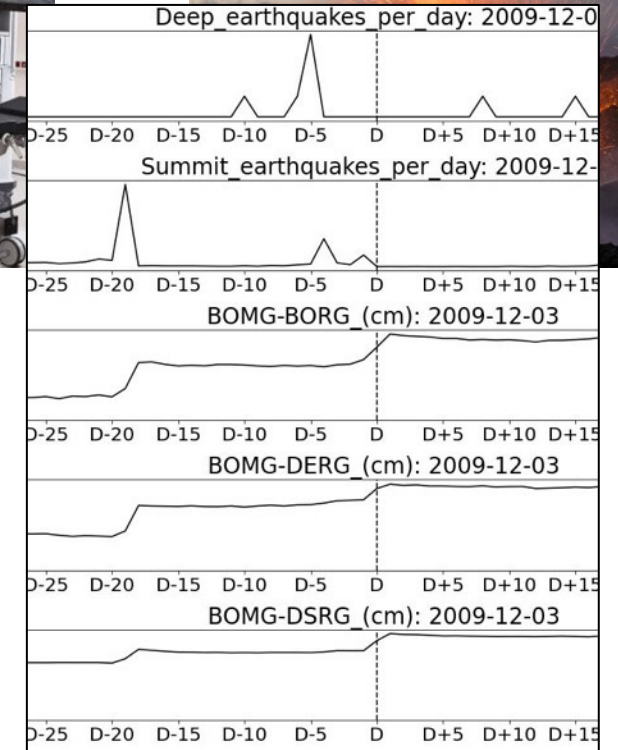
Medicine

Sensor measurements of the Da Vinci surgery robot



Volcanology

Sensor measurements on le Piton de la Fournaise

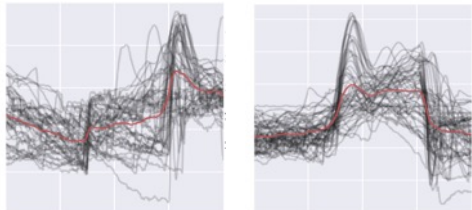


Introduction: *with Important Challenges*

Energy Production

Secondary circuit sensor measurements

Identification of precursors of feed-water pumps vibrations



Astrophysics

Fiber-acoustic sensors in the VIRGO north building

Noise detection in VIRGO interferometer north building



Medicine

Sensor measurements of the Da Vinci surgery robot

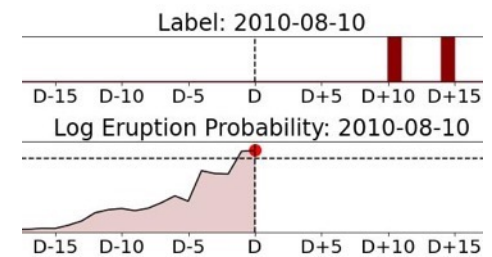
Unusual surgeons gestures detection



Volcanology

Sensor measurements on le Piton de la Fournaise

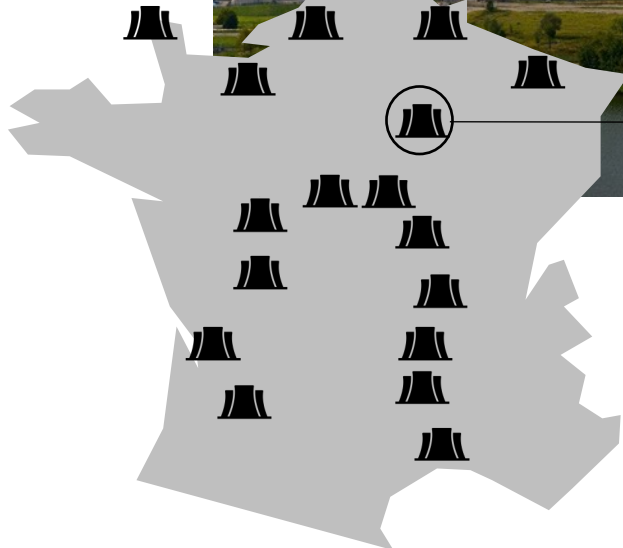
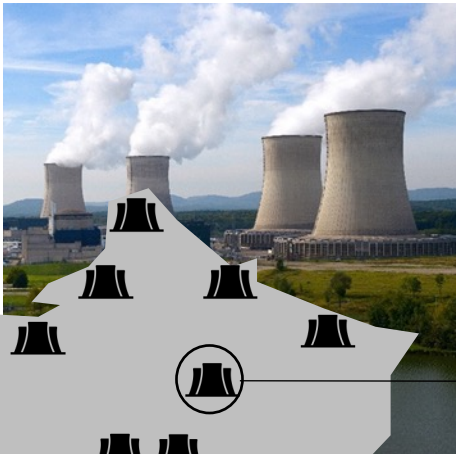
Detection of abnormal events on the volcano for predicting eruptions



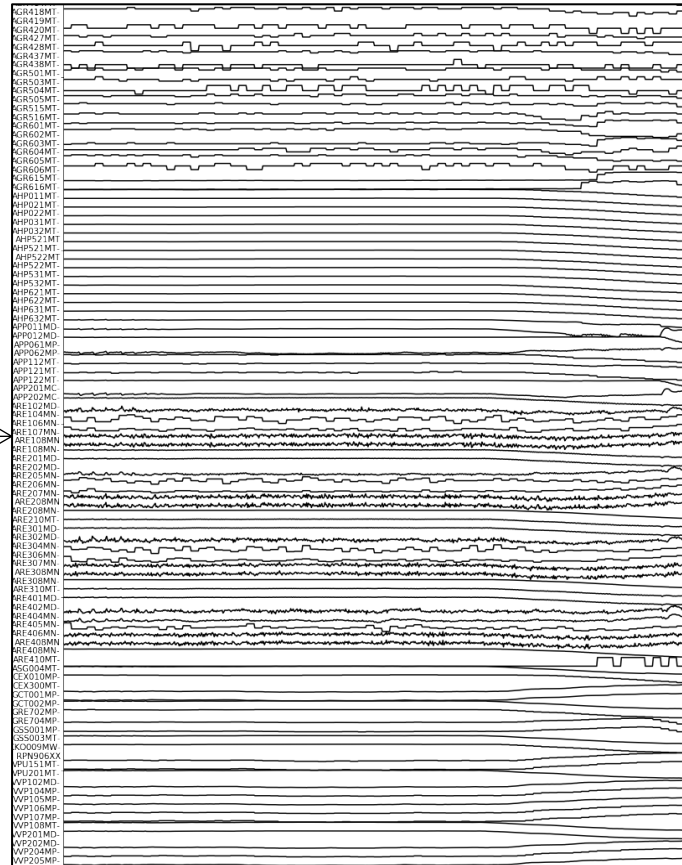
Introduction: *with Important Challenges*

Large-scale time series database

Energy Production



Edf.fr: tinyurl.com/yc7x5xje



Example of Nuclear production

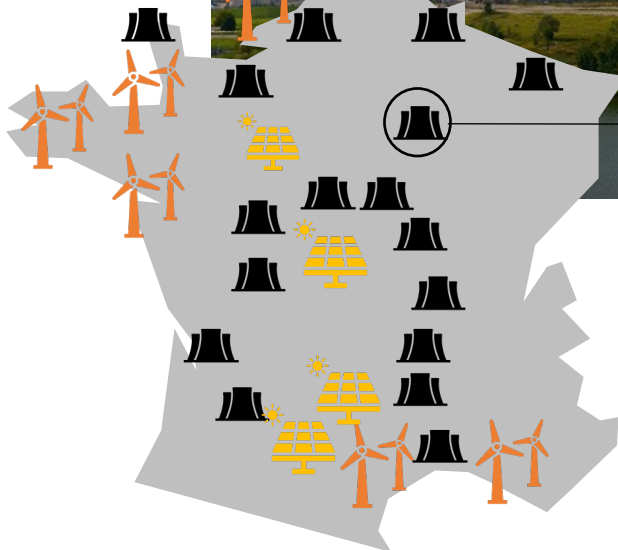
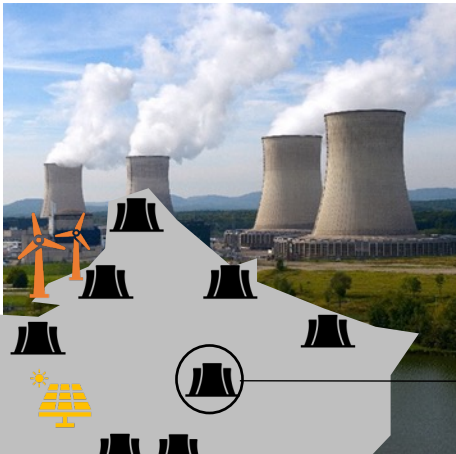
- 58 nuclear power plants across France
- 2000+ sensors per power plant
- 30 years of data collections

A total of 500 TeraBytes

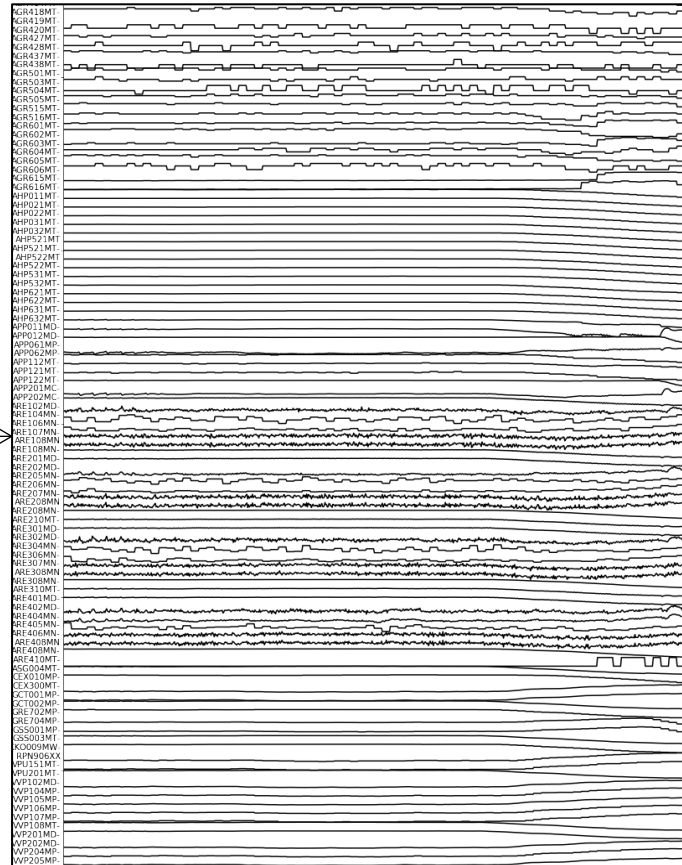
Introduction: *with Important Challenges*

Large-scale time series database

Energy Production



Edf.fr: tinyurl.com/yc7x5xje



Example of Nuclear production

- 58 nuclear power plants across France
- 2000+ sensors per power plant
- 30 years of data collections

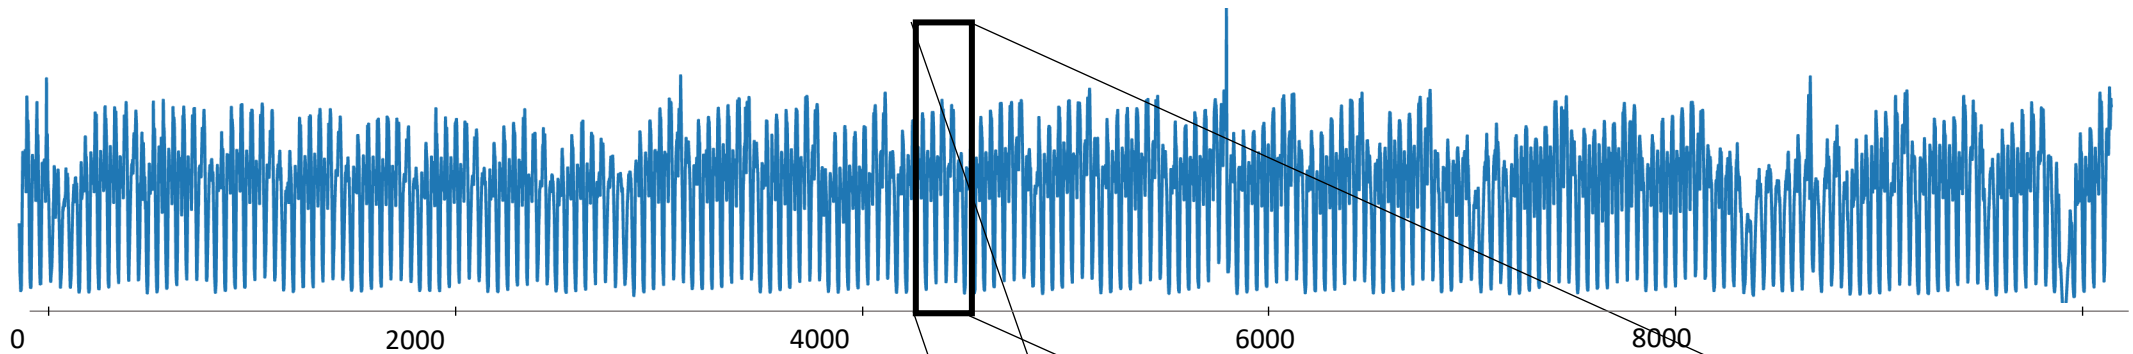
A total of 500 TeraBytes

Other source of production

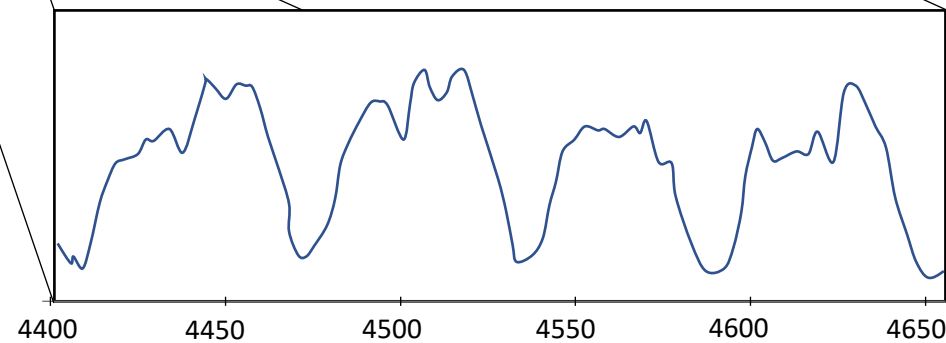
- New sensors with higher acquisition rate

Introduction: *Anomaly Detection in Time Series*

- Time series T (example : number of taxi passengers in New York City)

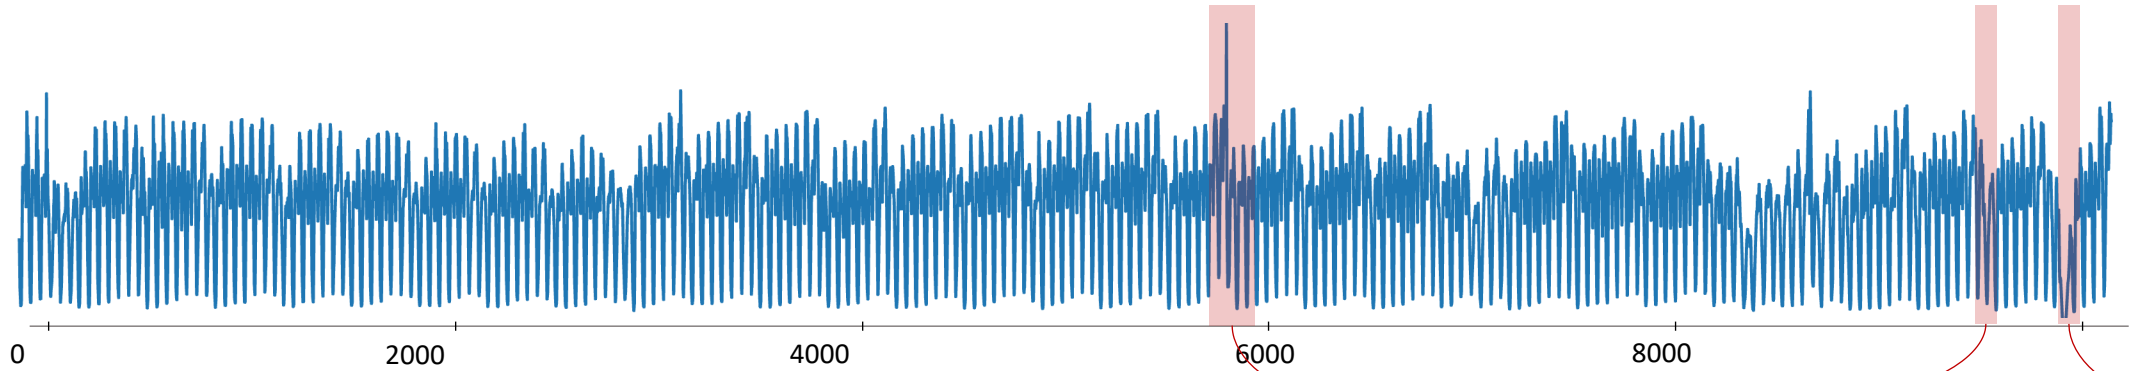


- Subsequence $T_{i,\ell}$
with $i = 4400, \ell = 250$



Introduction: *Anomaly Detection in Time Series*

- Time series T (example : number of taxi passengers in New York City)



- Anomaly: *rare* point or sequence (of a given length) potentially *non-desired*

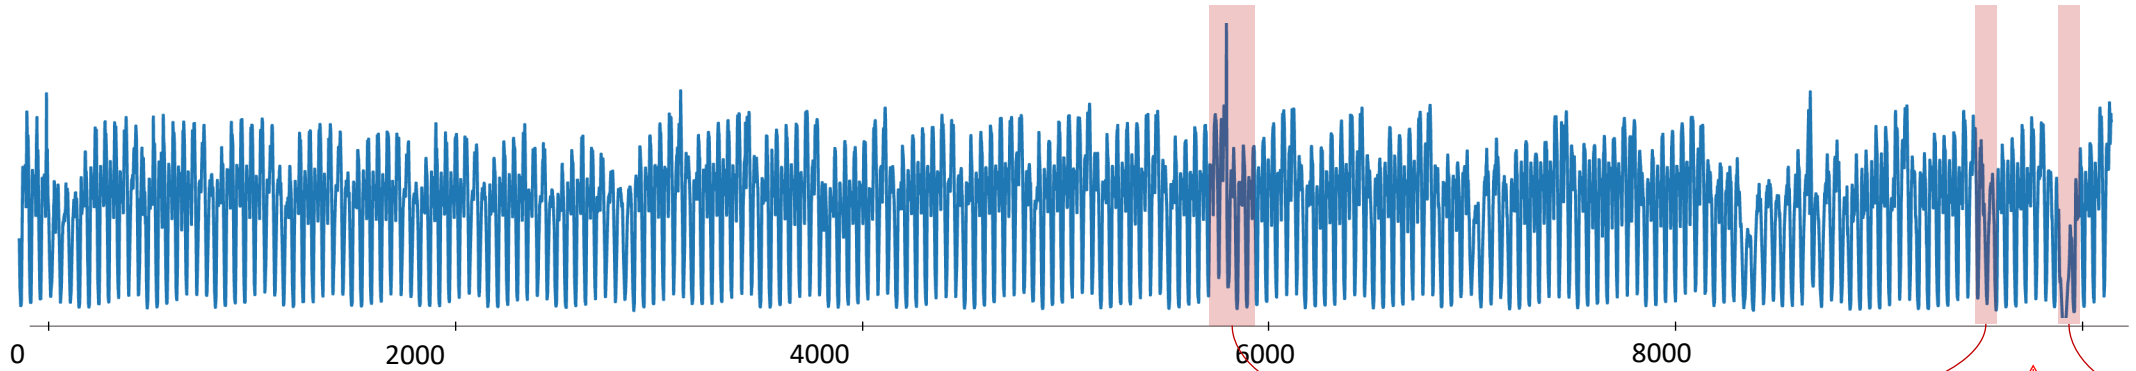
Daylight
Saving Time
(DST)

Flooding

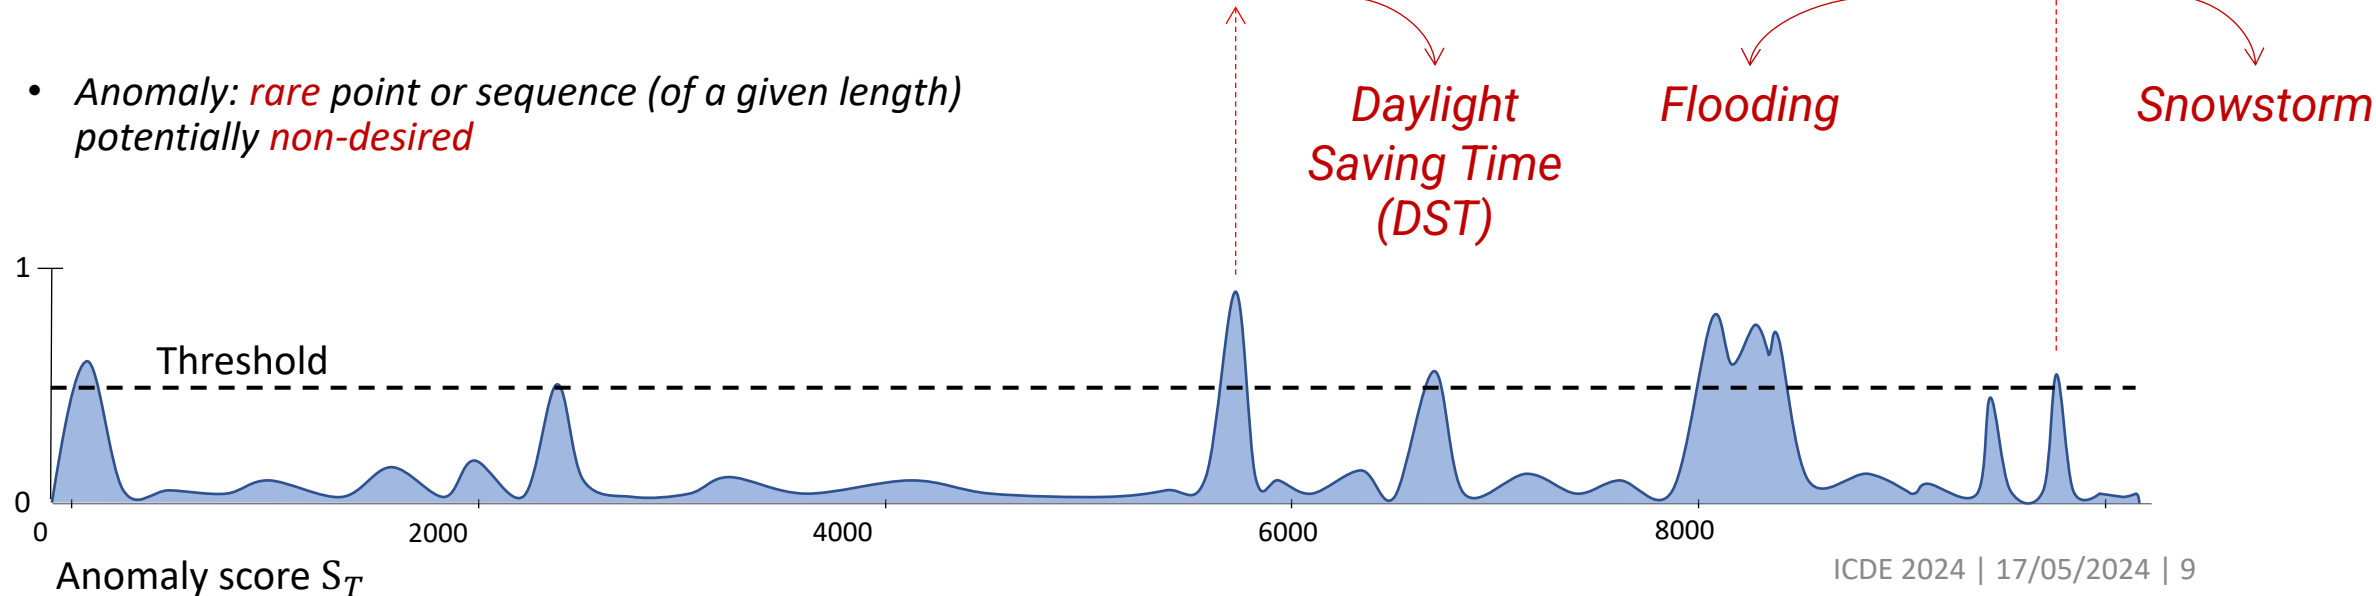
Snowstorm

Introduction: *Anomaly Detection in Time Series*

- Time series T (example : number of taxi passengers in New York City)

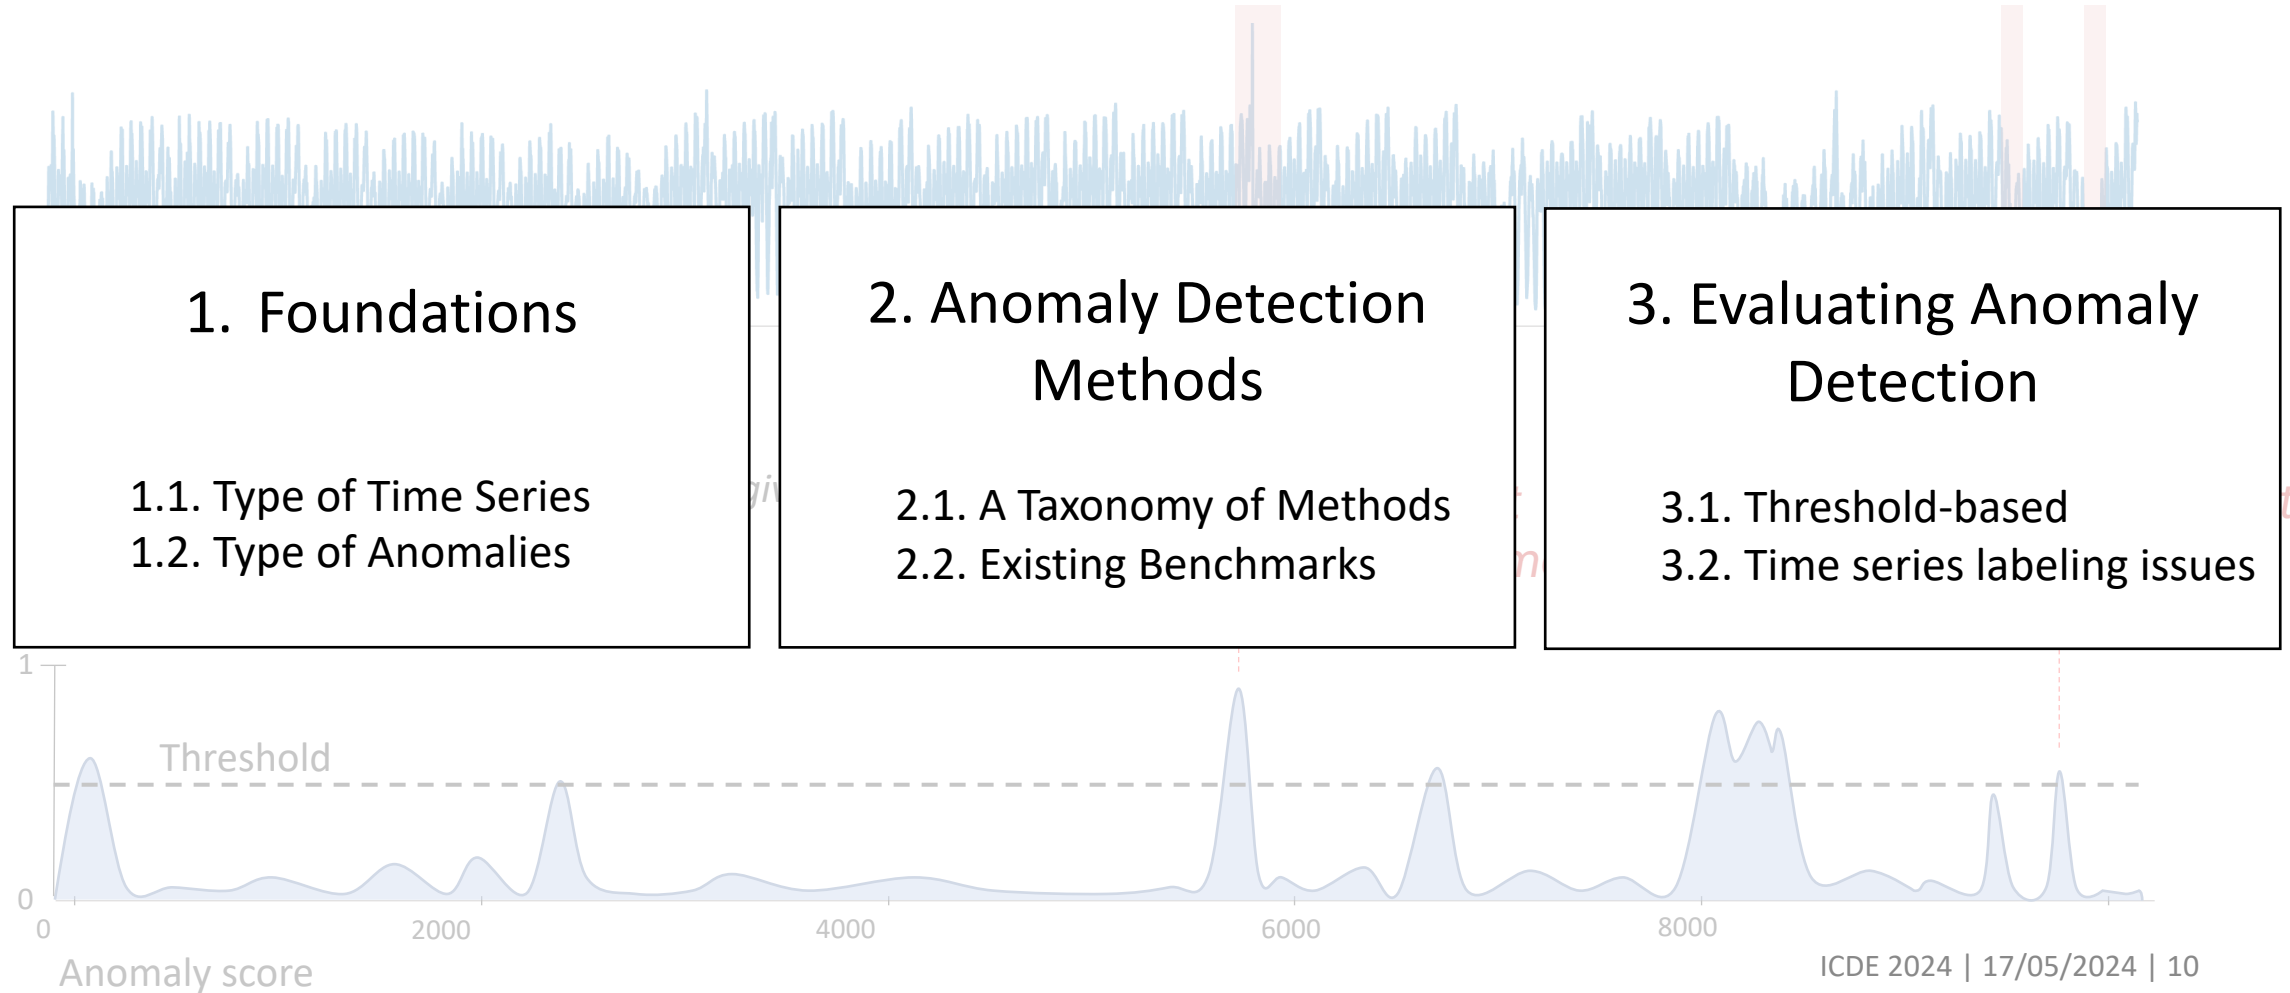


- Anomaly: *rare* point or sequence (of a given length) potentially *non-desired*



Introduction: *Outline*

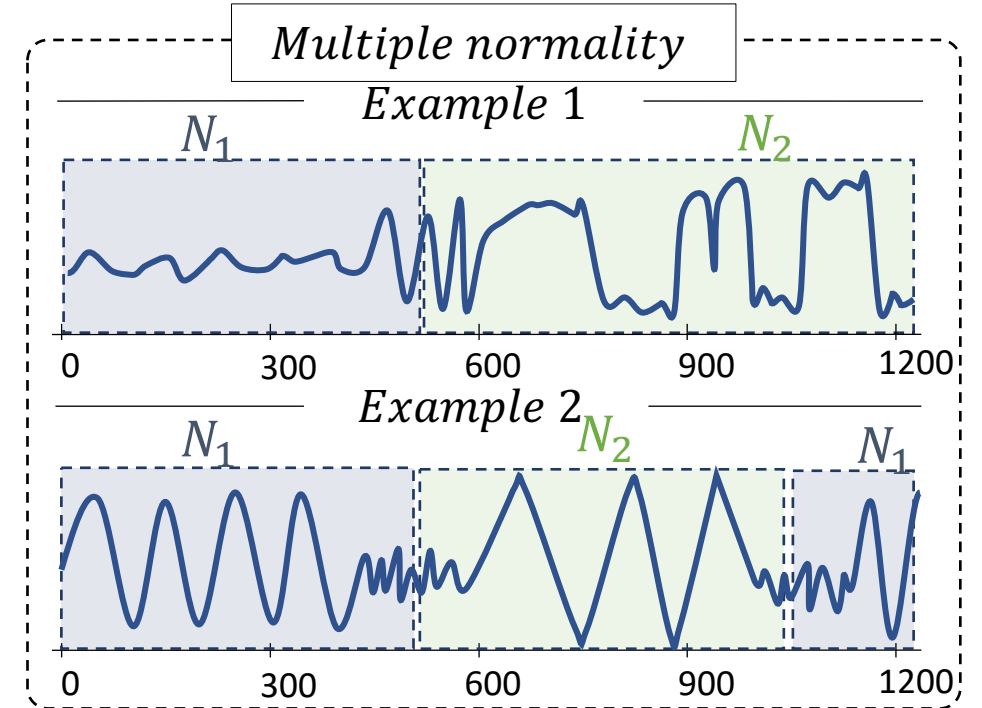
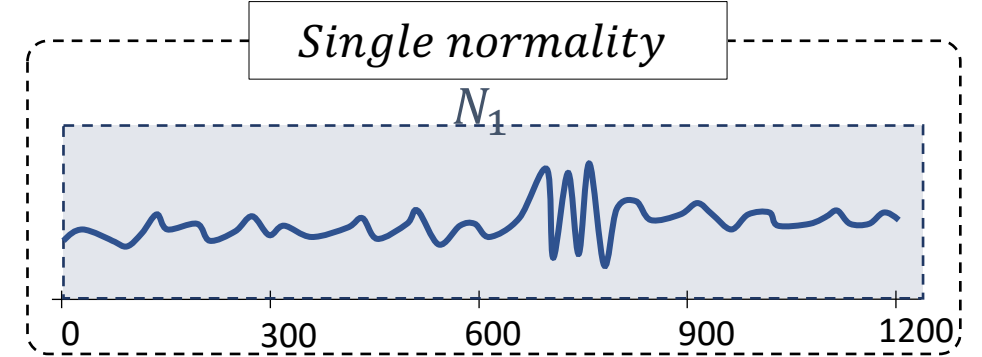
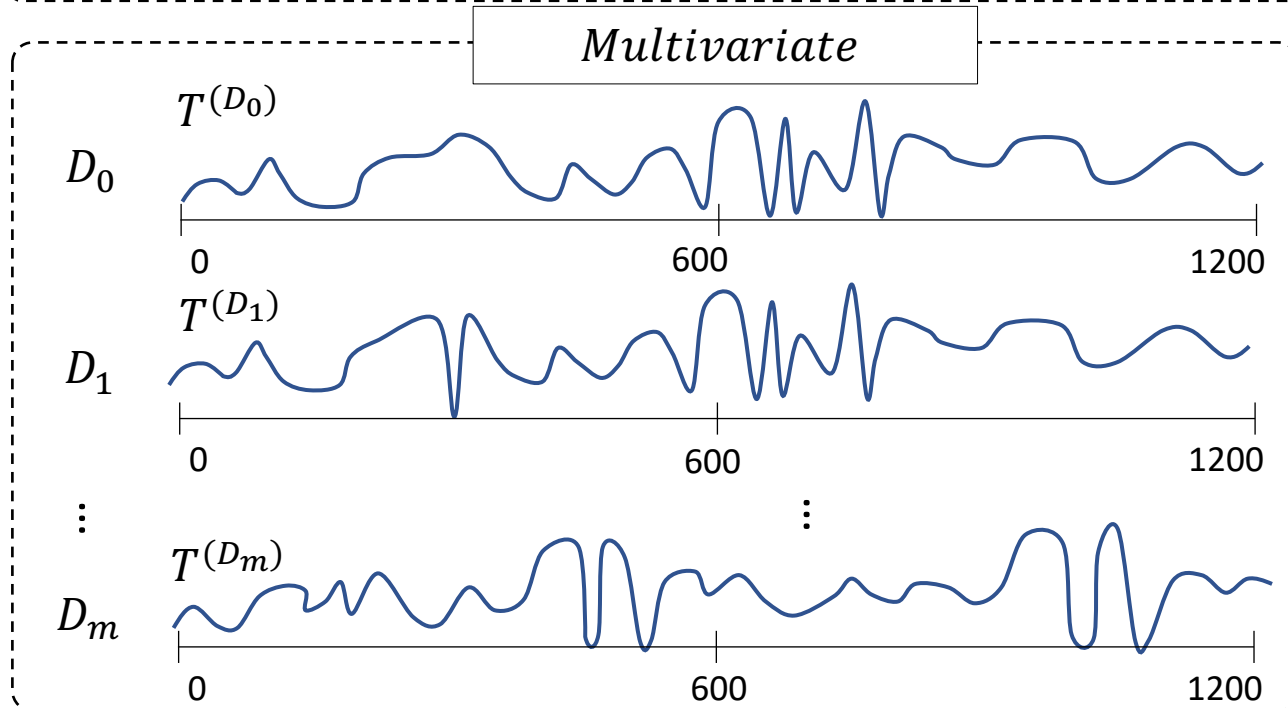
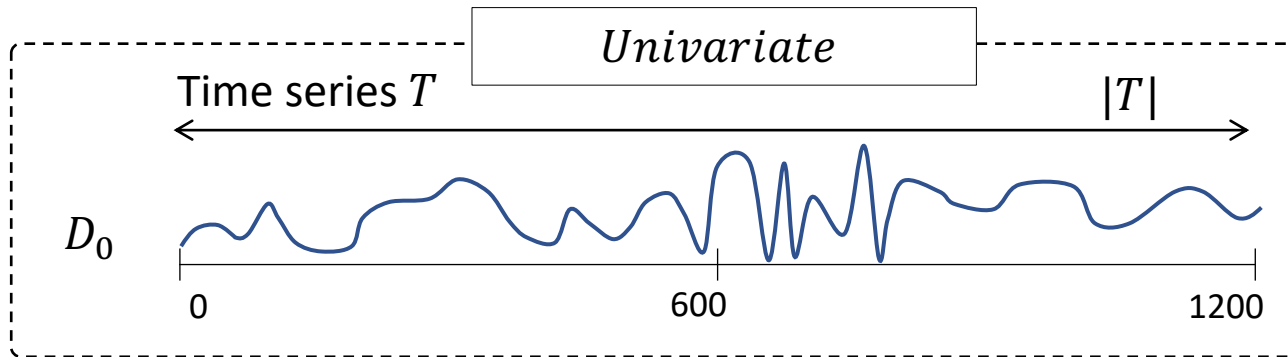
- Time series (example : number of taxi passengers in New York City)



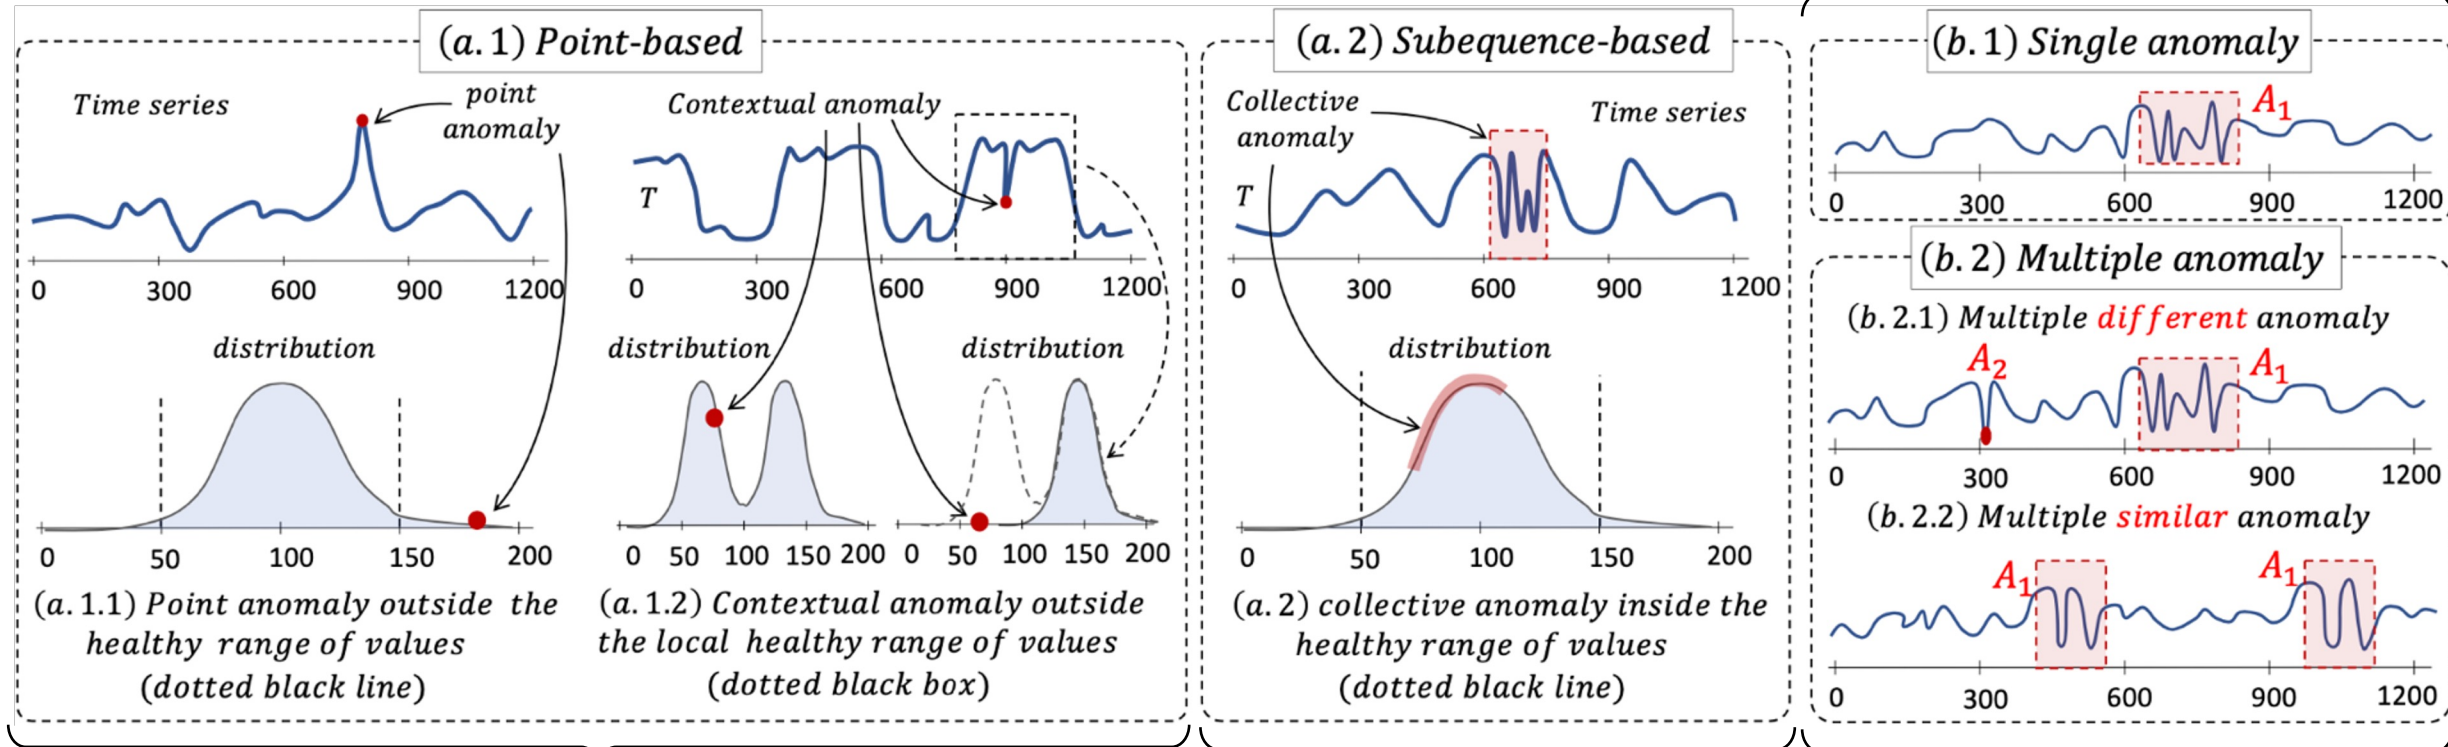
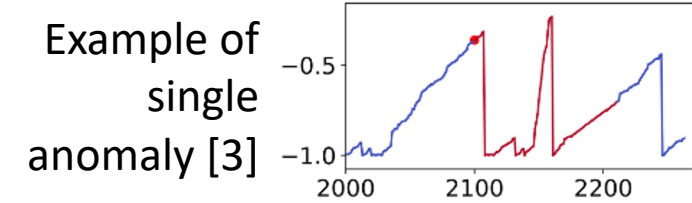


Foundations

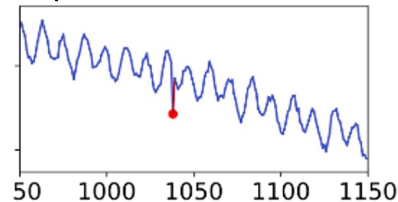
Foundations: *Type of time series*



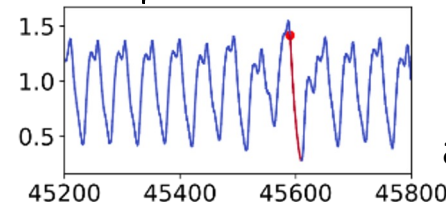
Foundations: *Type of anomalies*



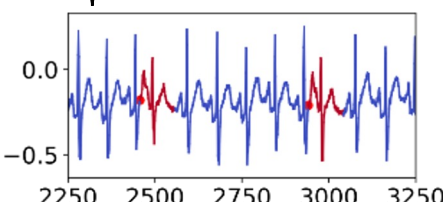
Example of point-based anomaly [1]



Example of subsequence-based anomaly [2]

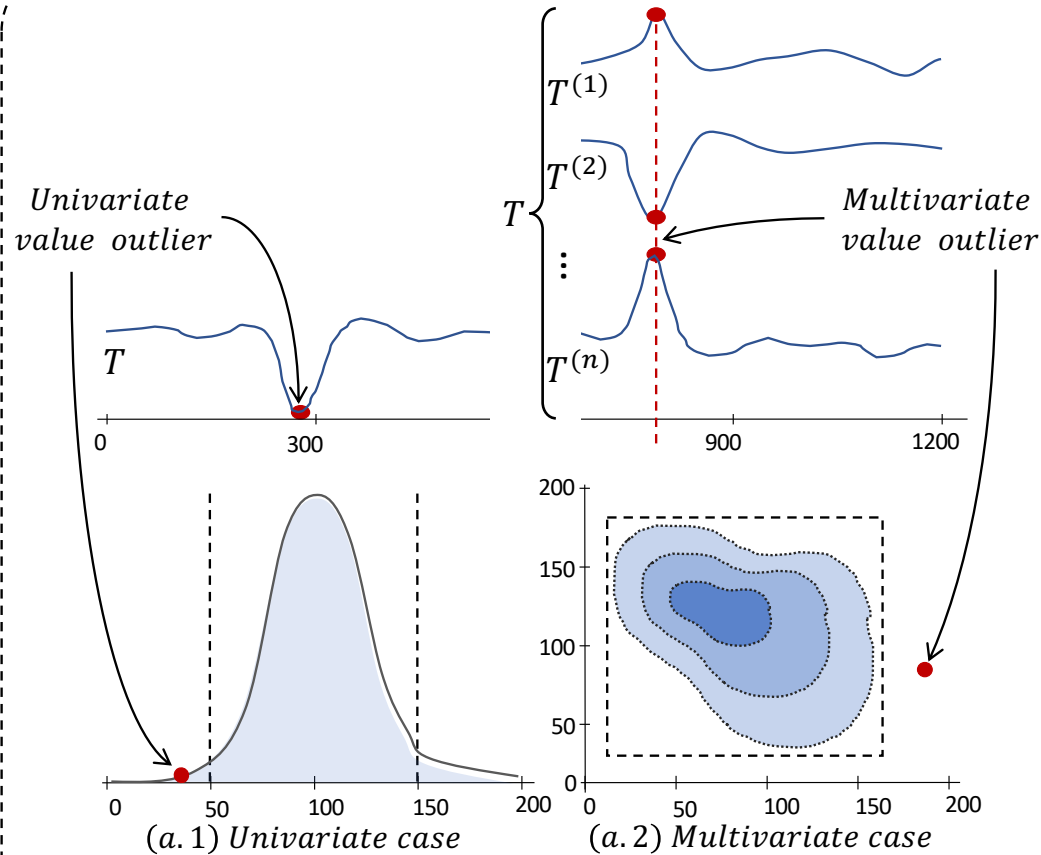


Example of multiple anomaly [4]



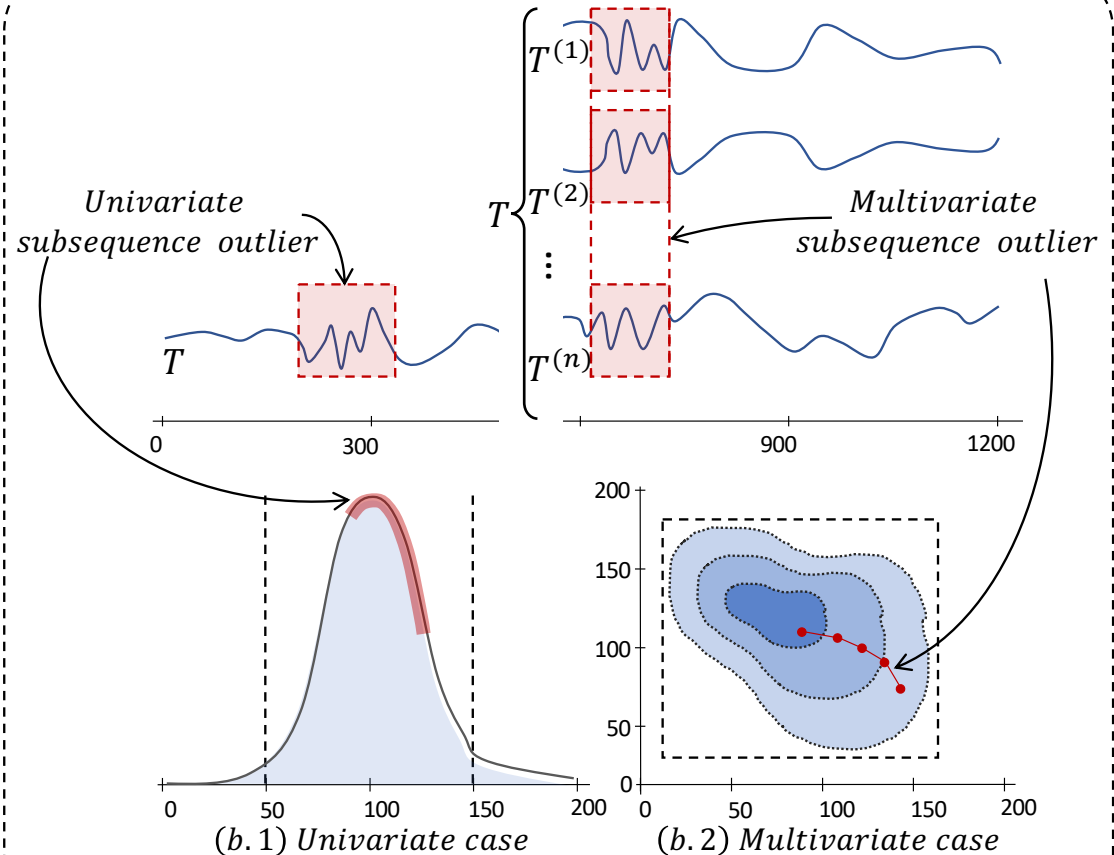
Foundations: *Type of anomalies*

Univariate and Multivariate point anomalies



(a) Point outlier outside the healthy range of values (dotted black line)

Univariate and Multivariate sequence anomalies

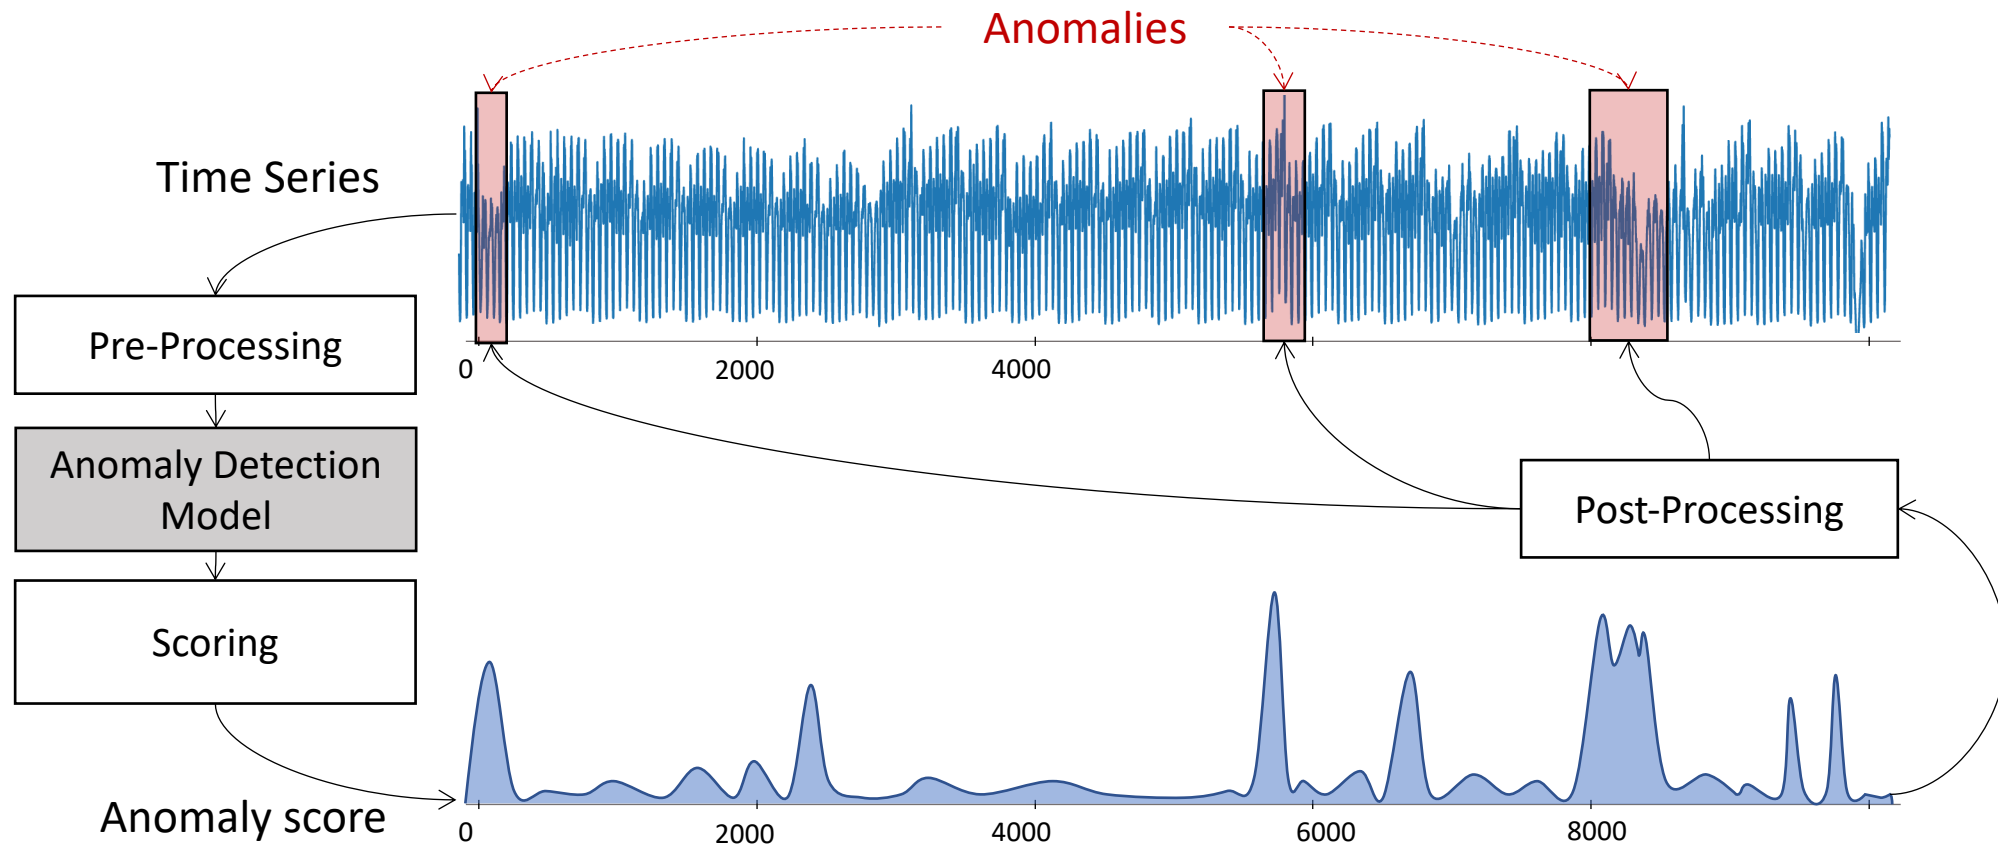


(b) Subsequence outlier inside the healthy range of values (dotted black line)



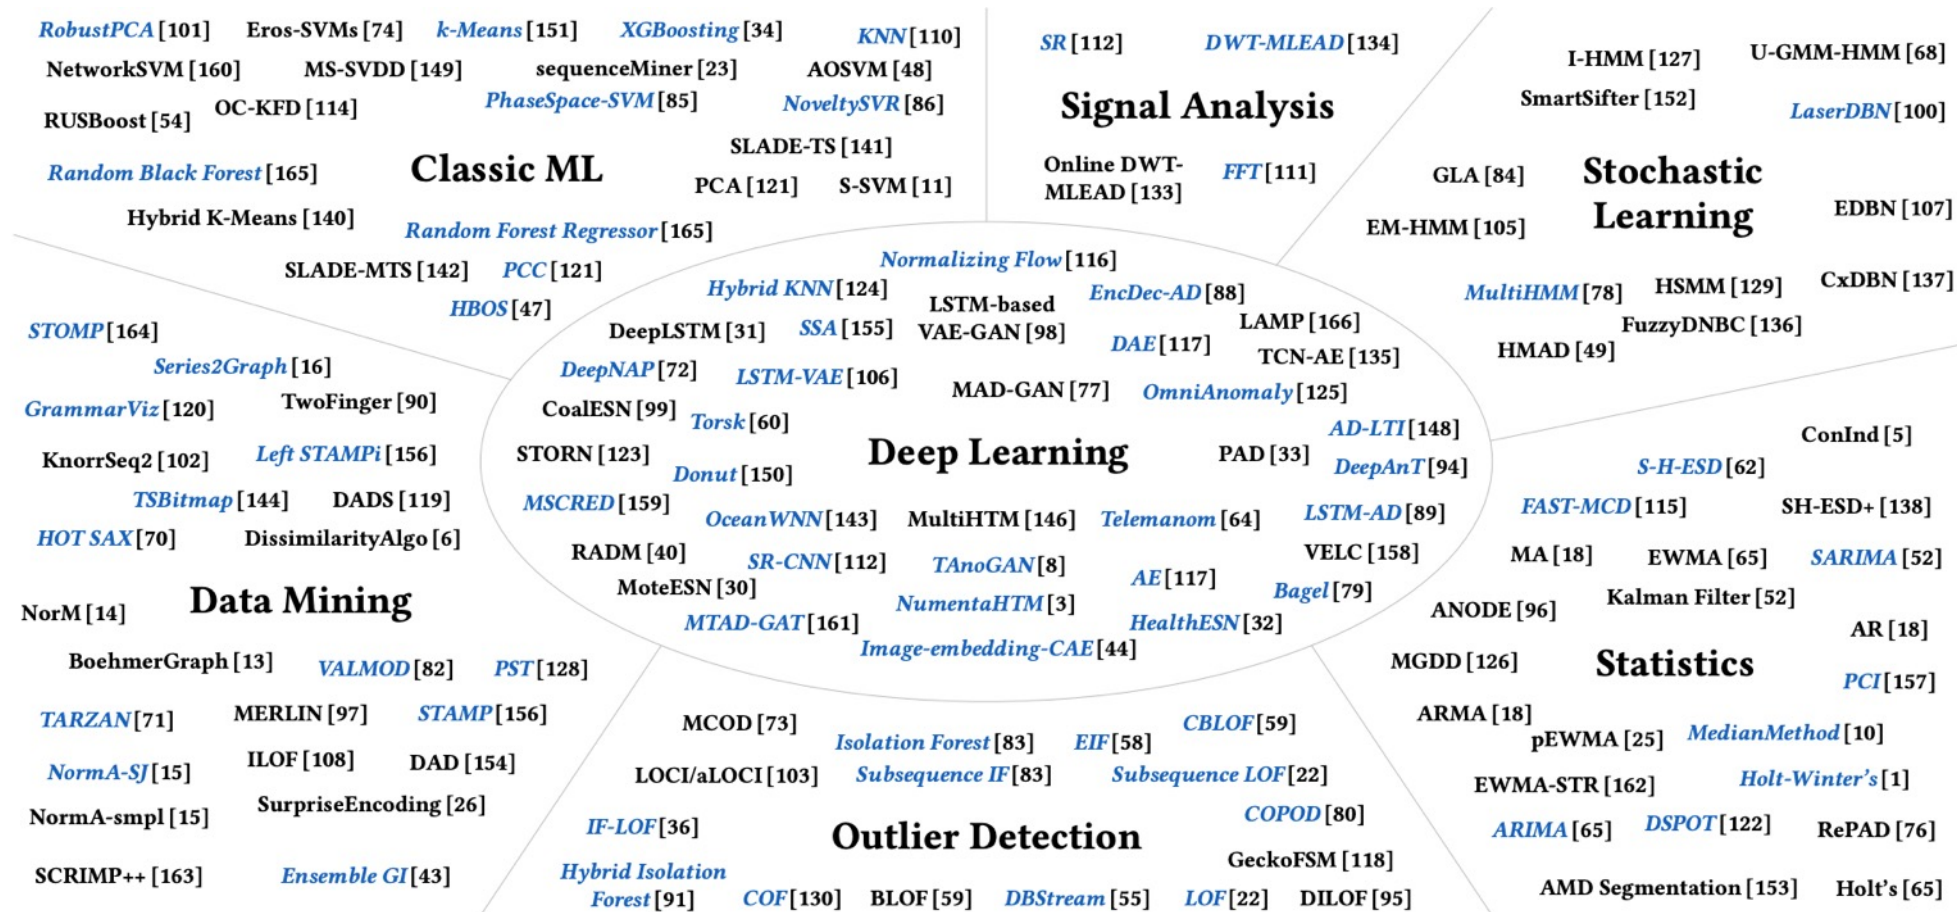
Anomaly Detection Methods

Anomaly Detection methods: *A taxonomy*



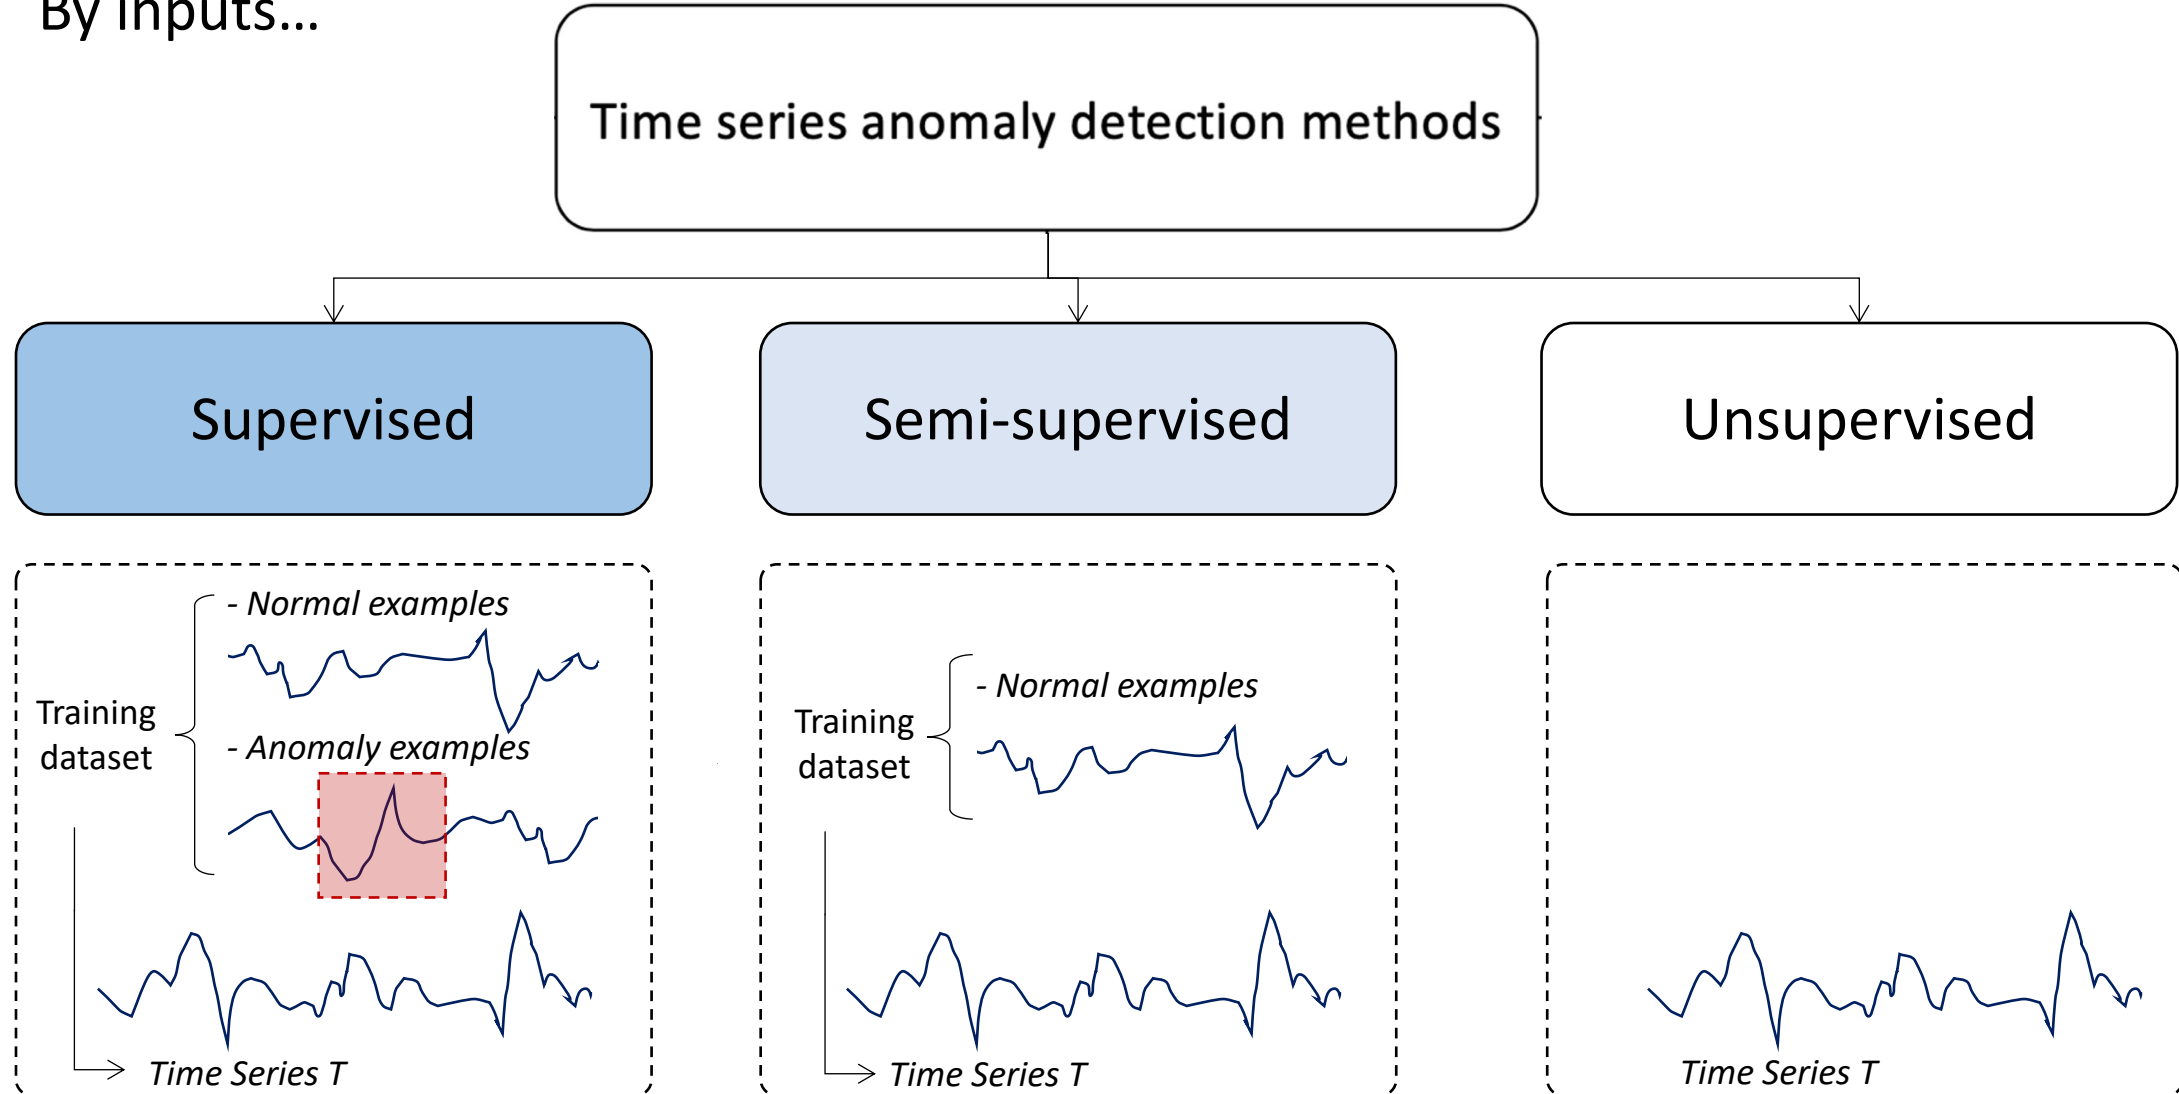
Anomaly Detection methods: *A taxonomy*

By domains [5] ...



Anomaly Detection methods: *A taxonomy*

By inputs...



Anomaly Detection

By inputs...

Time

Supervised

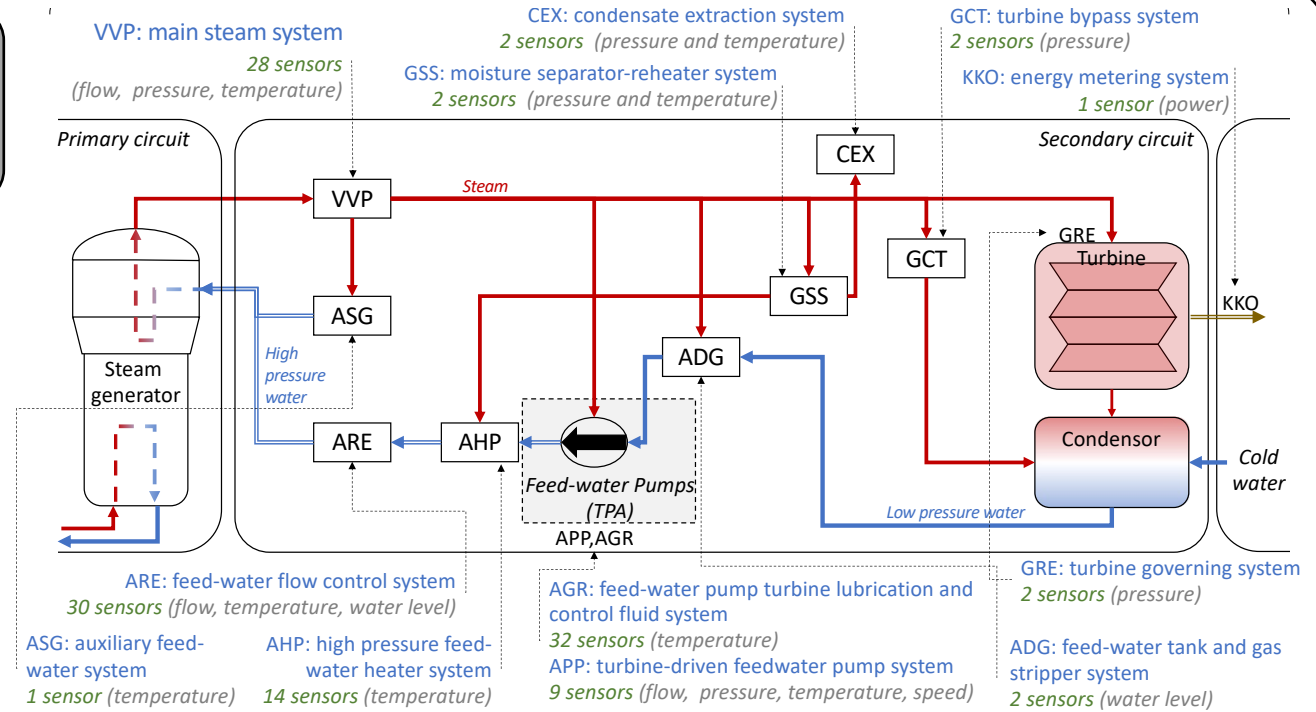
Training dataset

- Normal examples

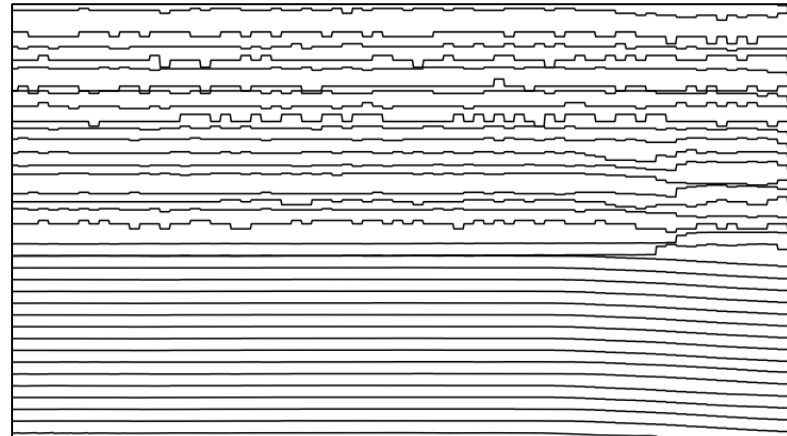
- Anomaly examples

Time Series T

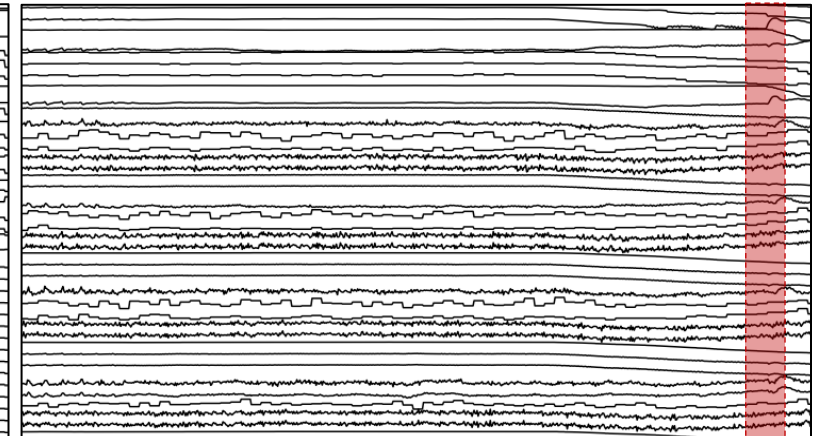
Supervised anomaly detection (e.g., classification)



Class 1: Time series without any vibrations



Class 2: Time series with a vibrations



Vibration

Anomaly Detection

By inputs...

Time

Supervised

Supervised anomaly detection (e.g., classification)

Explanation of the detection

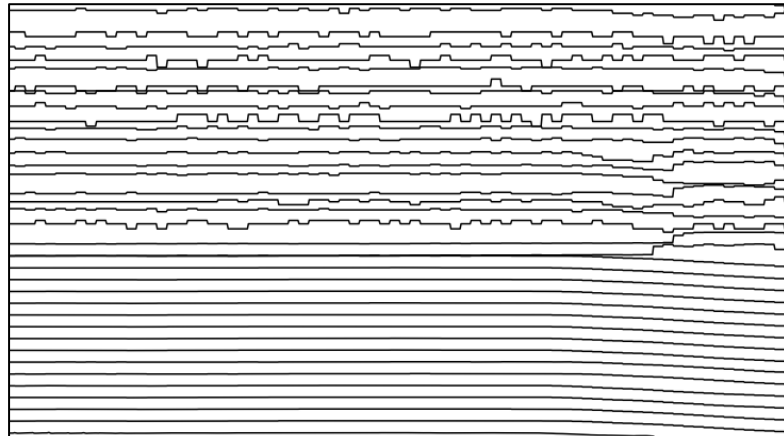
Training dataset

- Normal examples

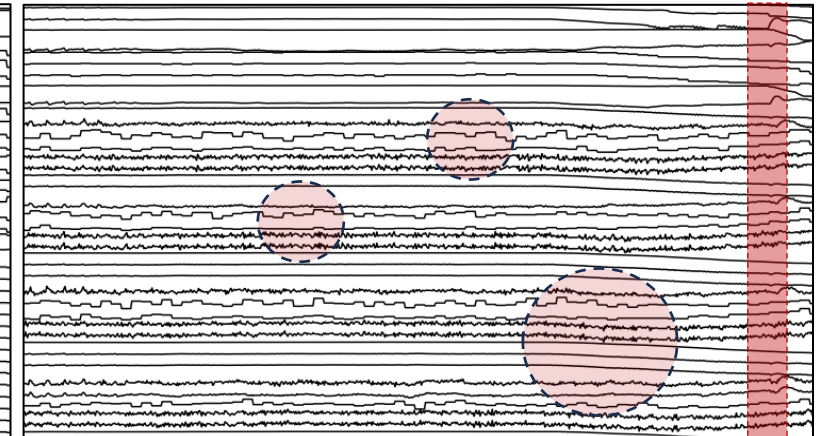
- Anomaly examples

Time Series T

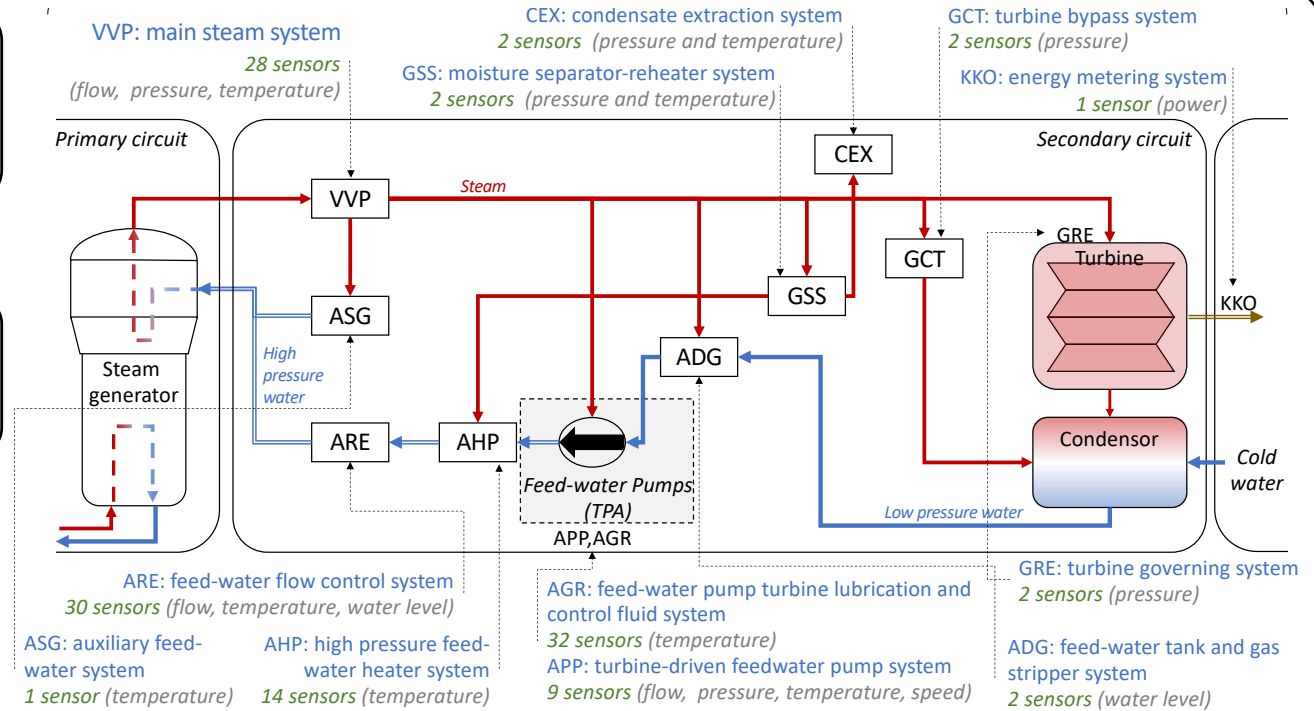
Class 1: Time series without any vibrations



Class 2: Time series with a vibrations



Vibration



Anomaly Detection

By inputs...

Time

Supervised

Training dataset

- Normal examples

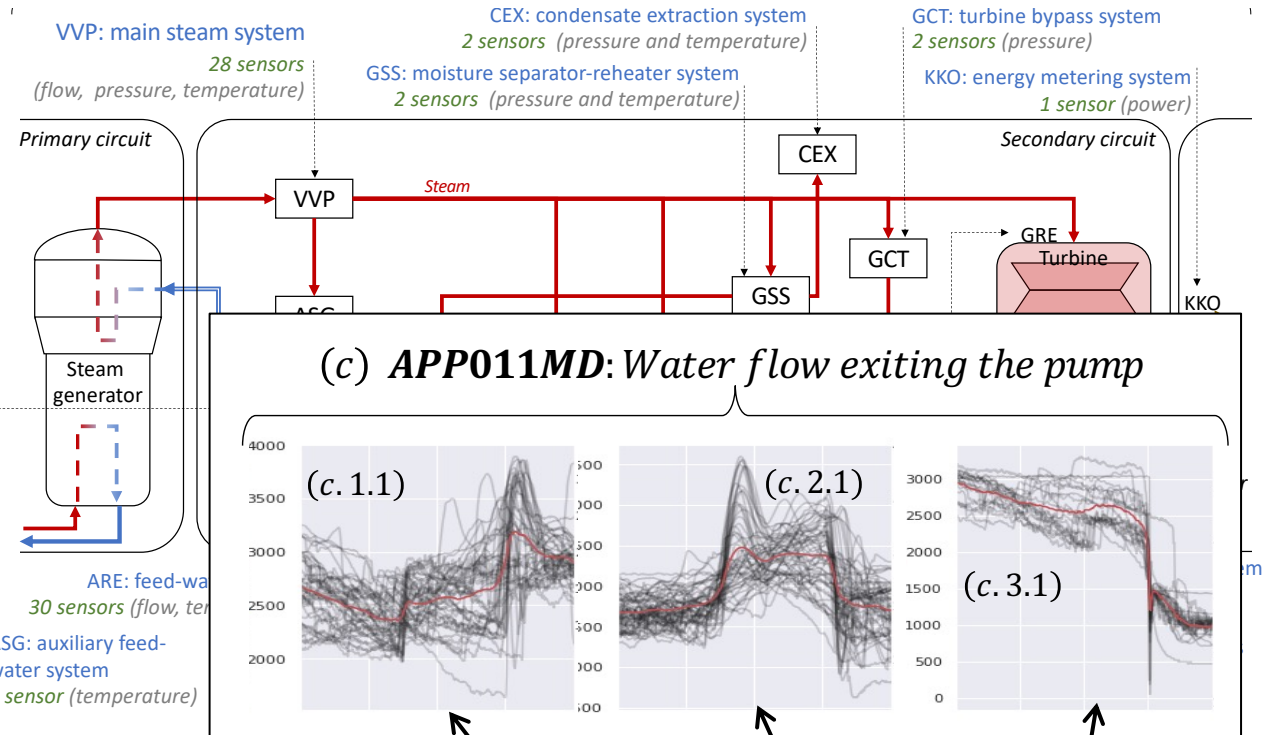
- Anomaly examples

Time Series T

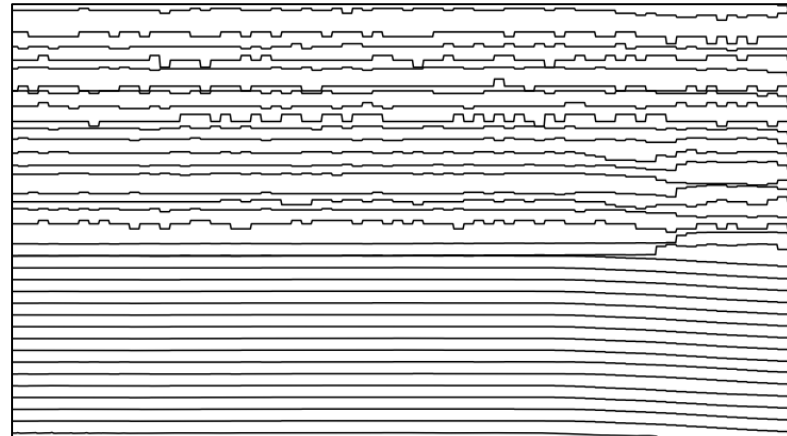
Supervised anomaly detection (e.g., classification)

Explanation of the detection

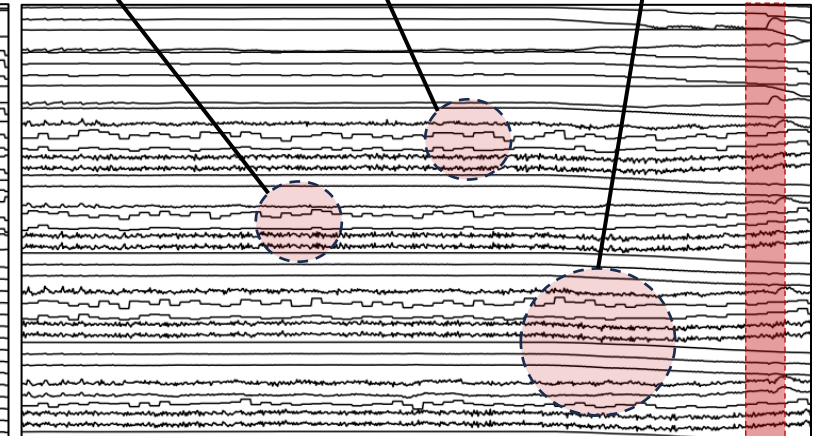
Identification of precursors



Class 1: Time series without any vibrations



Class 2: Time series with a vibrations



Vibration

Anomaly Detection

By inputs...

Time

Supervised

Training dataset

- Normal examples

- Anomaly examples

Time Series T

Supervised anomaly detection (e.g., classification)

Explanation of the detection

Identification of precursors

Class 1: Turbine

More info :

On the use case



DCE journal 2023

On the method



SIGMOD 2022



pump

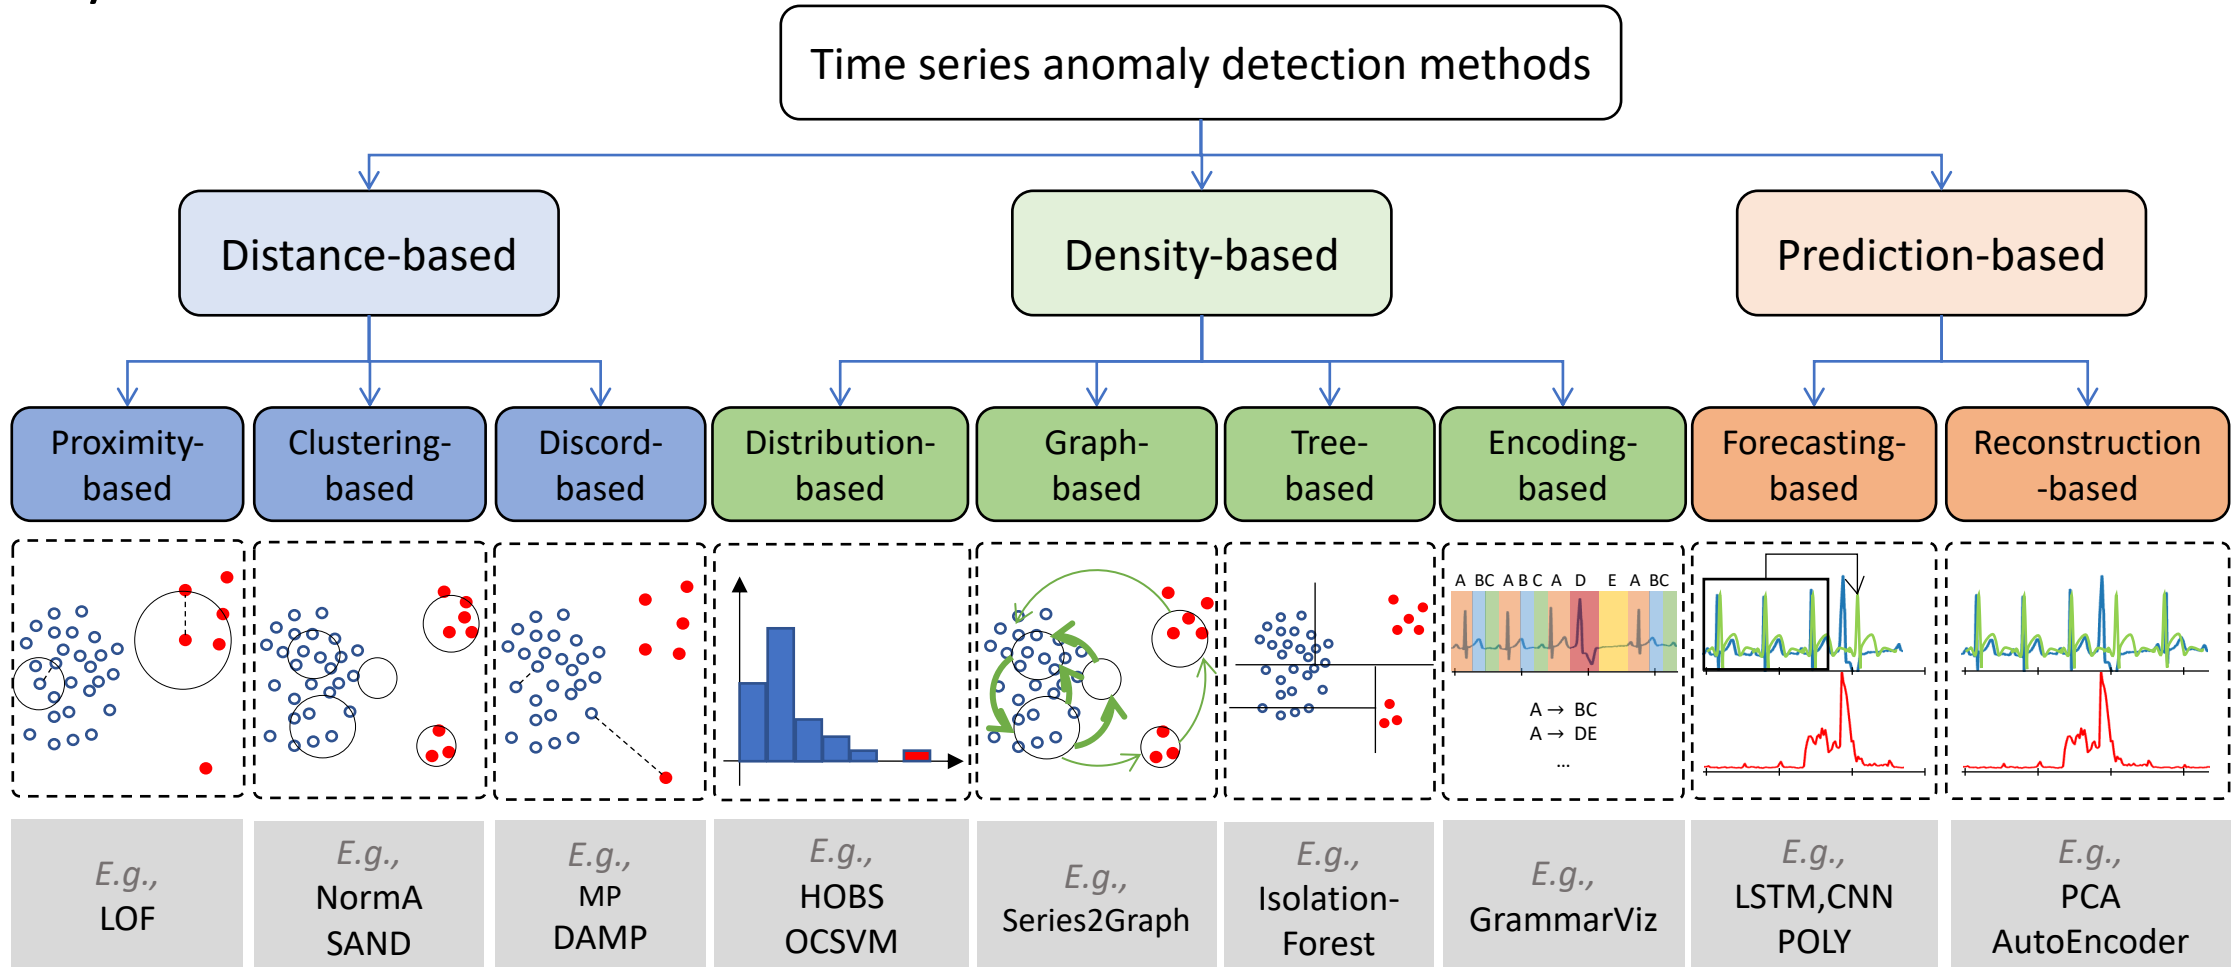
1)

vibrations

Vibration

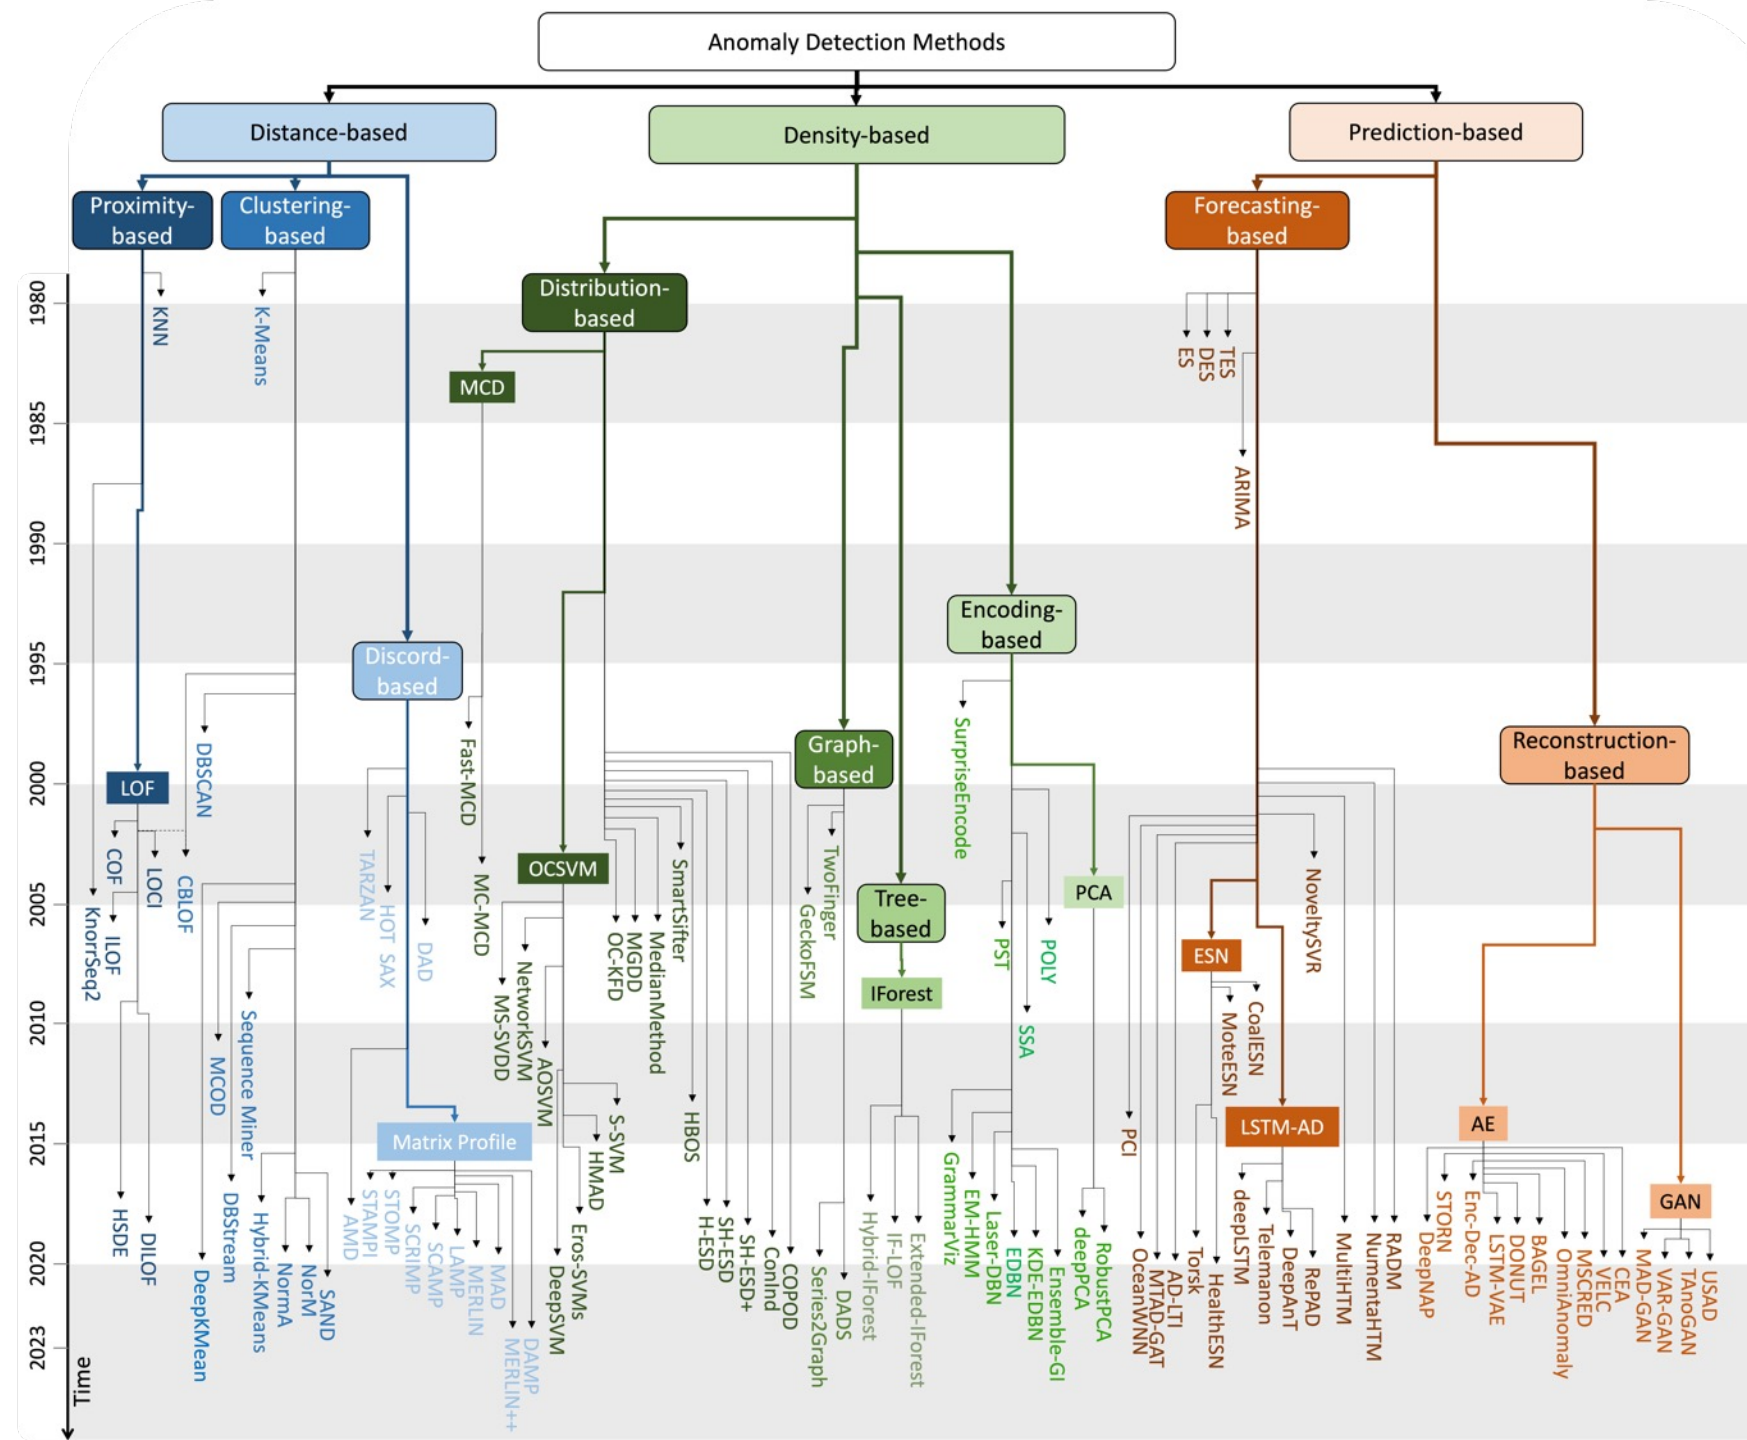
Anomaly Detection methods: *A taxonomy*

By methods...



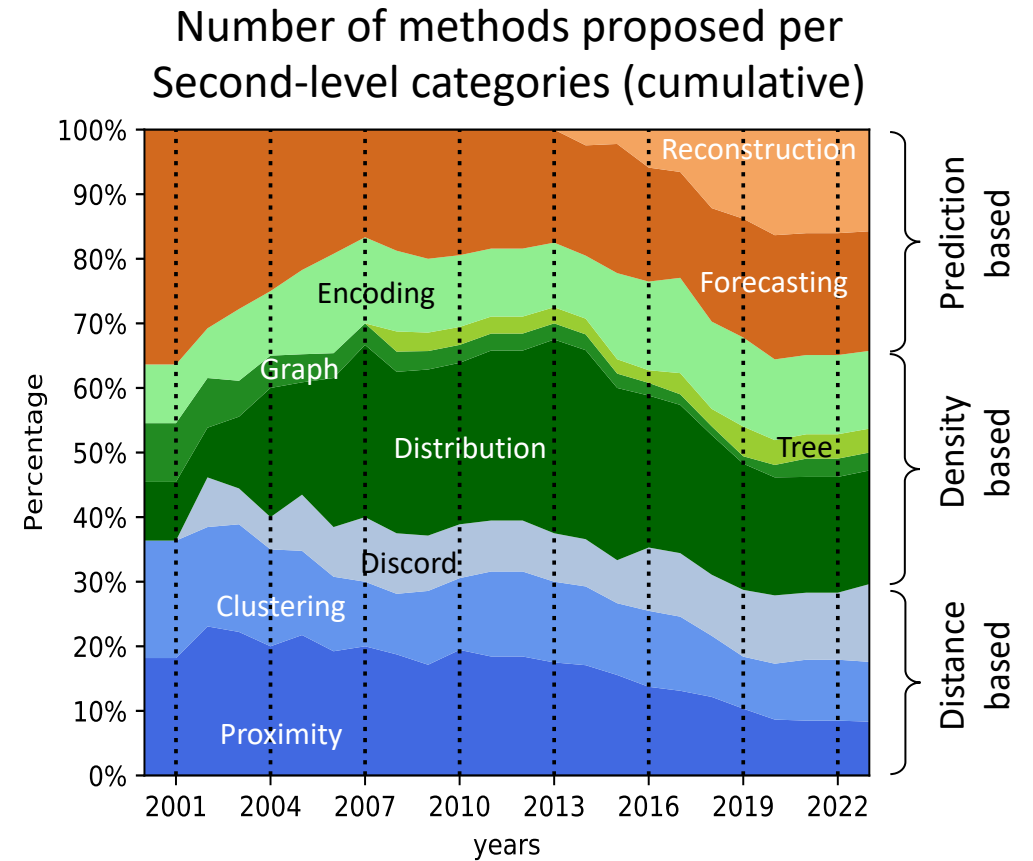
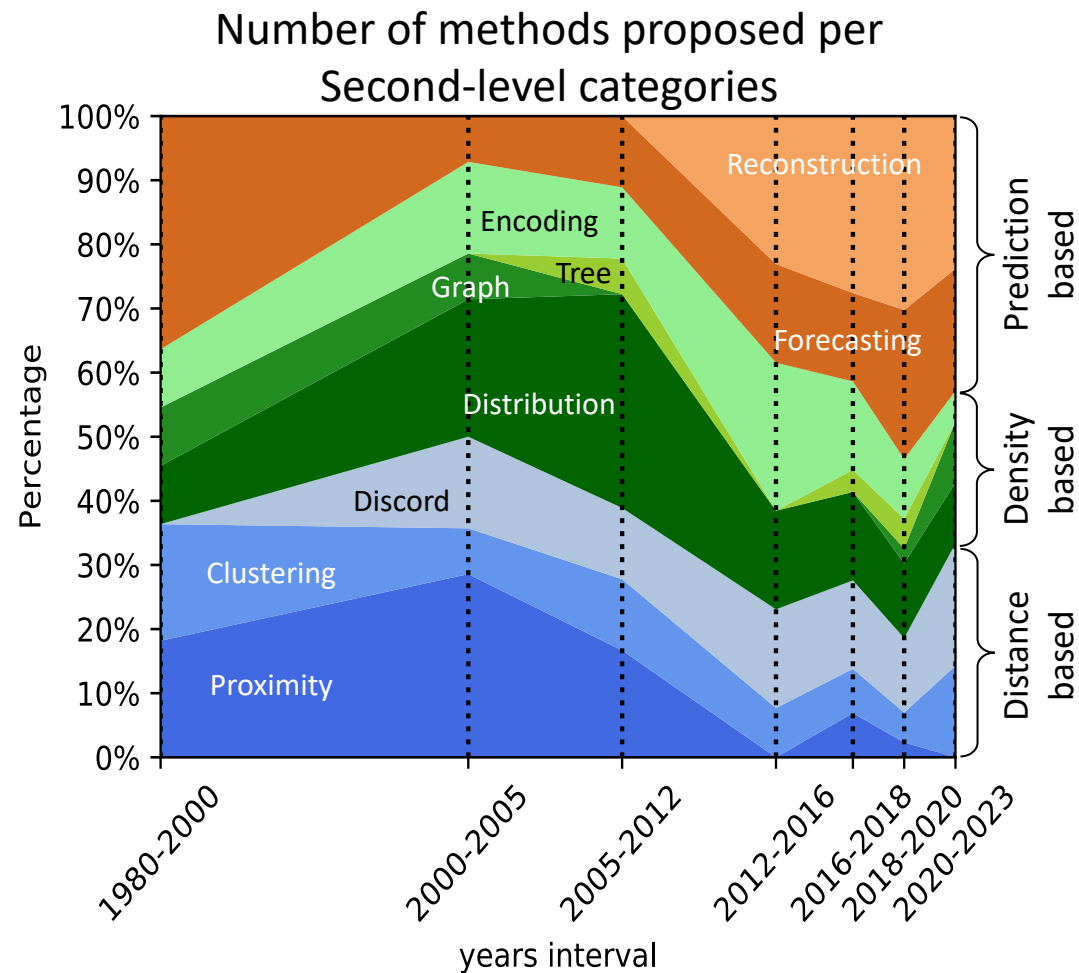
Anomaly Detection methods: *A taxonomy*

By time...



Anomaly Detection methods: *A taxonomy*

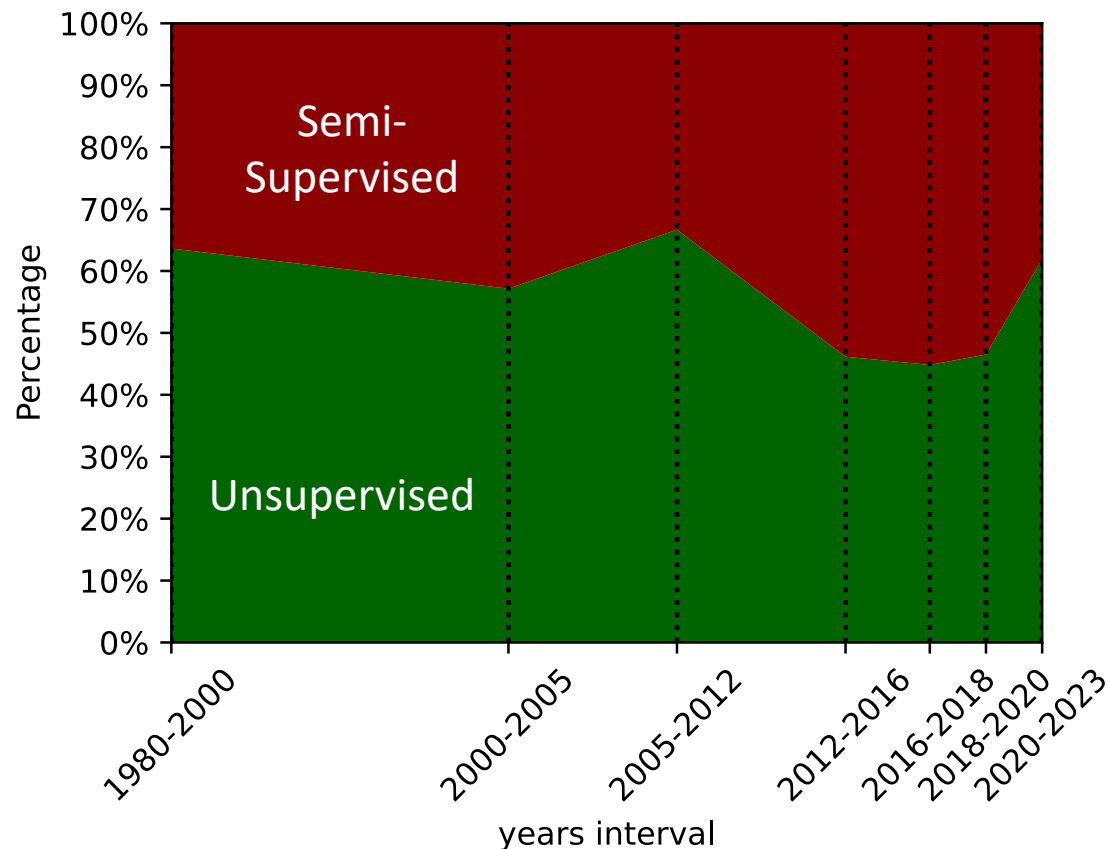
By time...



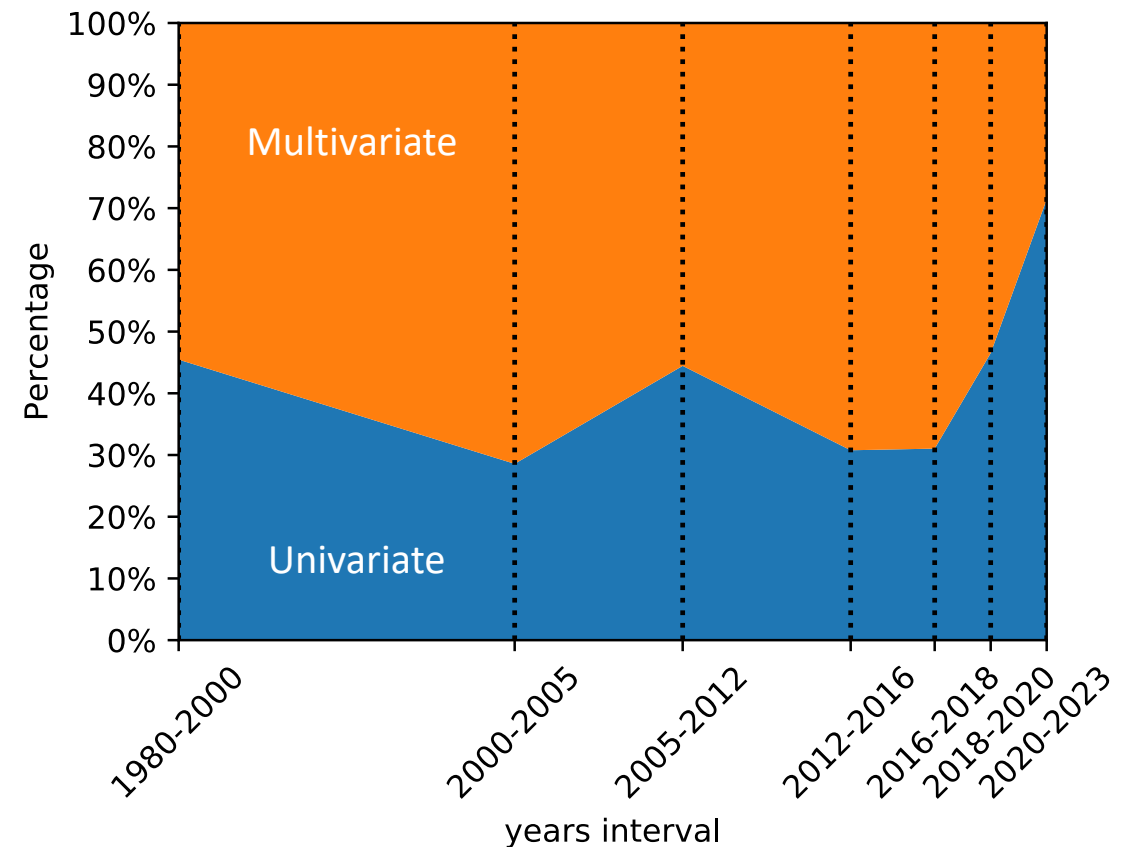
Anomaly Detection methods: *A taxonomy*

By time...

Number of methods proposed that are
Unsupervised or *Semi-Supervised*

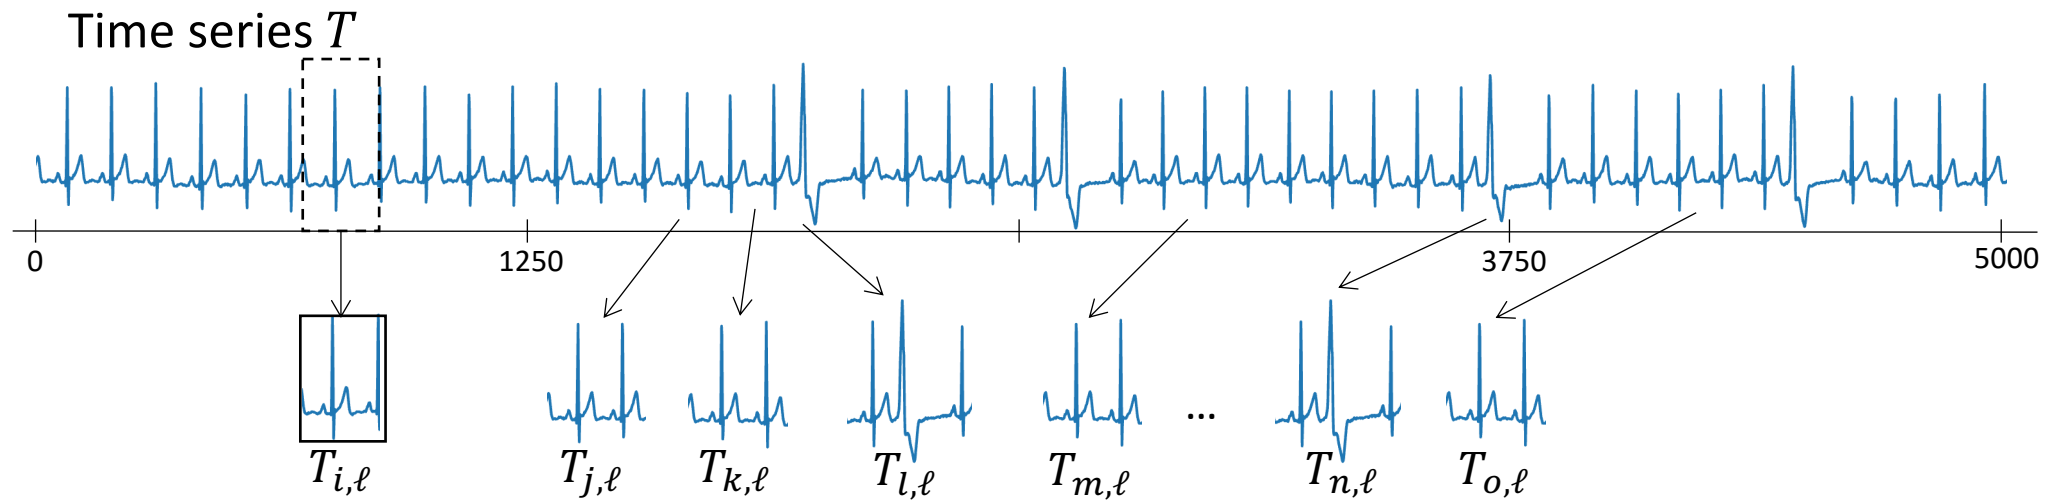


Number of methods proposed that can handle
Univariate or *Multivariate* time series



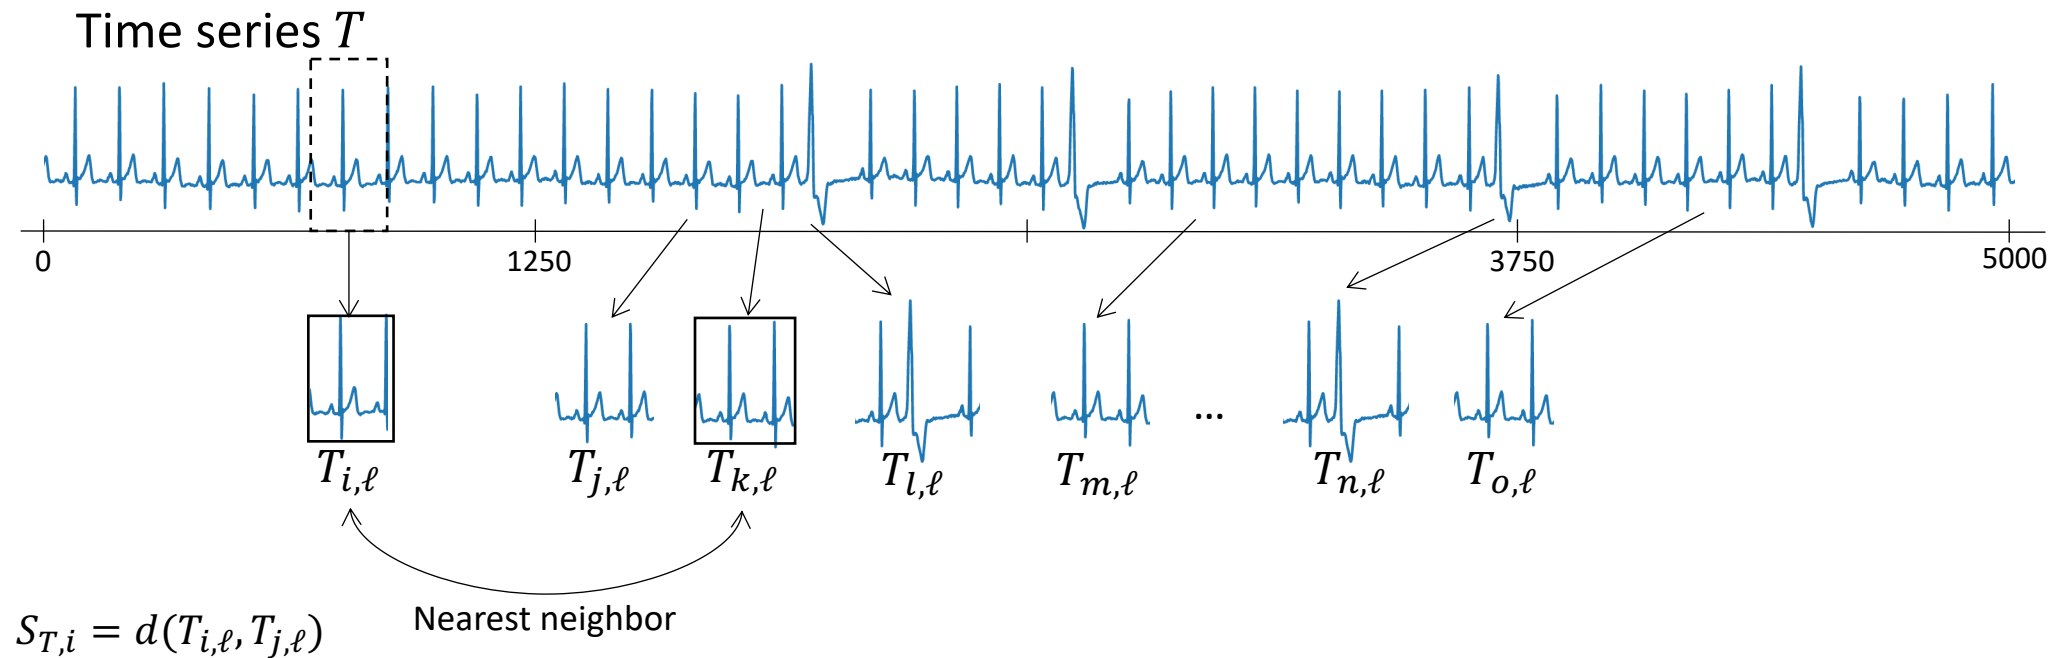
Anomaly Detection methods: *Distance-based*

Methods that use **distance computation** between subsequences (or group of subsequences) to detect anomalies.



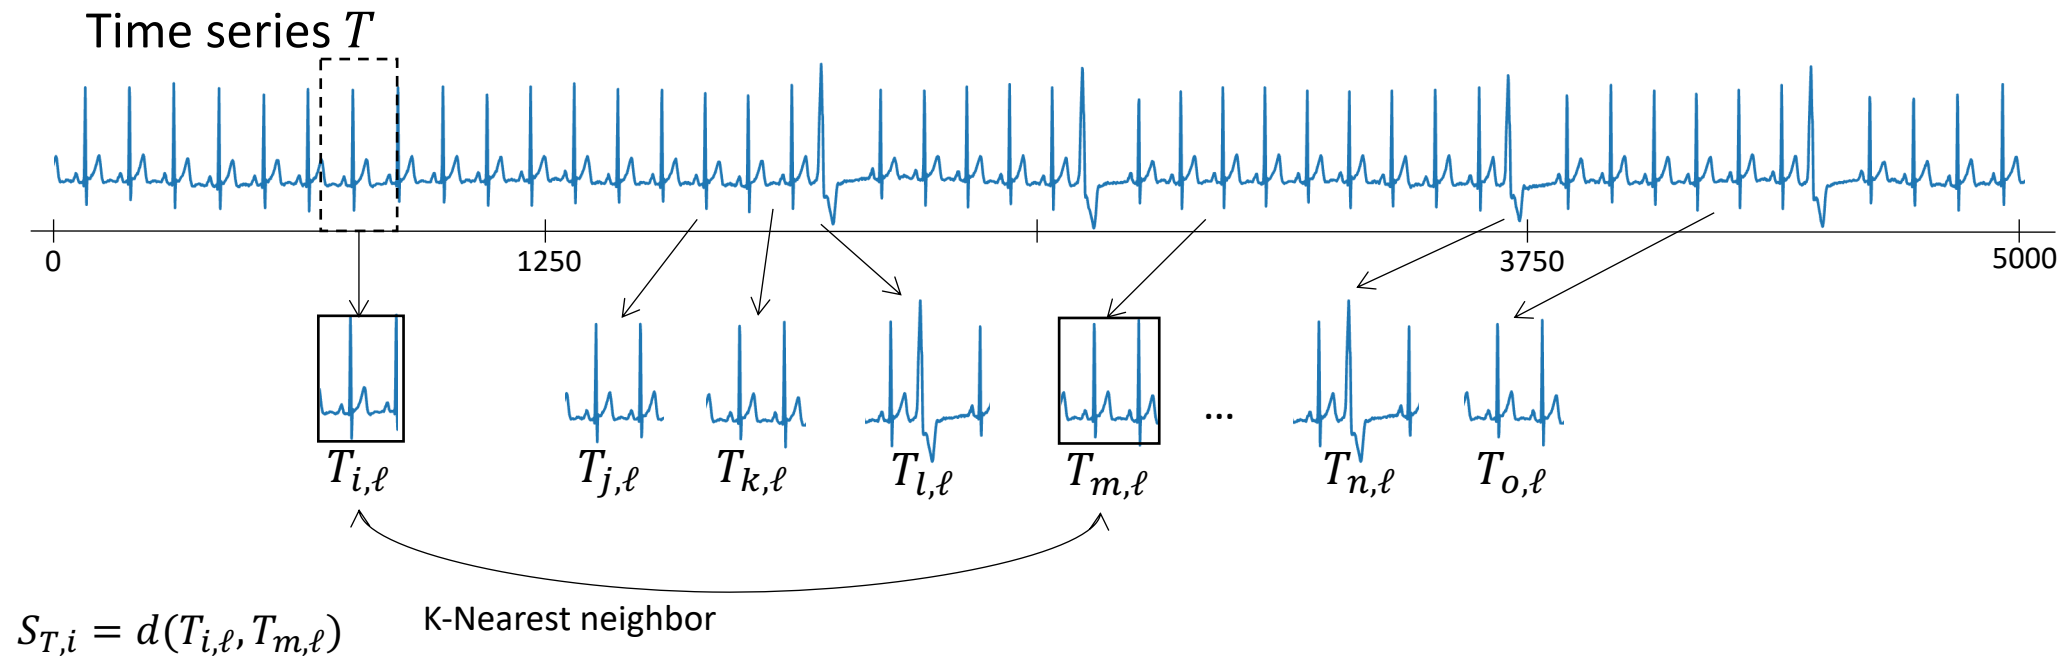
Anomaly Detection methods: *Distance-based*

Methods that use **distance computation** between subsequences (or group of subsequences) to detect anomalies.



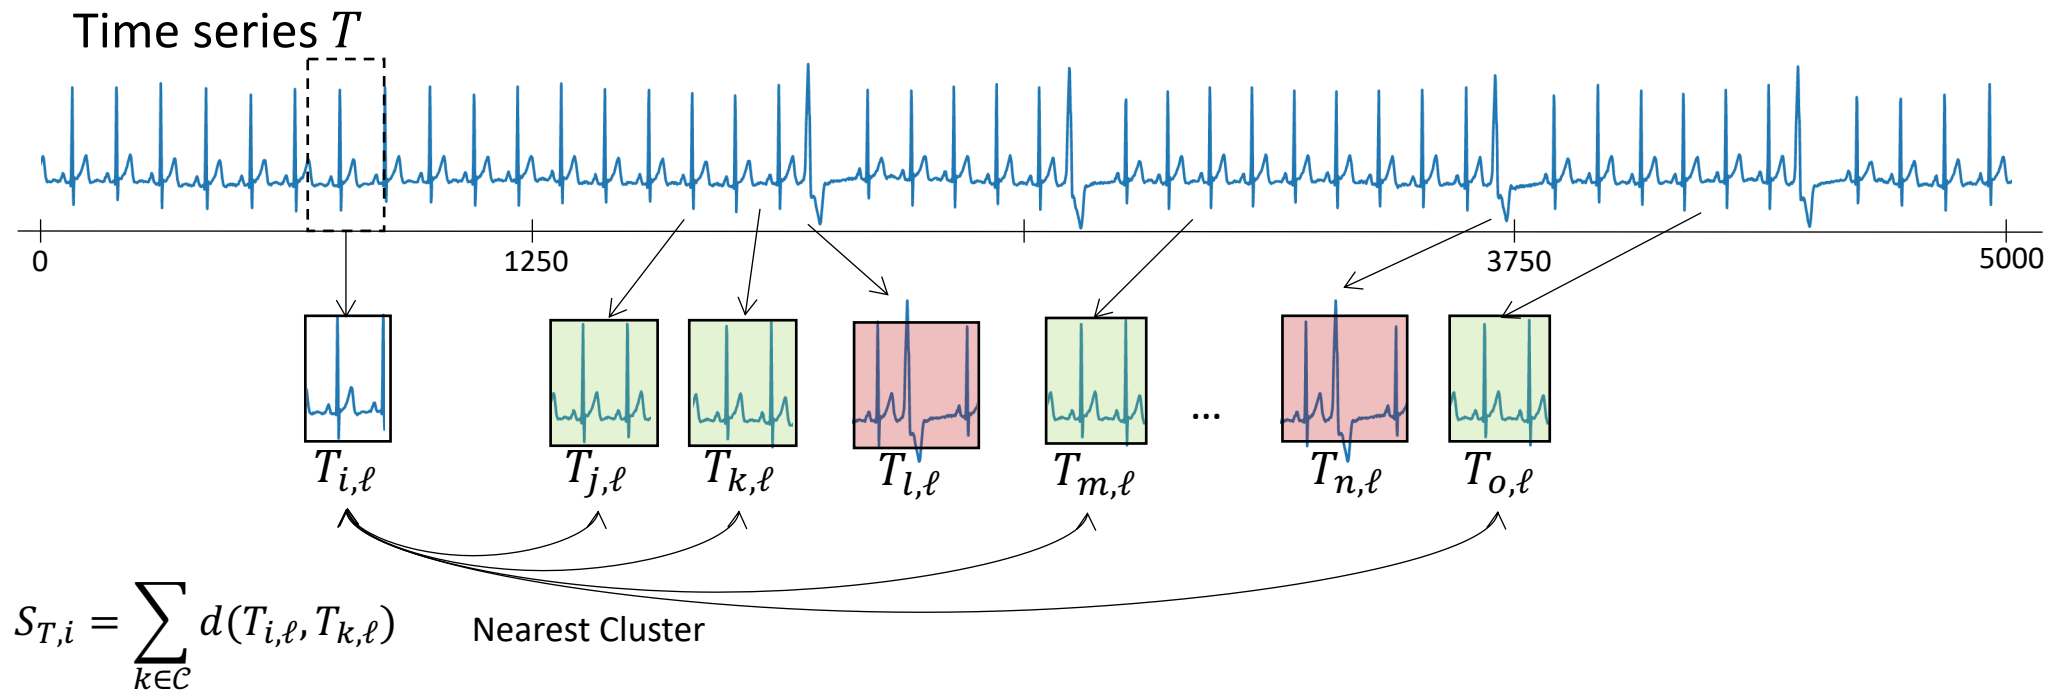
Anomaly Detection methods: *Distance-based*

Methods that use **distance computation** between subsequences (or group of subsequences) to detect anomalies.



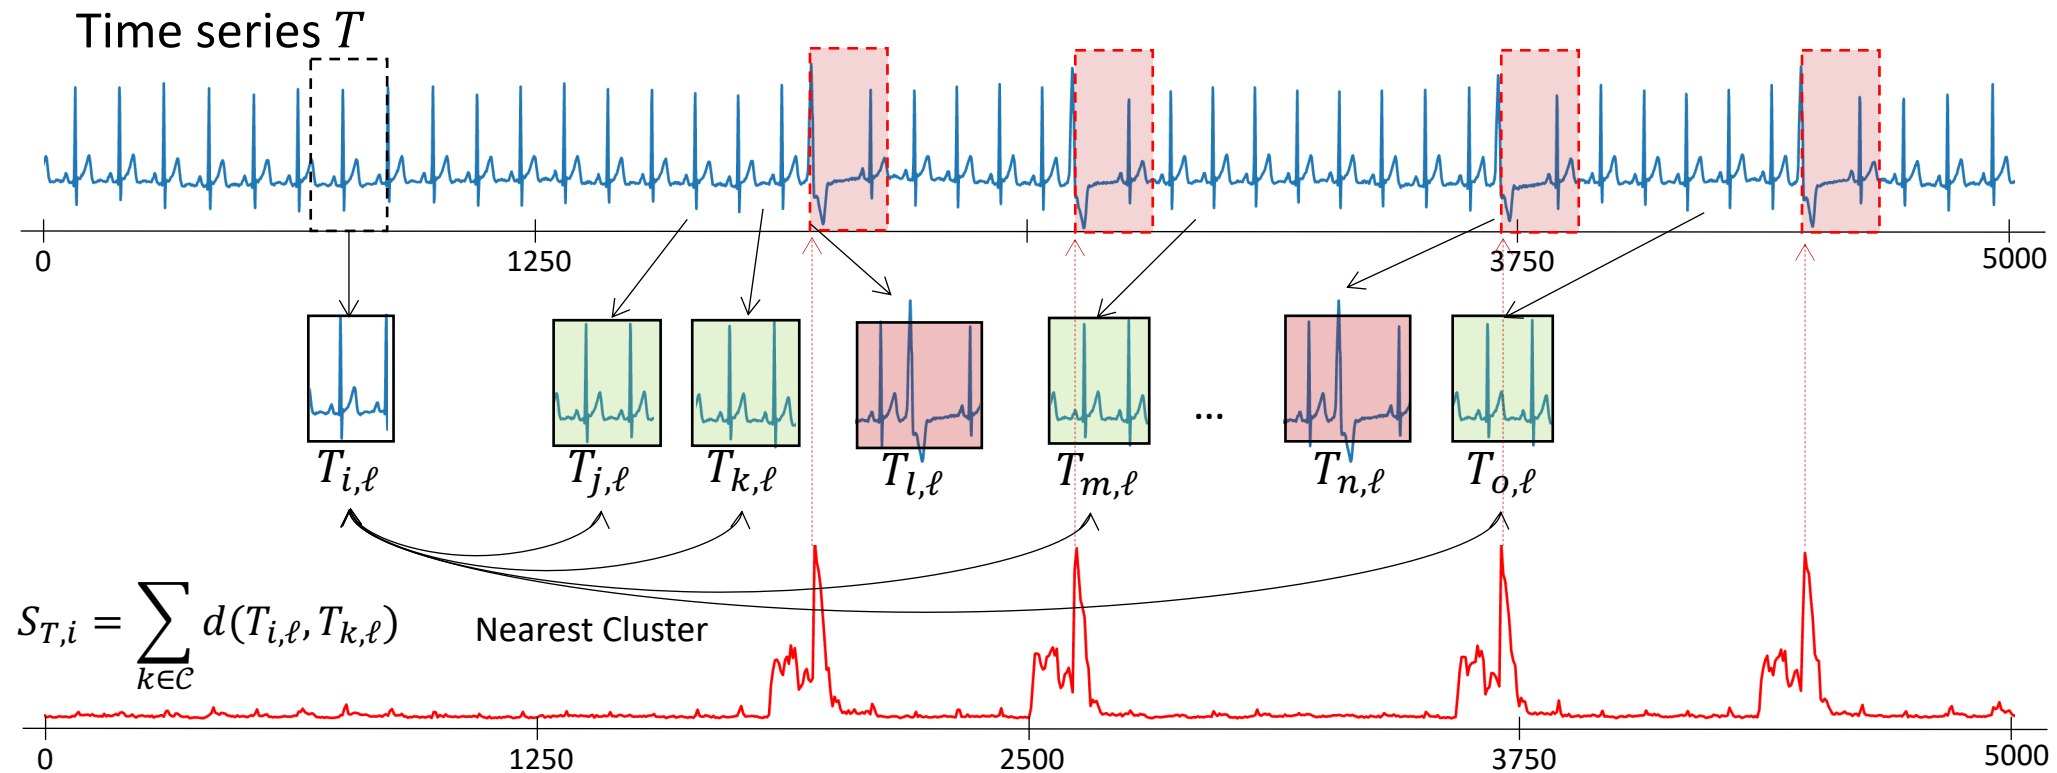
Anomaly Detection methods: *Distance-based*

Methods that use **distance computation** between subsequences (or group of subsequences) to detect anomalies.



Anomaly Detection methods: *Distance-based*

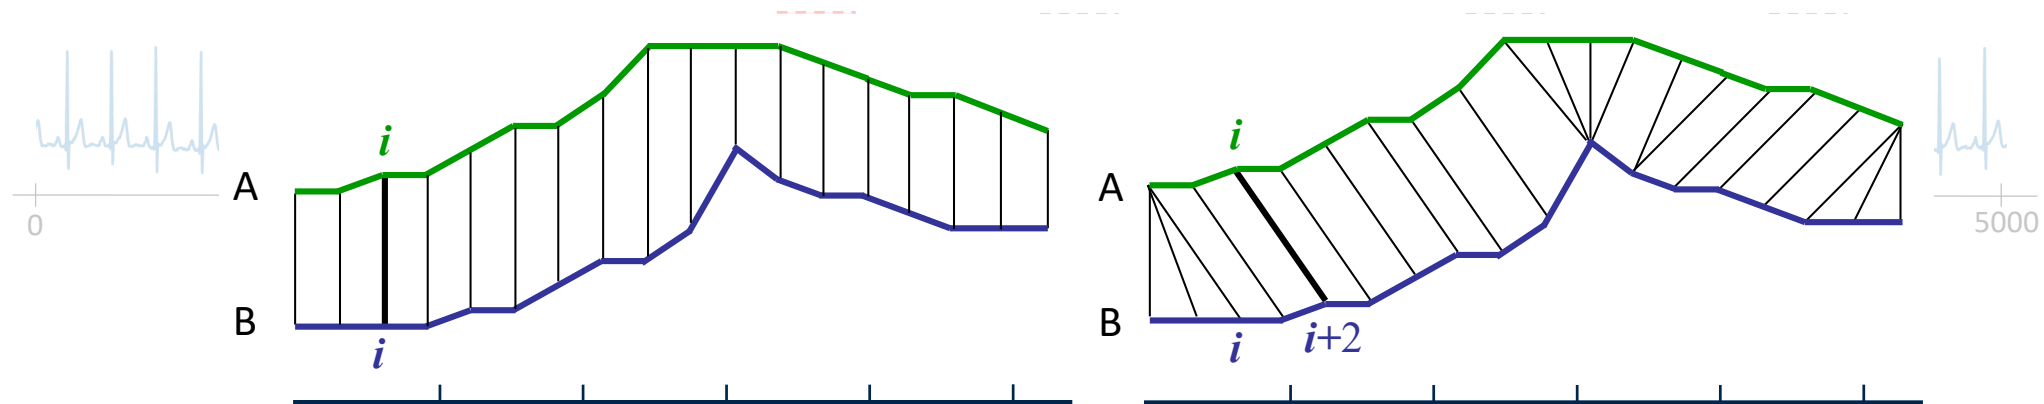
Methods that use **distance computation** between subsequences (or group of subsequences) to detect anomalies.



Anomaly Detection methods: *Distance-based*

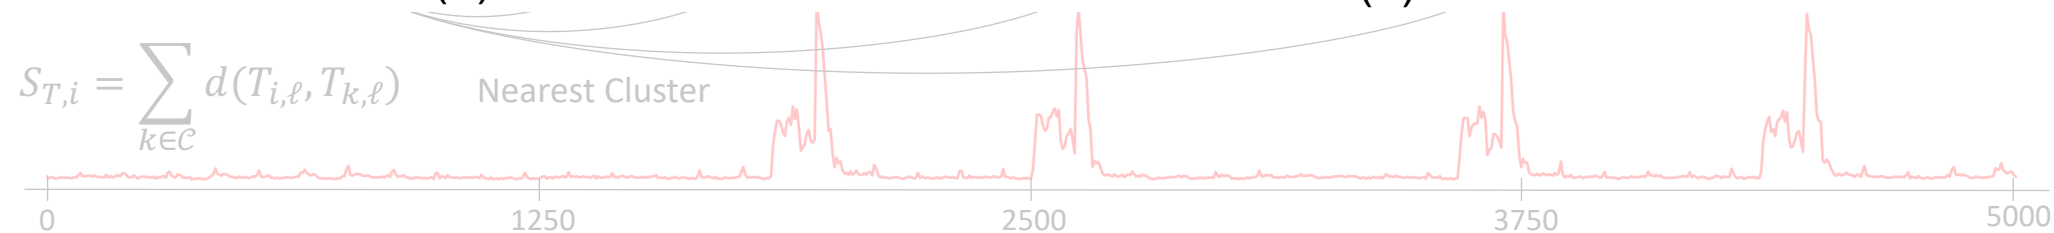
Methods that use **distance computation** between subsequences (or group of subsequences) to detect anomalies.

Example of distance computation

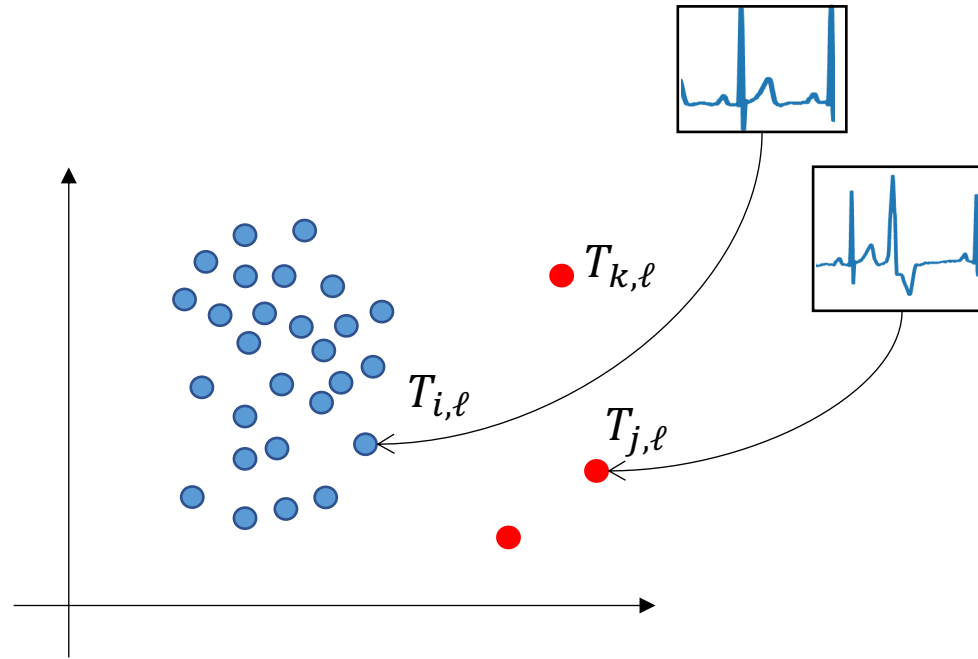


(a) Euclidian Distance

(b) DTW distance



Anomaly Detection methods: *an Example*



Matrix Profile [6] (MP)

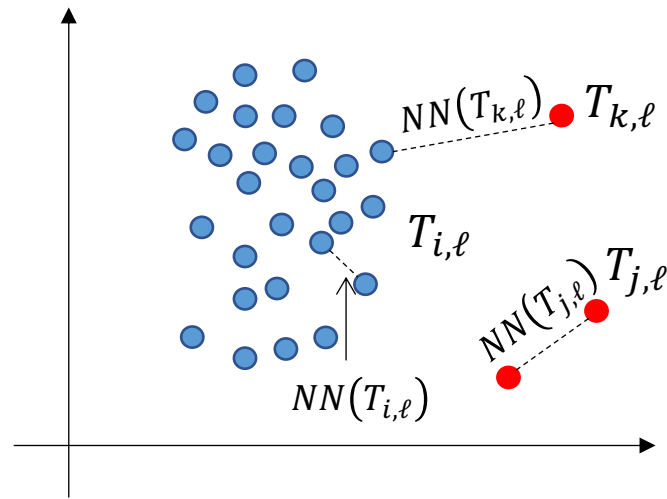
Compute the **distance to the nearest neighbor** (using the MASS algorithm z-norm Euclidean distance computation) and use it as anomaly score

Unsupervised

Univariate

sequence

Anomaly Detection methods: *an Example*



The matrix Profile is computed as follows:

$$S_T = [NN(T_{0,\ell}), NN(T_{1,\ell}), \dots, NN(T_{|T|-\ell,\ell})]$$

Matrix Profile [6] (MP)

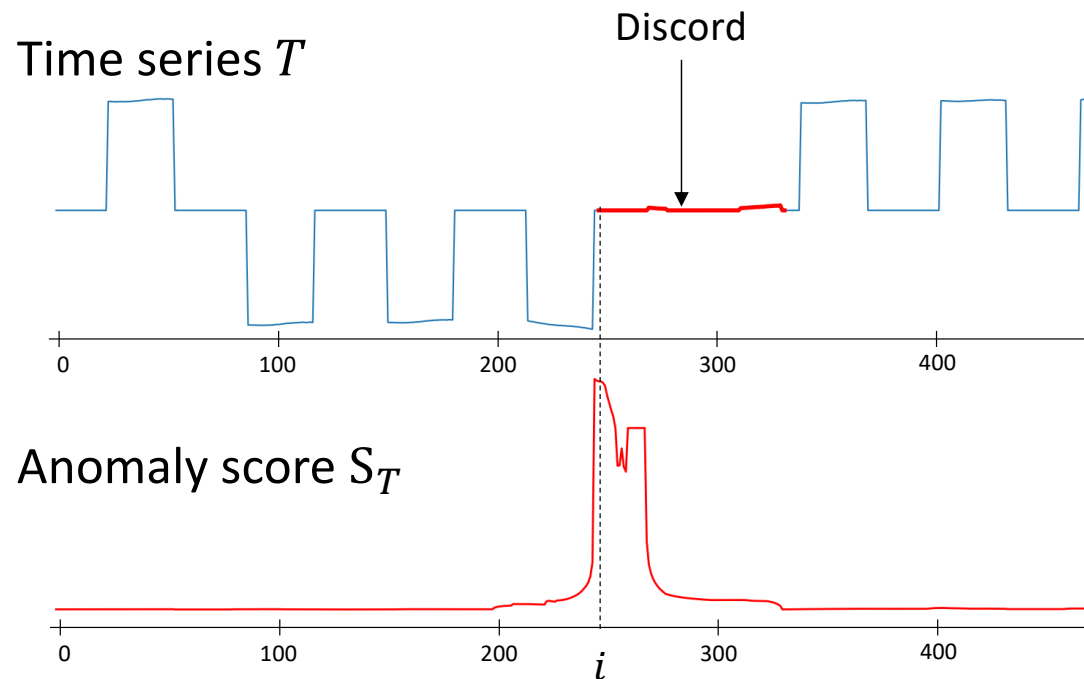
Compute the **distance to the nearest neighbor** (using the MASS algorithm z-norm Euclidean distance computation) and use it as anomaly score

Unsupervised

Univariate

sequence

Anomaly Detection methods: *an Example*



Matrix Profile [6] (MP)

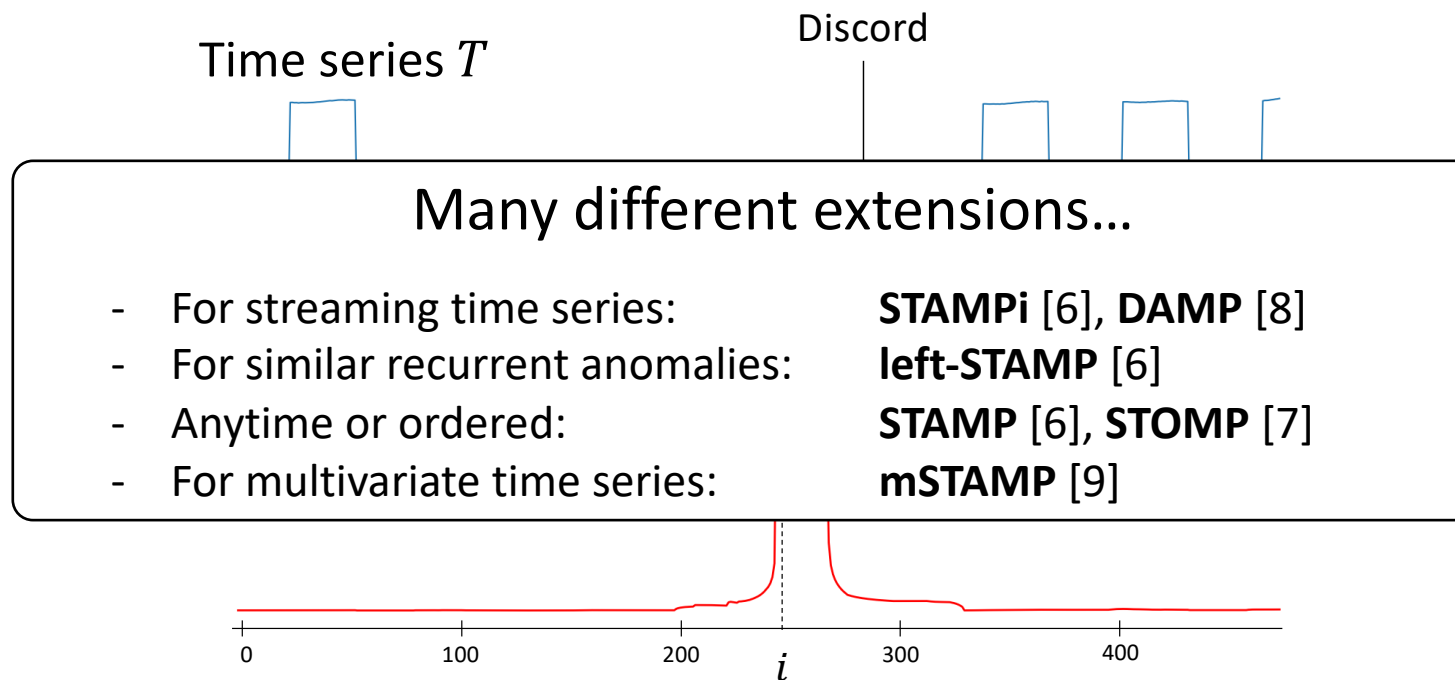
Compute the **distance to the nearest neighbor** (using the MASS algorithm z-norm Euclidean distance computation) and use it as anomaly score

Unsupervised

Univariate

sequence

Anomaly Detection methods: *an Example*



Matrix Profile [6] (MP)

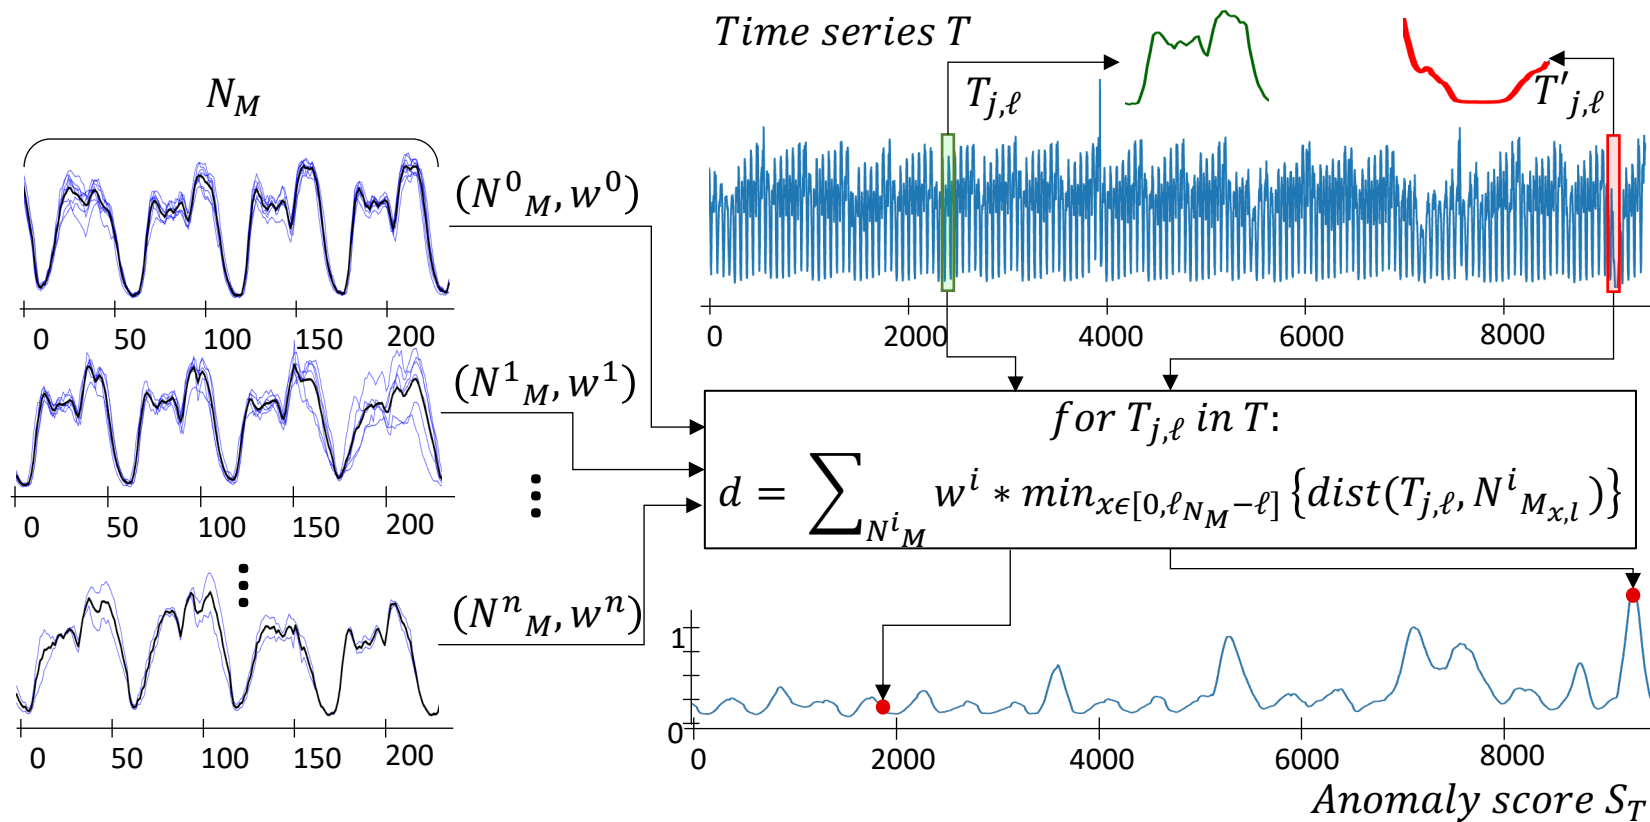
Compute the **distance to the nearest neighbor** (using the MASS algorithm z-norm Euclidean distance computation) and use it as anomaly score

Unsupervised

Univariate

sequence

Anomaly Detection methods: *an Example*



NormA [10]

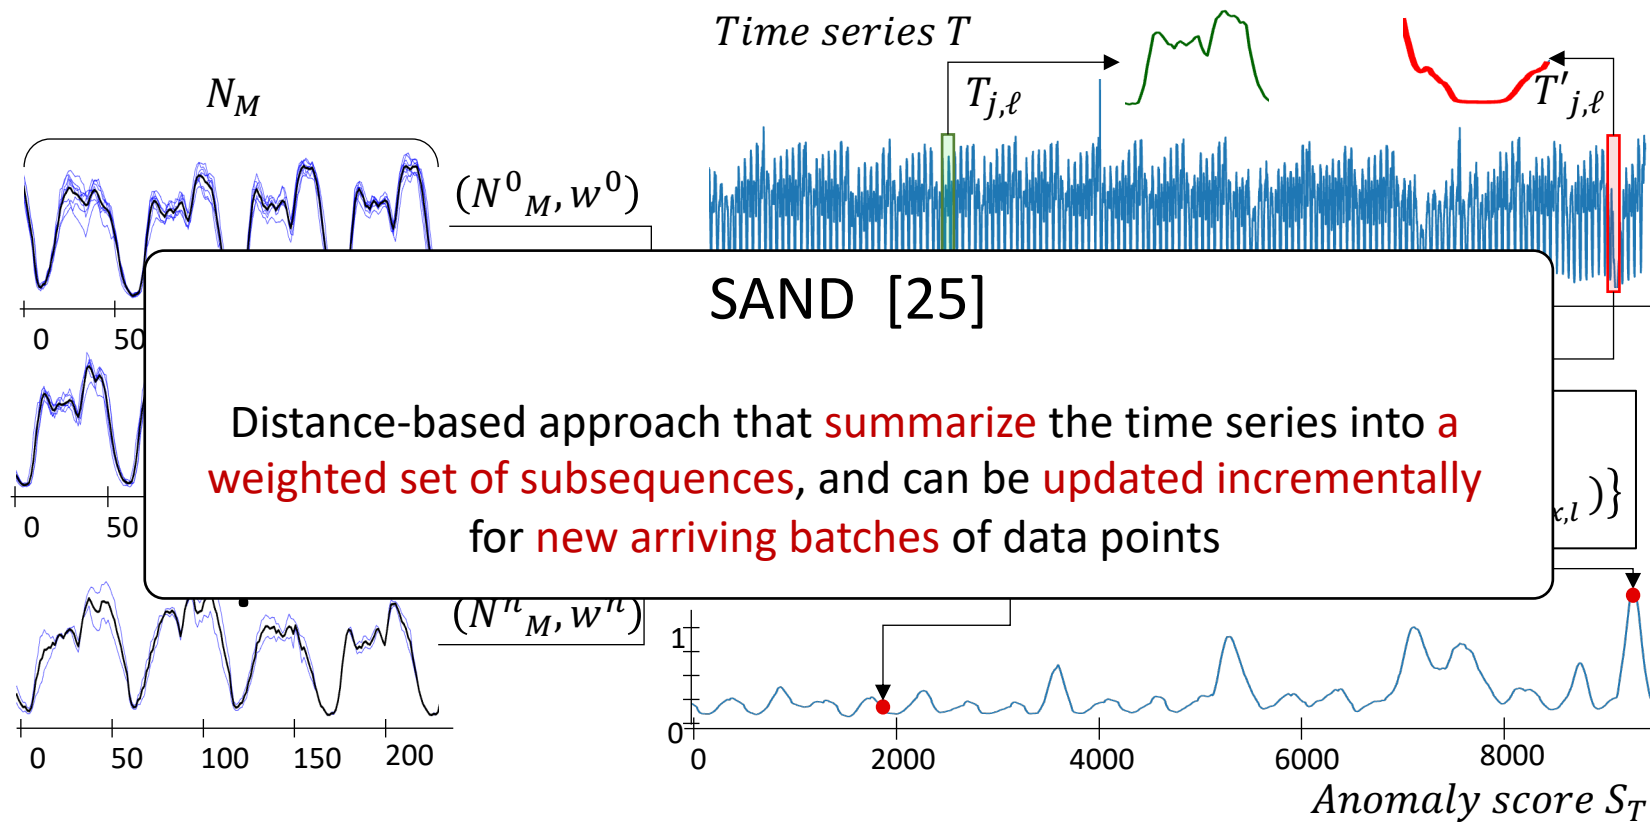
Distance-based approach that **summarize** the time series into **a weighted set of subsequences** and use the distance to them as anomaly score

Unsupervised

Univariate

sequence

Anomaly Detection methods: *an Example*



NormA [10]

Distance-based approach that **summarize** the time series into a **weighted set of subsequences** and use the distance to them as anomaly score

Unsupervised

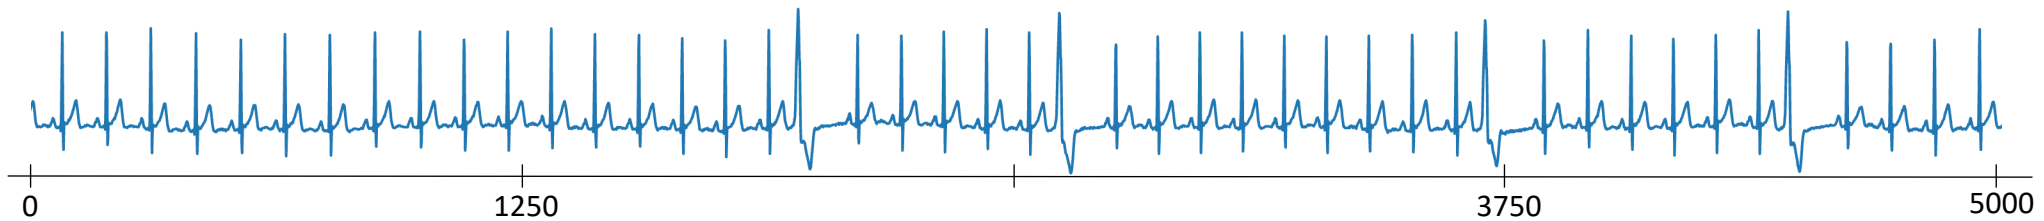
Univariate

sequence

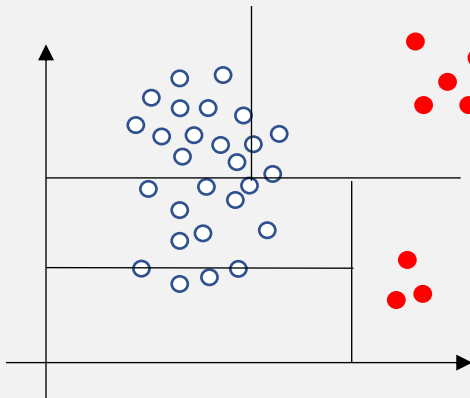
Anomaly Detection methods: *Density-based*

Methods that **estimate the density** of the space (points or subsequences) and identify as anomalies points (or sequences) that are in **low-density subspace**.

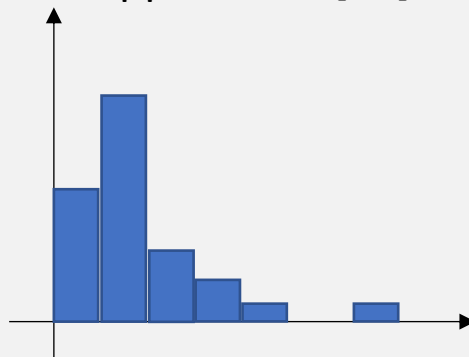
Time series T



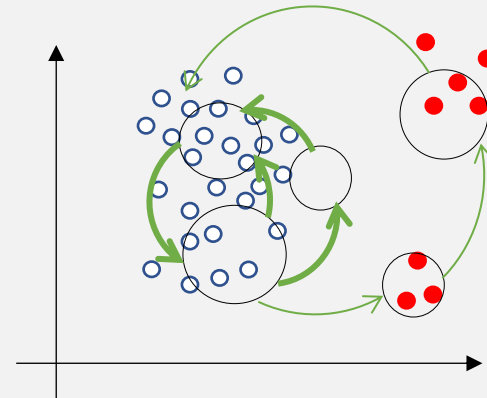
Tree-based approaches [11]



Distribution-based
Approaches [12]

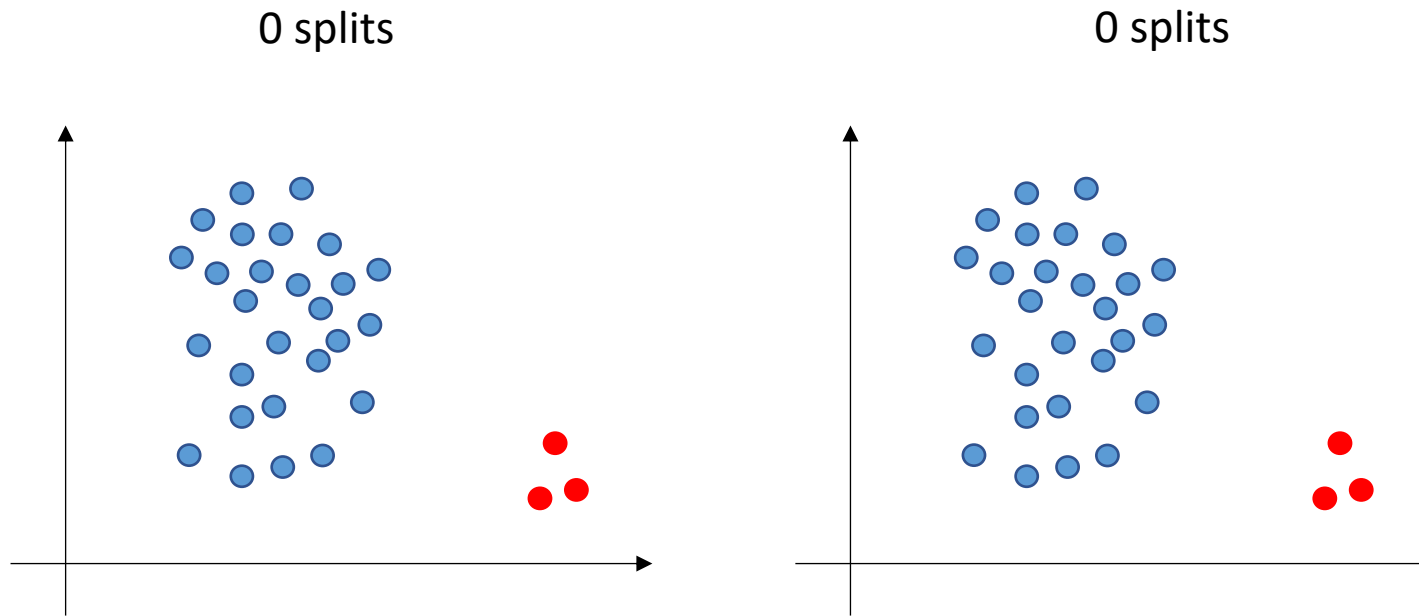


Graph-based approaches [13]



...

Anomaly Detection methods: *an Example*



Isolation Forest [11]

Density-based approach that **split the space randomly** and using the **depth of the trees** to identify anomalies

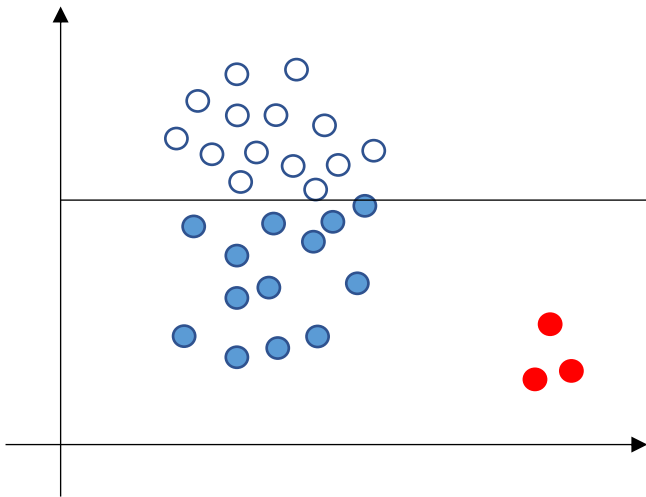
Unsupervised

Univariate/Multivariate

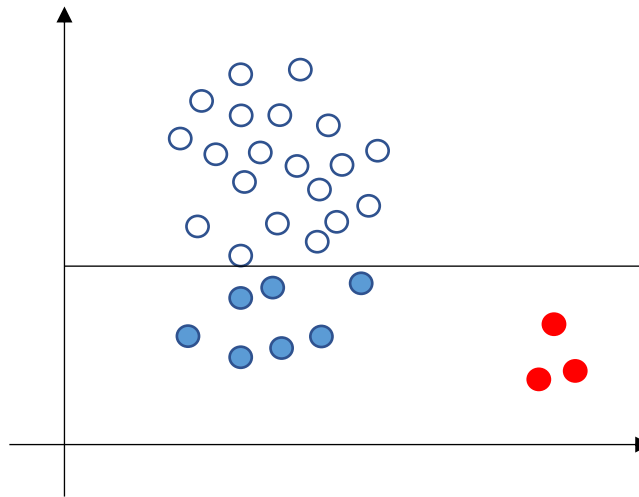
Point/sequence

Anomaly Detection methods: *an Example*

1 splits



1 splits



Isolation Forest [11]

Density-based approach that **split the space randomly** and using the **depth of the trees** to identify anomalies

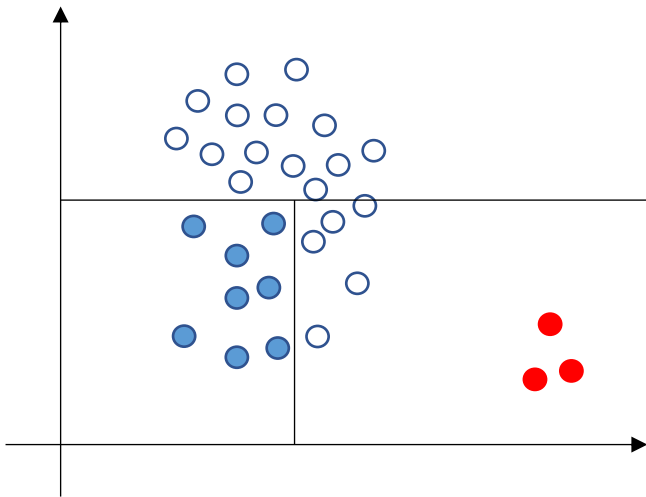
Unsupervised

Univariate/Multivariate

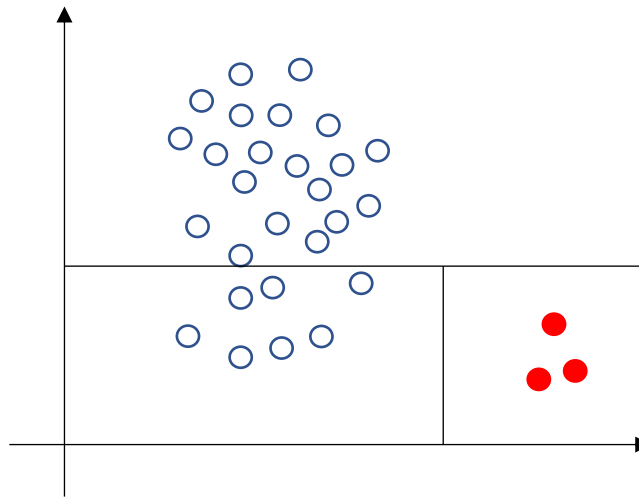
Point/sequence

Anomaly Detection methods: *an Example*

2 splits



2 splits



Isolation Forest [11]

Density-based approach that **split the space randomly** and using the **depth of the trees** to identify anomalies

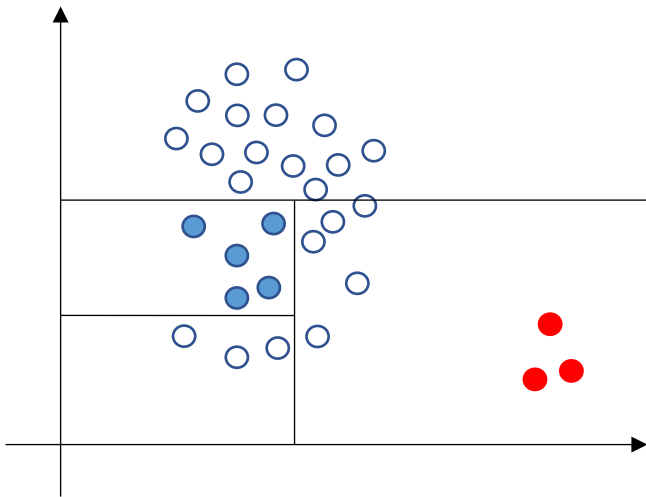
Unsupervised

Univariate/Multivariate

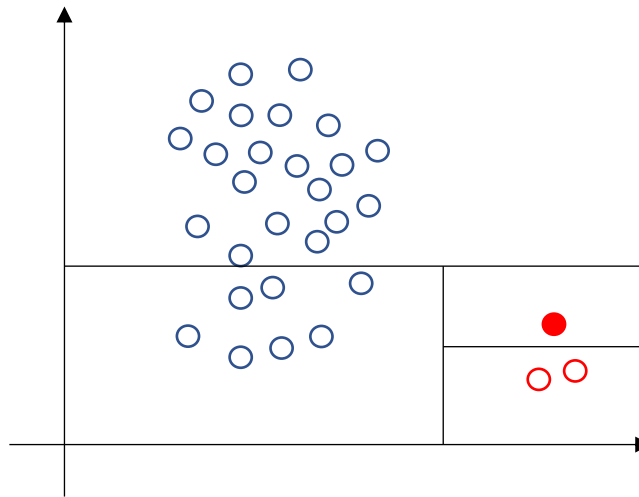
Point/sequence

Anomaly Detection methods: *an Example*

3 splits



3 splits



Isolation Forest [11]

Density-based approach that **split the space randomly** and using the **depth of the trees** to identify anomalies

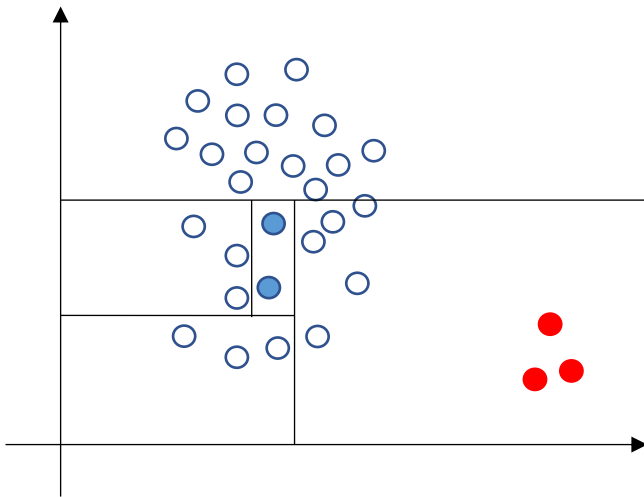
Unsupervised

Univariate/Multivariate

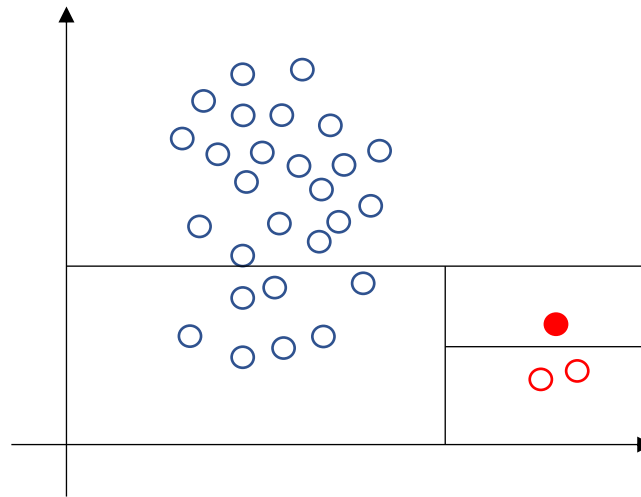
Point/sequence

Anomaly Detection methods: *an Example*

4 splits



3 splits



Isolation Forest [11]

Density-based approach that **split the space randomly** and using the **depth of the trees** to identify anomalies

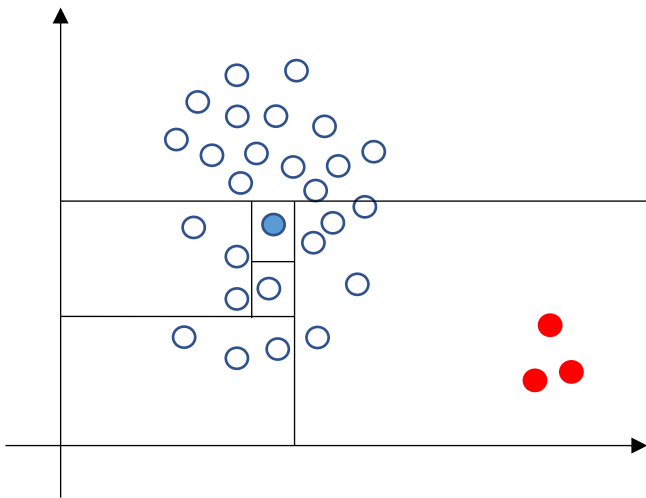
Unsupervised

Univariate/Multivariate

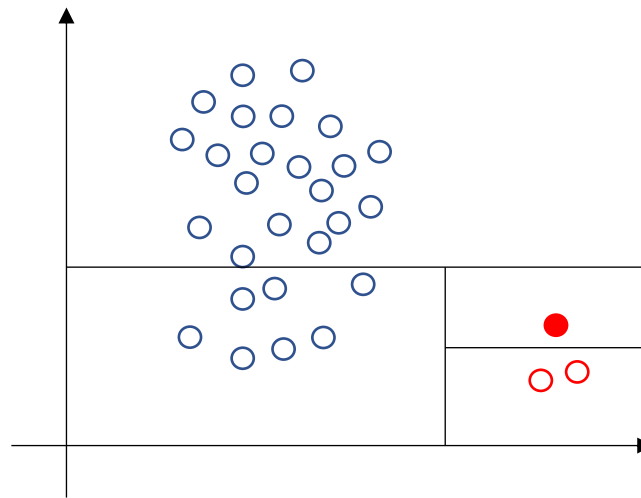
Point/sequence

Anomaly Detection methods: *an Example*

5 splits



3 splits



Isolation Forest [11]

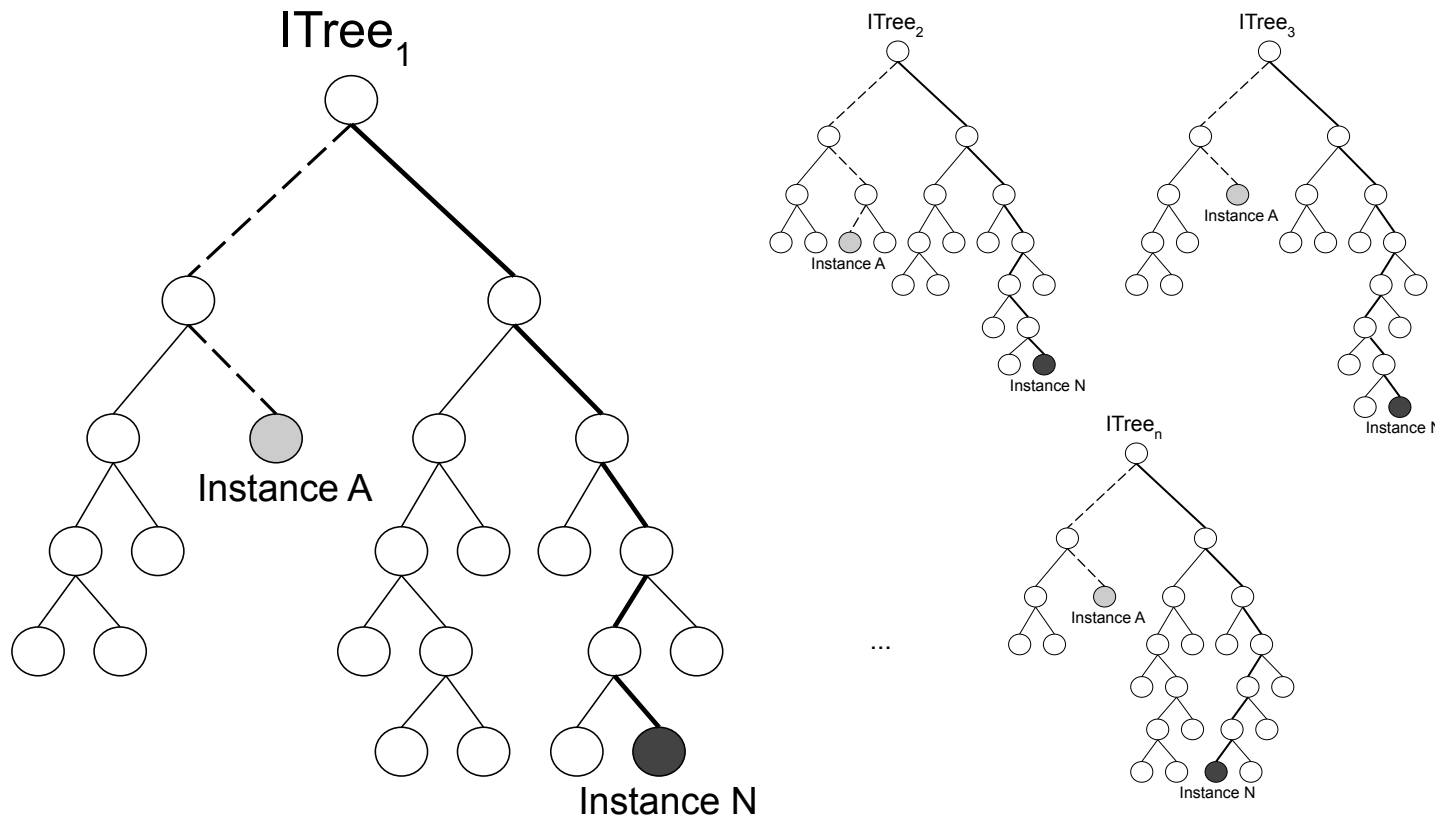
Density-based approach that **split the space randomly** and using the **depth of the trees** to identify anomalies

Unsupervised

Univariate/Multivariate

Point/sequence

Anomaly Detection methods: *an Example*



Isolation Forest [11]

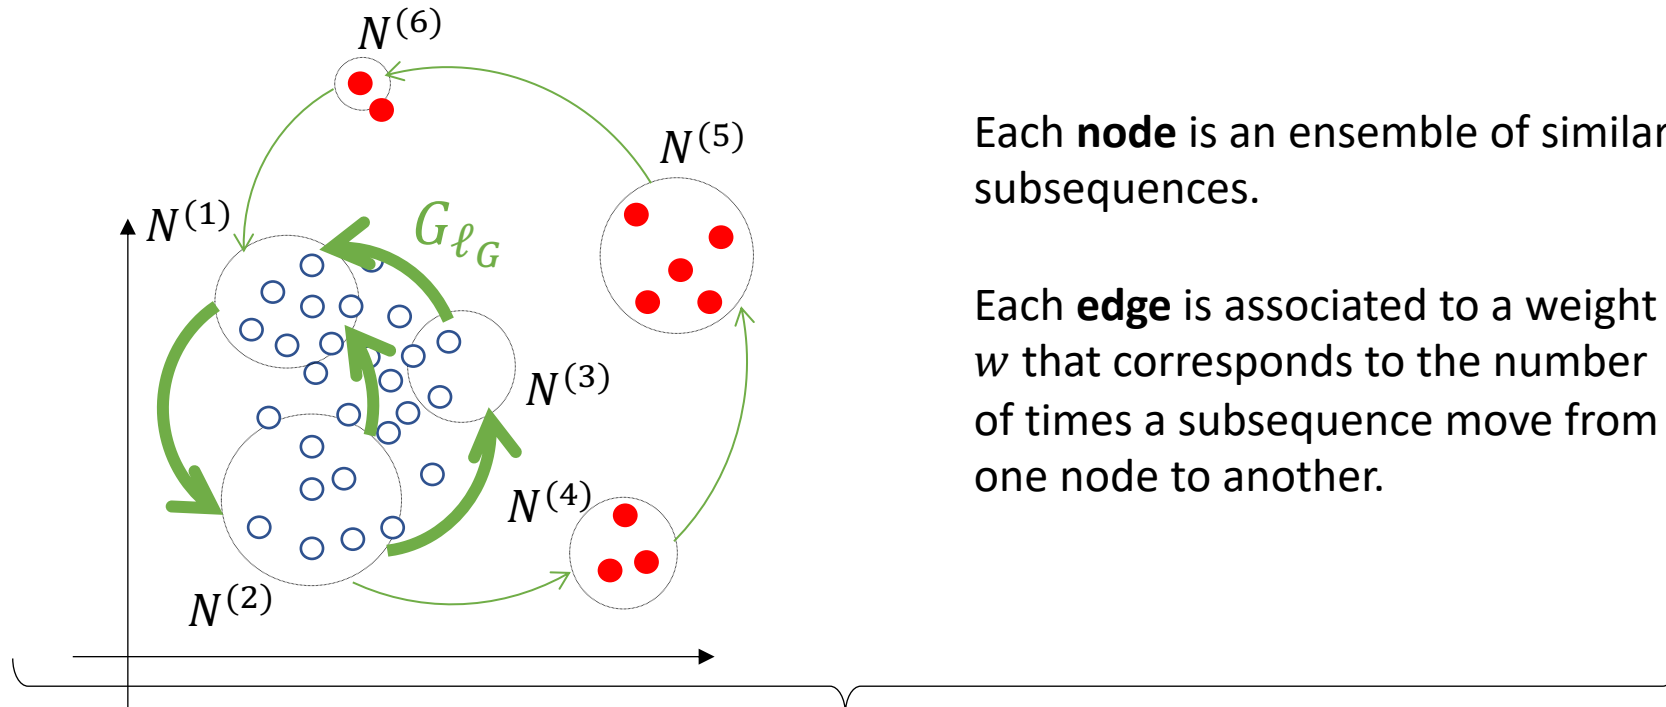
Density-based approach that **split the space randomly** and using the **depth of the trees** to identify anomalies

Unsupervised

Univariate/Multivariate

Point/sequence

Anomaly Detection methods: *an Example*



For a given subsequence $T_{i,\ell}$ and its corresponding path $P_{th} = \langle N^{(i)}, N^{(i+1)}, \dots, N^{(i+\ell)} \rangle$, we define the normality score as follows:

$$Norm(P_{th}) = \sum_{j=i}^{i+\ell-1} \frac{w(N^{(j)}, N^{(j+1)}) \deg(N^{(j)} - 1)}{\ell}$$

Series2Graph [13]

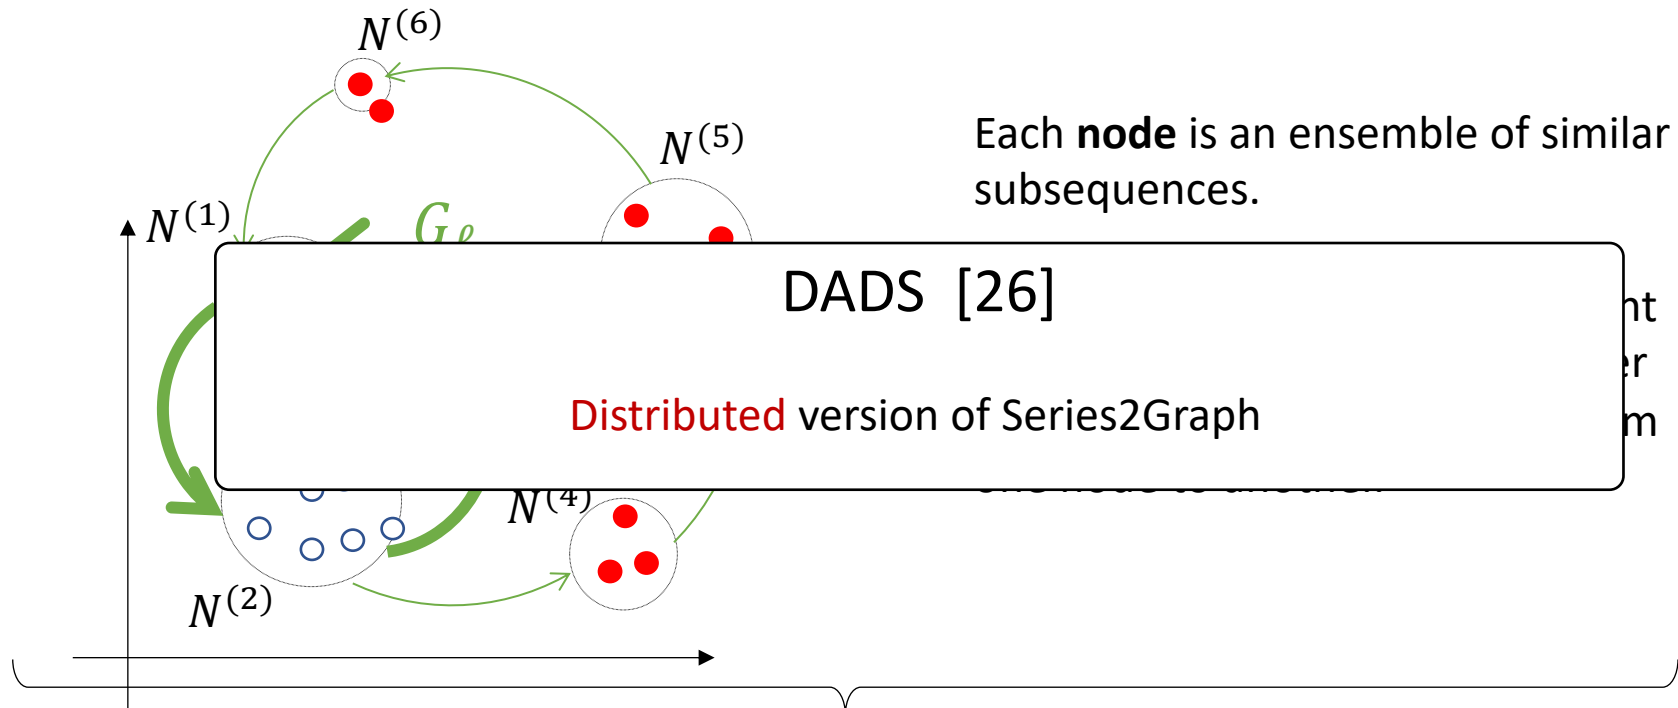
Density-based approach that **convert** the time series into a **graph** and detect **unusual trajectories**

Unsupervised

Univariate

subsequence

Anomaly Detection methods: *an Example*



For a given subsequence $T_{i,\ell}$ and its corresponding path $P_{th} = \langle N^{(i)}, N^{(i+1)}, \dots, N^{(i+\ell)} \rangle$, we define the normality score as follows:

$$Norm(P_{th}) = \sum_{j=i}^{i+\ell-1} \frac{w(N^{(j)}, N^{(j+1)}) \deg(N^{(j)} - 1)}{\ell}$$

Series2Graph [13]

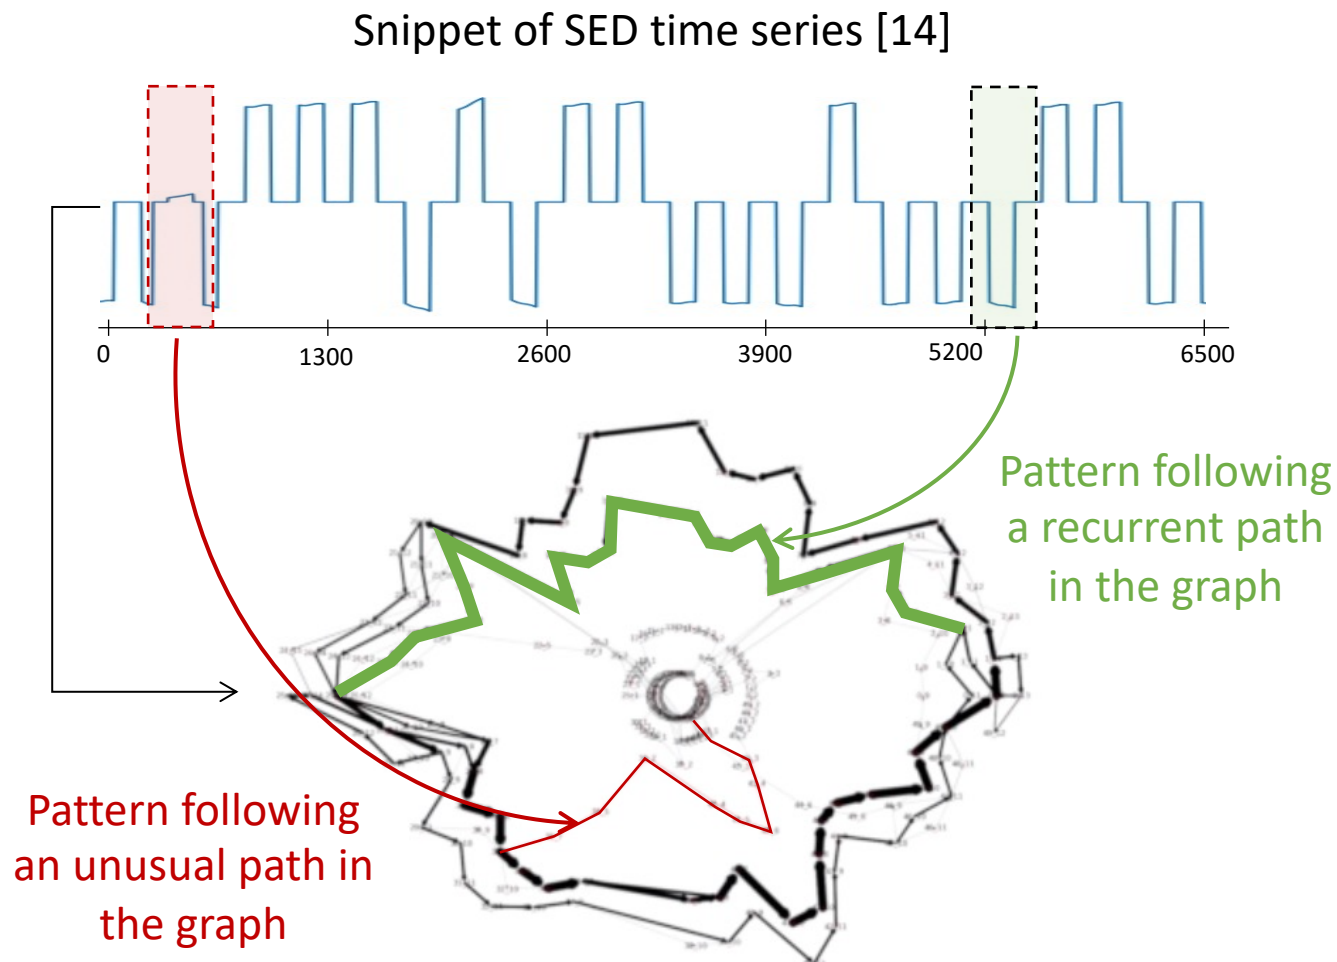
Density-based approach that **convert** the time series into a **graph** and detect **unusual trajectories**

Unsupervised

Univariate

subsequence

Anomaly Detection methods: *an Example*



Series2Graph [13]

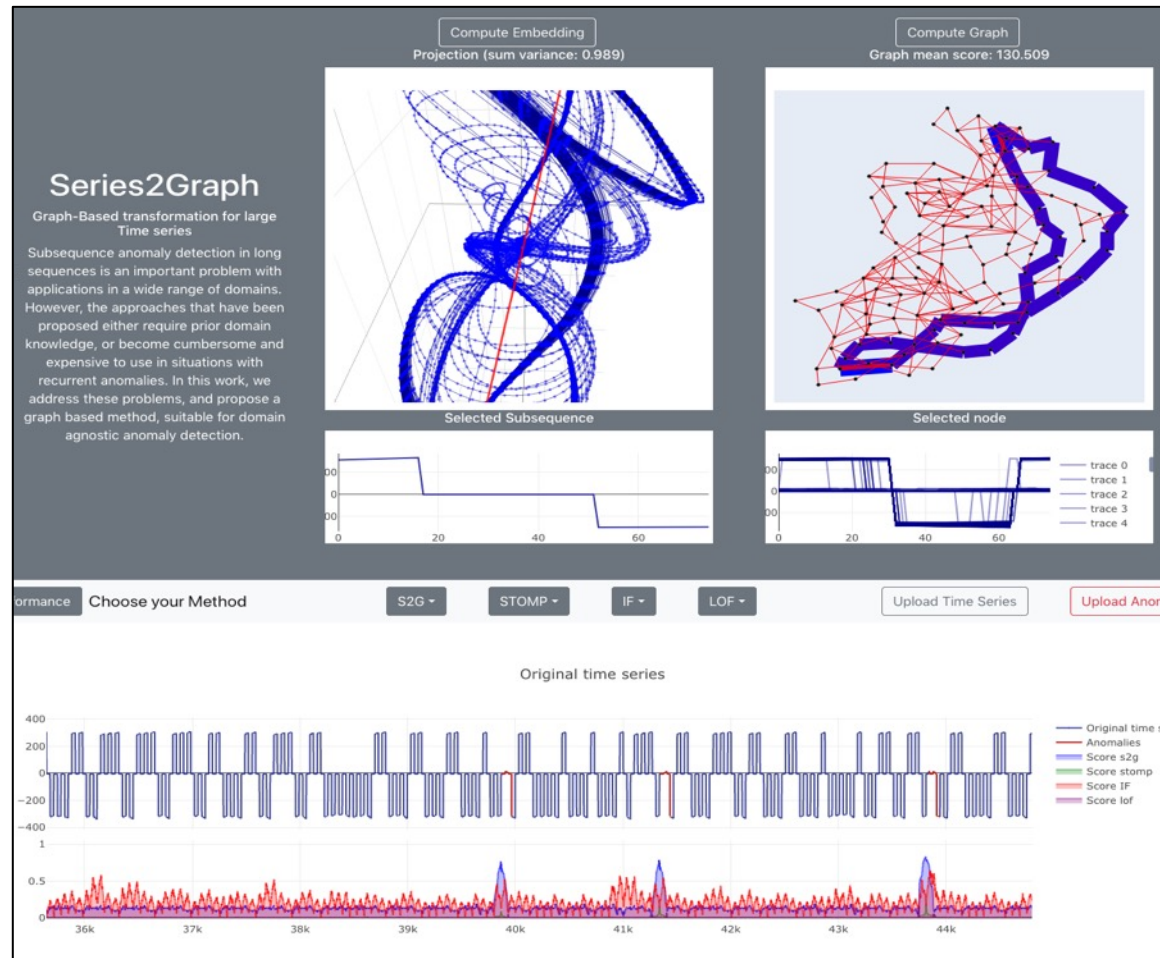
Density-based approach that **convert** the time series into a **graph** and detect **unusual trajectories**

Unsupervised

Univariate

subsequence

Anomaly Detection methods: *an Example*



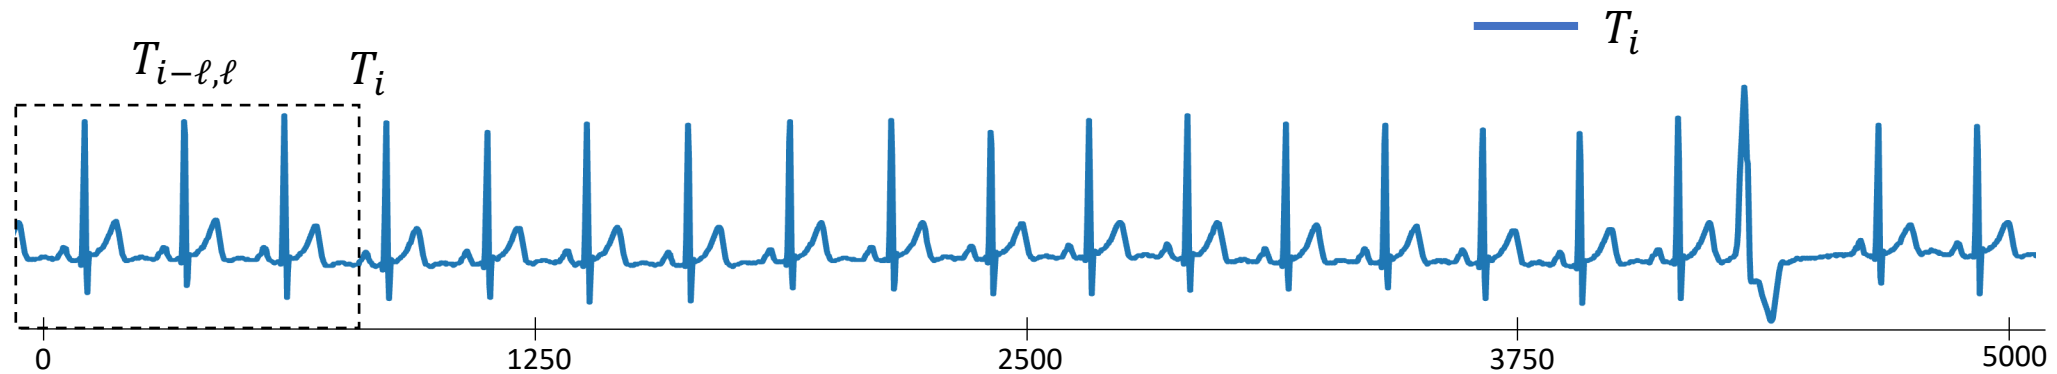
GraphAn [28]

An interactive tool to dive into the computation steps of Series2Graph :



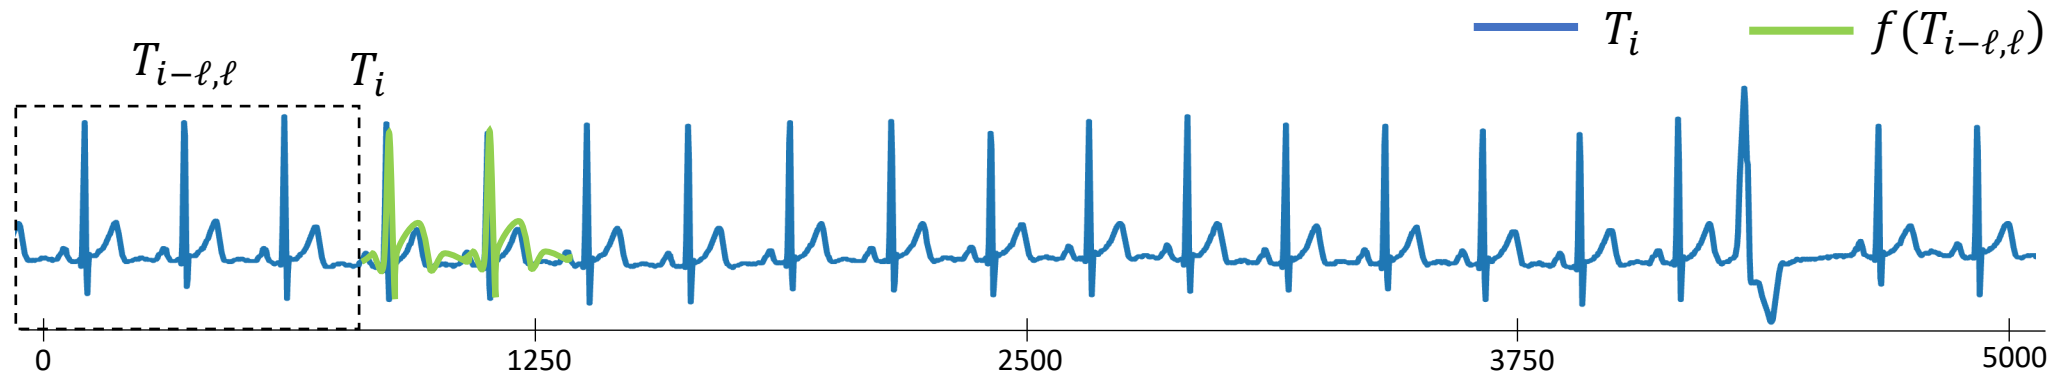
Anomaly Detection methods: *Forecasting-based*

Methods that aims to **predict the next points** based on the previous ones. The **prediction error** is used to detect if there is an anomaly or not.



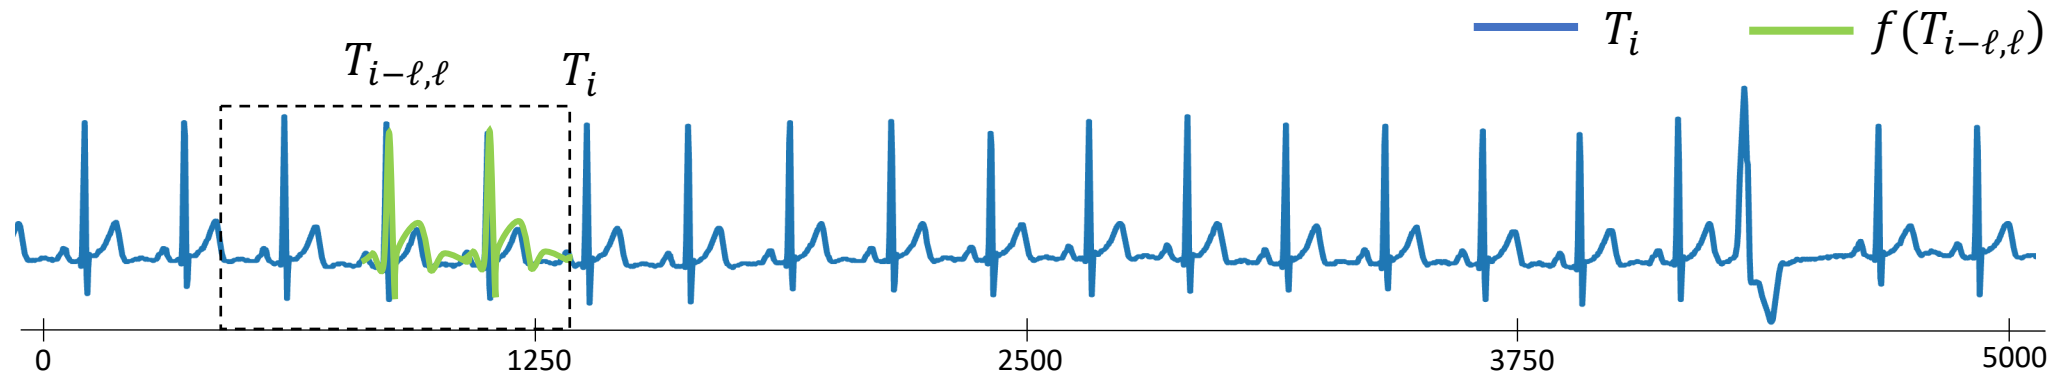
Anomaly Detection methods: *Forecasting-based*

Methods that aims to **predict the next points** based on the previous ones. The **prediction error** is used to detect if there is an anomaly or not.



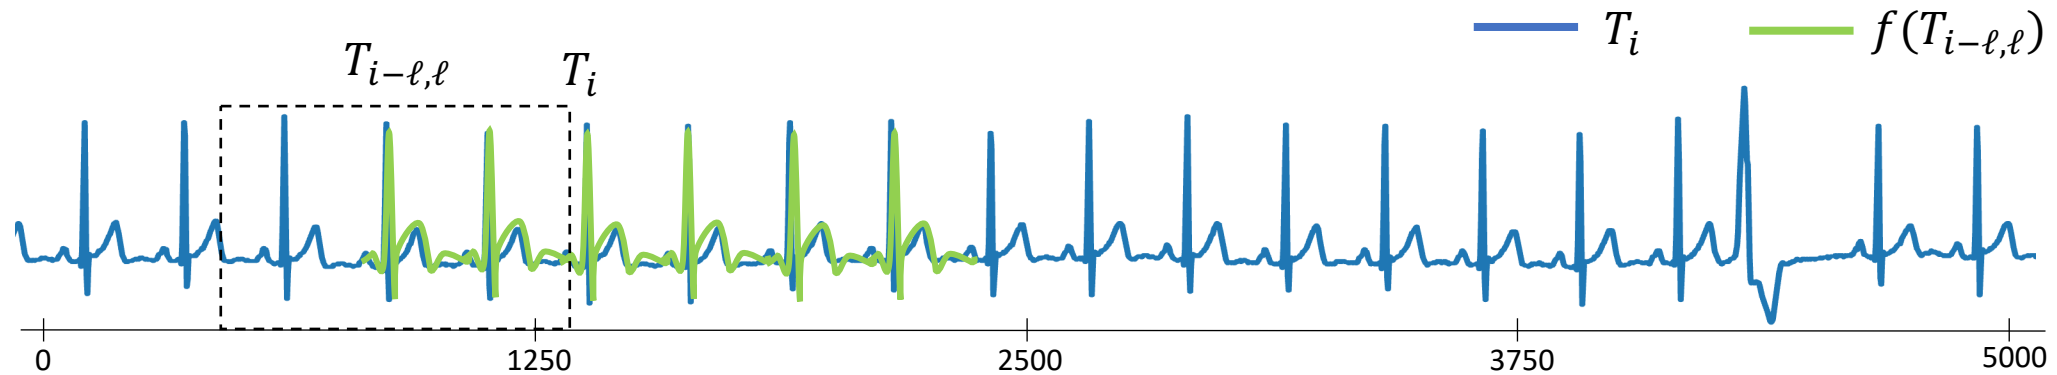
Anomaly Detection methods: *Forecasting-based*

Methods that aims to **predict the next points** based on the previous ones. The **prediction error** is used to detect if there is an anomaly or not.



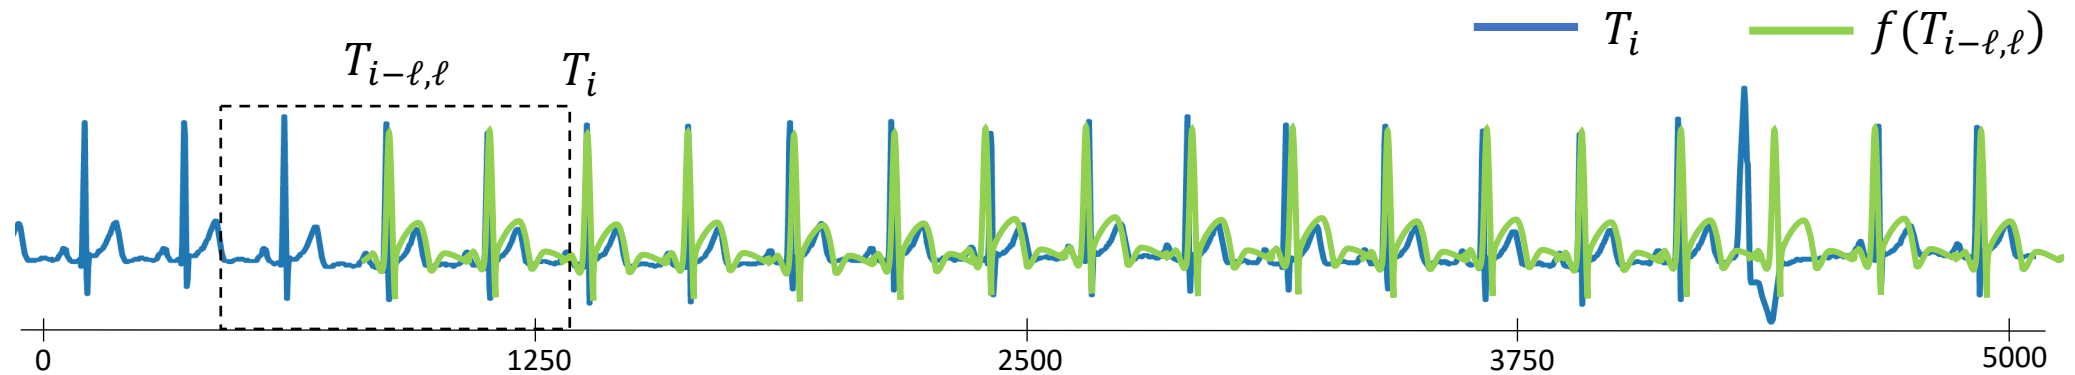
Anomaly Detection methods: *Forecasting-based*

Methods that aims to **predict the next points** based on the previous ones. The **prediction error** is used to detect if there is an anomaly or not.



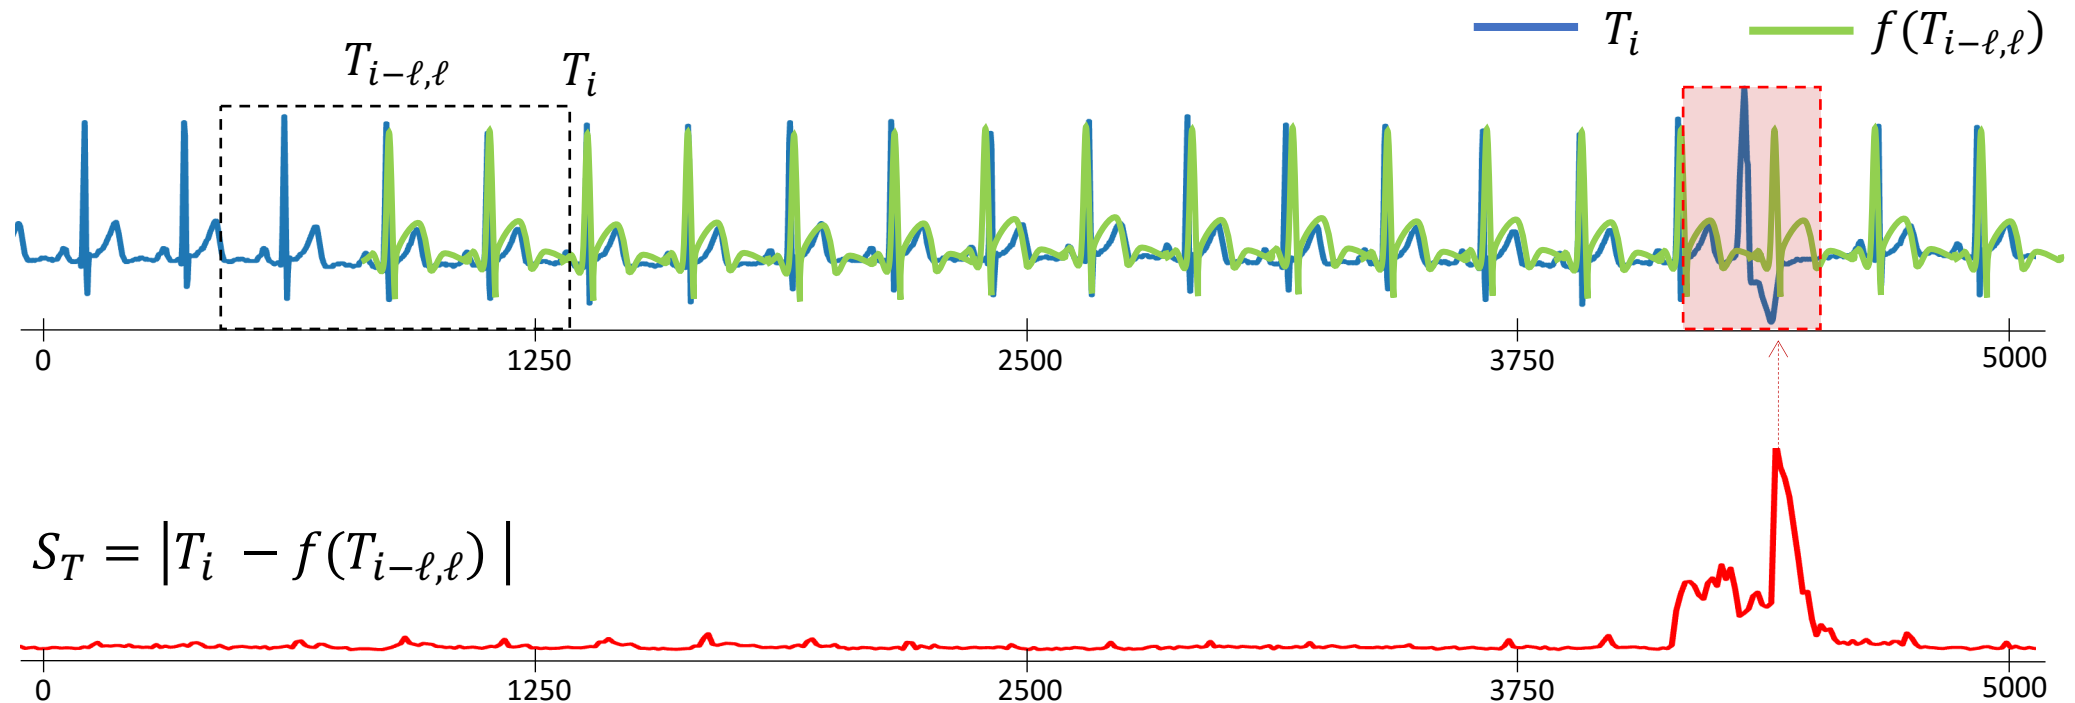
Anomaly Detection methods: *Forecasting-based*

Methods that aims to **predict the next points** based on the previous ones. The **prediction error** is used to detect if there is an anomaly or not.

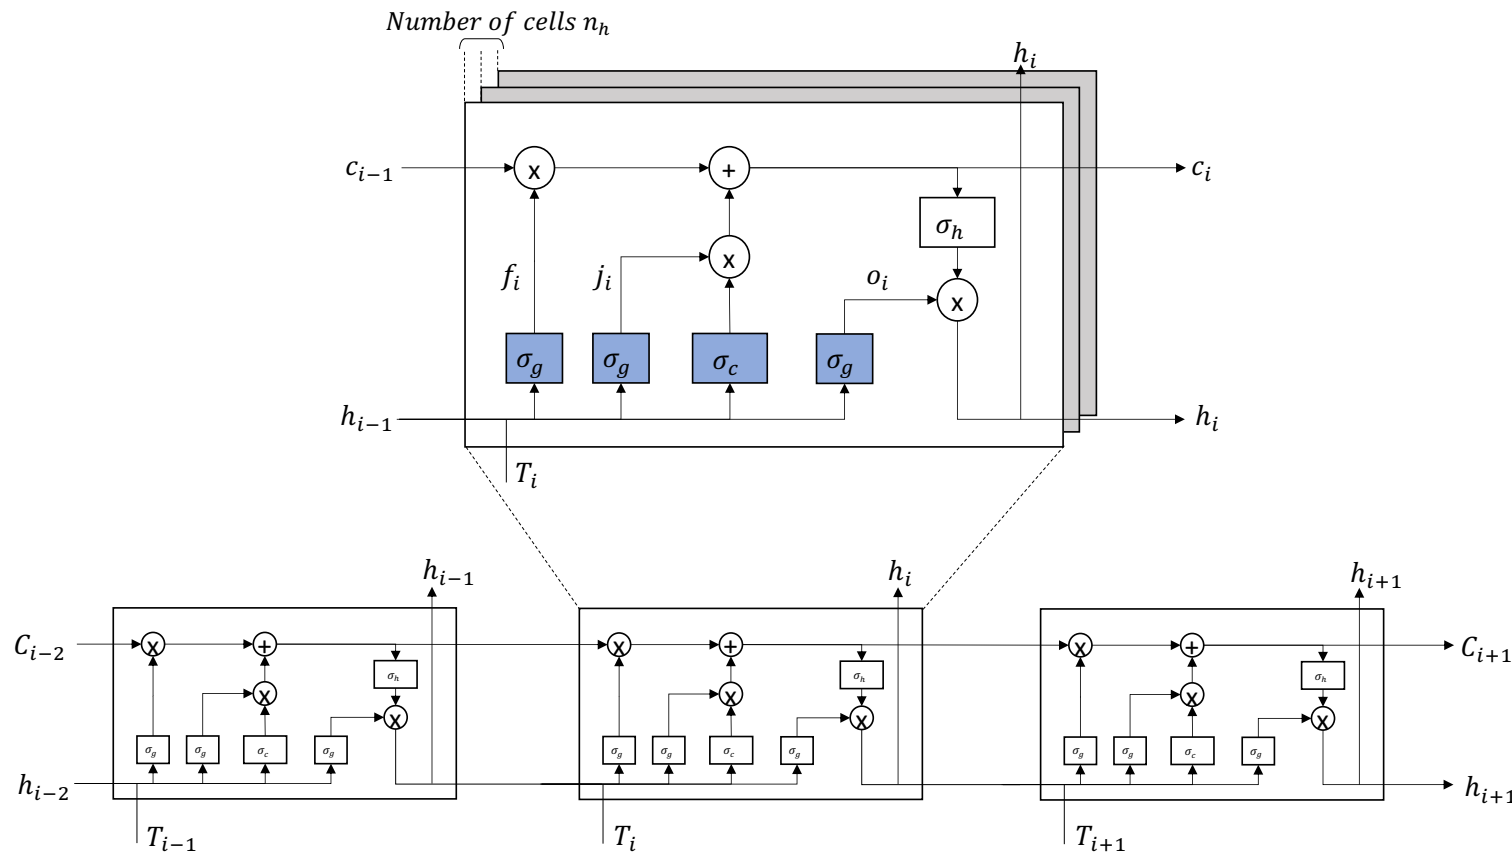


Anomaly Detection methods: *Forecasting-based*

Methods that aims to **predict the next points** based on the previous ones. The **prediction error** is used to detect if there is an anomaly or not.



Anomaly Detection methods: *an Example*



LSTM-AD [15]

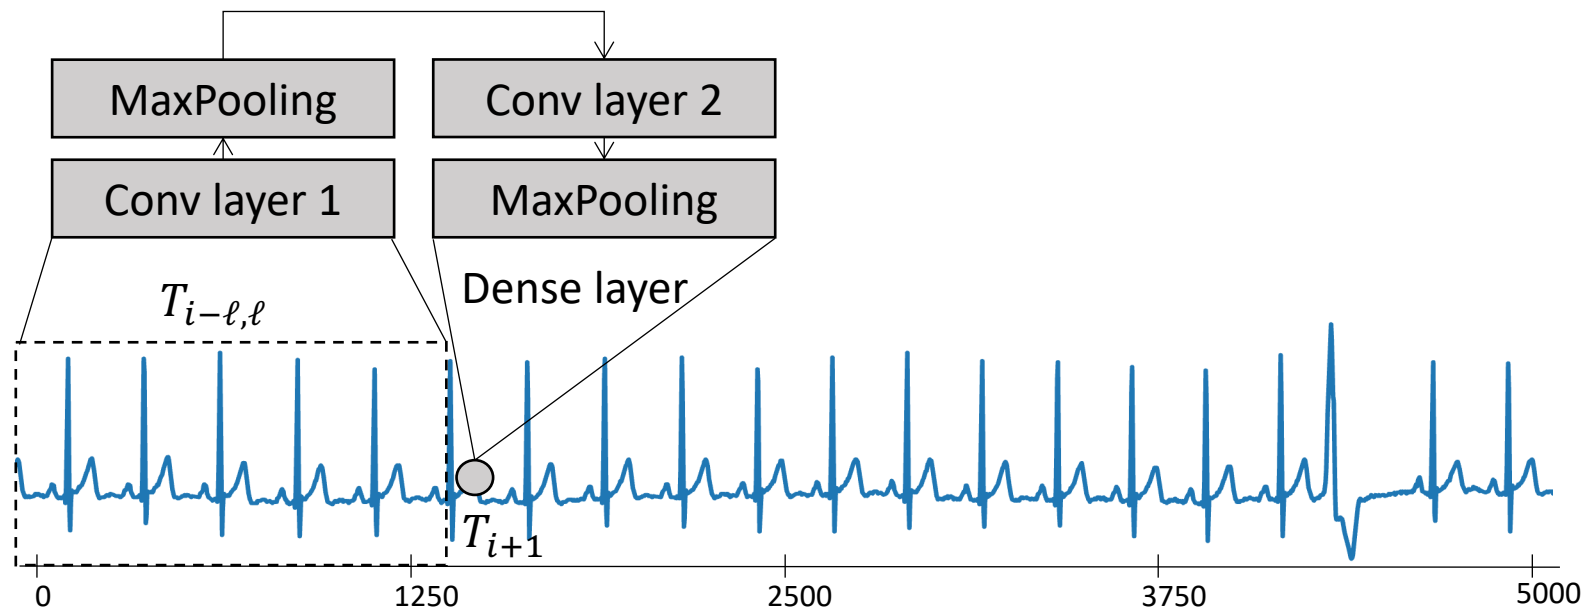
Model that stack multiple LSTM cell and use the output to predict the next value

Semi-supervised

Univariate/Multivariate

Point/sequence

Anomaly Detection methods: *an Example*



DeepAnT [16] (CNN)

Convolutional-based approach (2 convolutional layers) taking as input a sequence and aims to predict the next value.

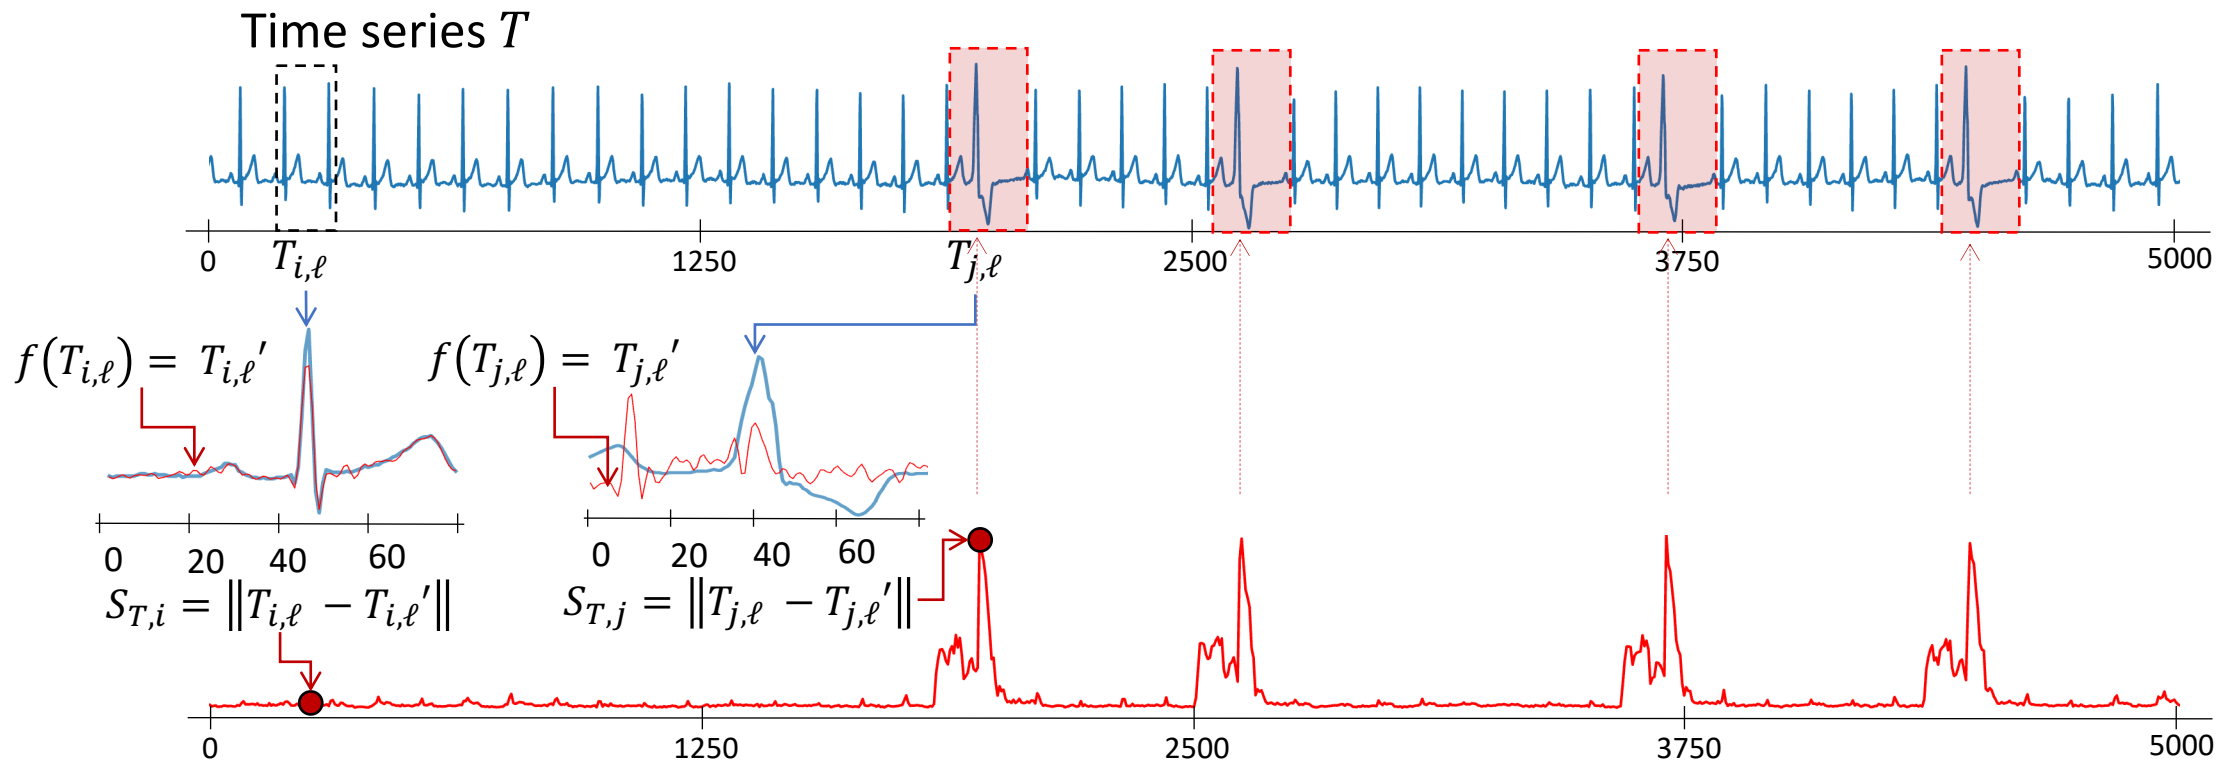
Semi-supervised

Univariate/Multivariate

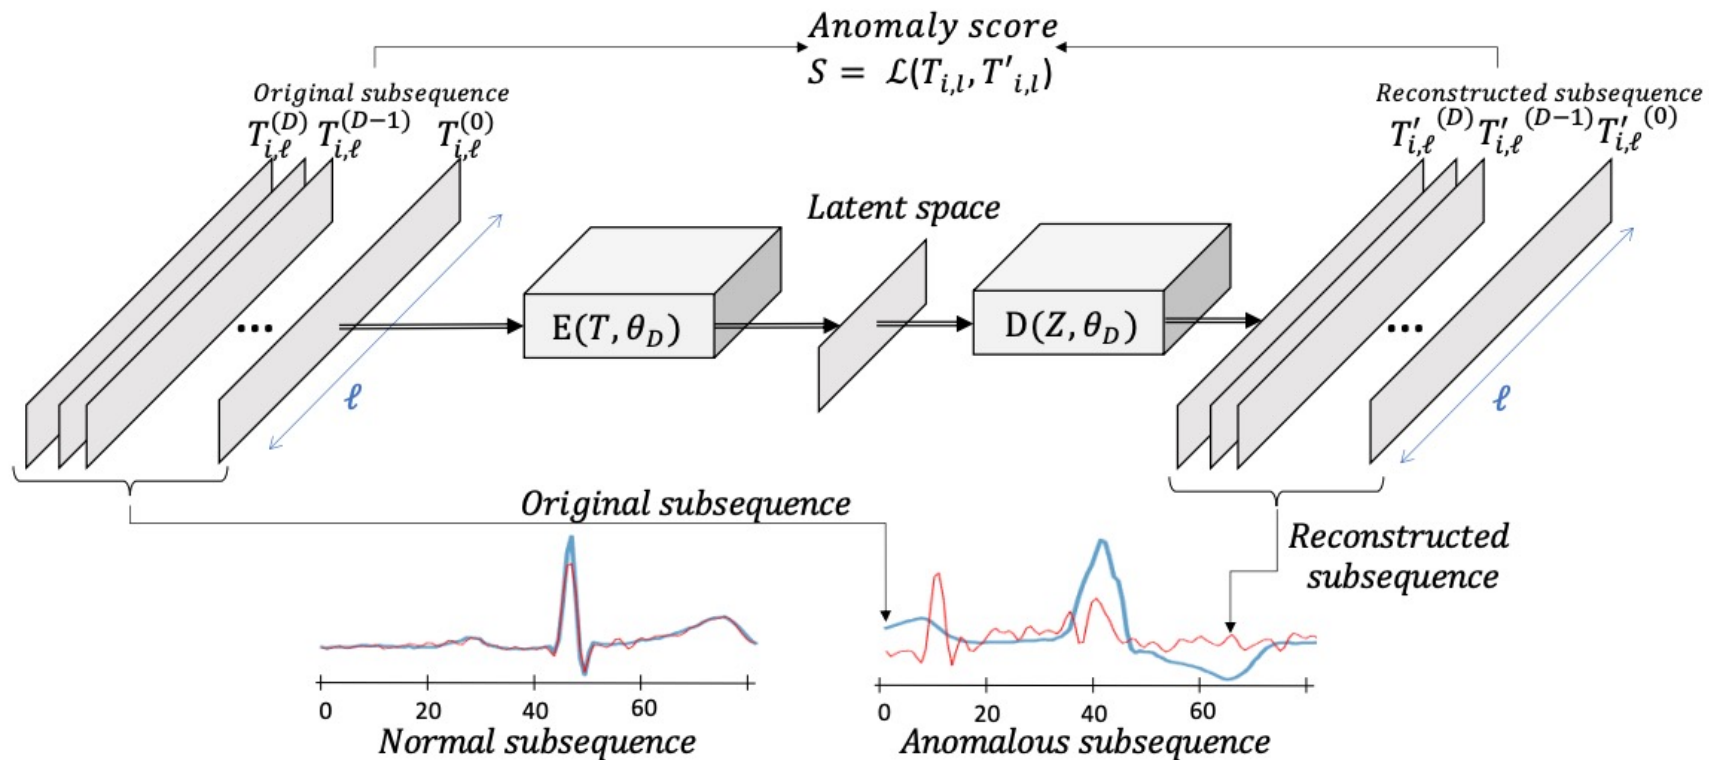
Point/sequence

Anomaly Detection methods: *Reconstruction-based*

Methods that aims to **reconstruct** the time series T and use the **reconstruction error** to detect if the time series is an anomaly or not.



Anomaly Detection methods: *an Example*



AutoEncoders [17] (AE)

Neural Network composed of an **encoder** (that reduce the dimensionality) and **decoder** that **reconstruct** the time series. The objective is to **minimize the reconstruction error**.

Semi-supervised

Univariate/Multivariate

Point/sequence

Anomaly Detection methods: *Existing benchmark*

HEX/UCR [18]

Set of **250 time series** with labels.

Details

- The labels have been manually checked and are reliable
- Each time series contains only 1 labeled anomaly

TimeEval [5]

Set of **976 time series** with labels.

Details

- New synthetic benchmark GutenTag used to tune parameters
- Only Time series with low contamination rate (< 0.1)
- Time series with at least one methods above 0.8 AUC-ROC

TSB-UAD [19]

Set of **2000 time series** with labels.

Details

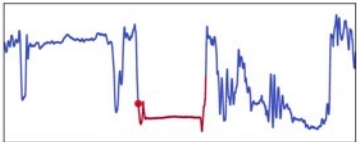
- Collected as proposed in the literature (no filtering based on contamination, size or label quality)
- Artificial and synthetic data generation methods for reliable labels

Anomaly Detection methods: *Existing benchmark*

HEX/UCR [18]

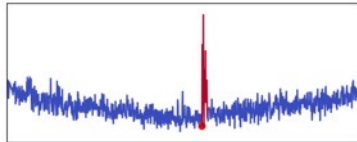
Set of 250 time series with

OPPORTUNITY

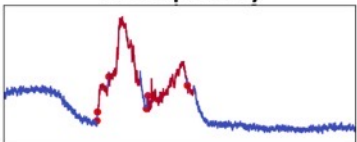


Occupancy

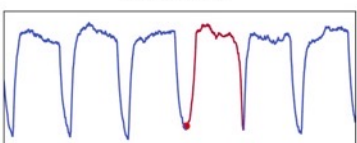
IOPS



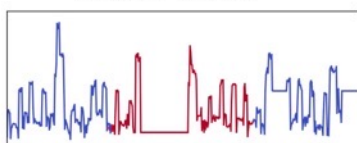
ECG



KDD21



NASA-SMAP

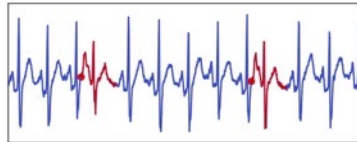


only 1 labeled anomaly

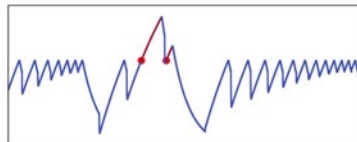
TimeEval [5]

Real datasets collection

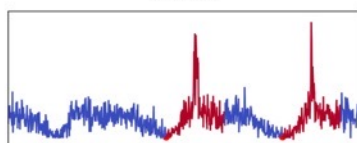
SVDB



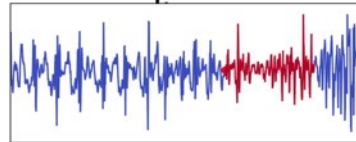
GHL



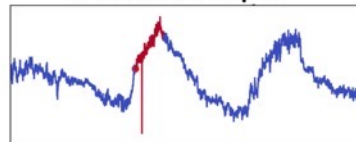
NAB



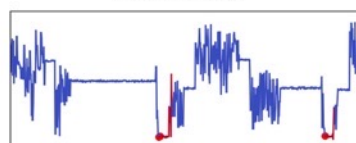
Daphnet



SensorScope



Genesis



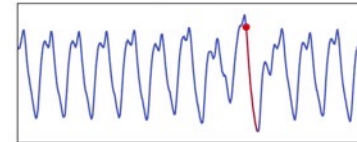
contamination rate (< 0.1)

- Time series with at least one methods above 0.8 AUC-ROC

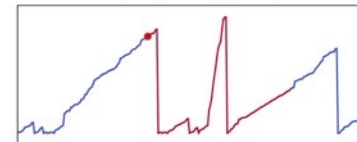
TSB-UAD [19]

Set of 2000 time series with

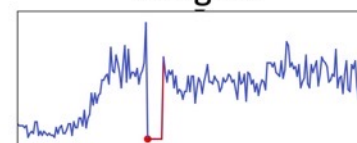
MGAB



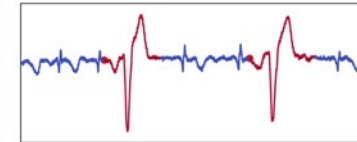
NASA-MSL



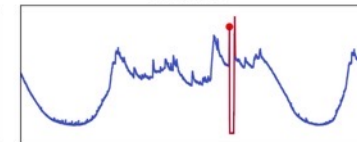
Dodgers



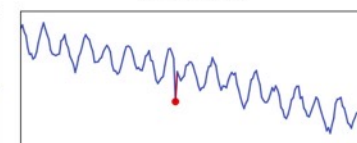
MITDB



SMD



YAHOO



contamination, size of label quality.

Anomaly Detection methods: *Existing benchmark*

HEX/UCR [18]

TimeEval [5]

TSB-UAD [19]

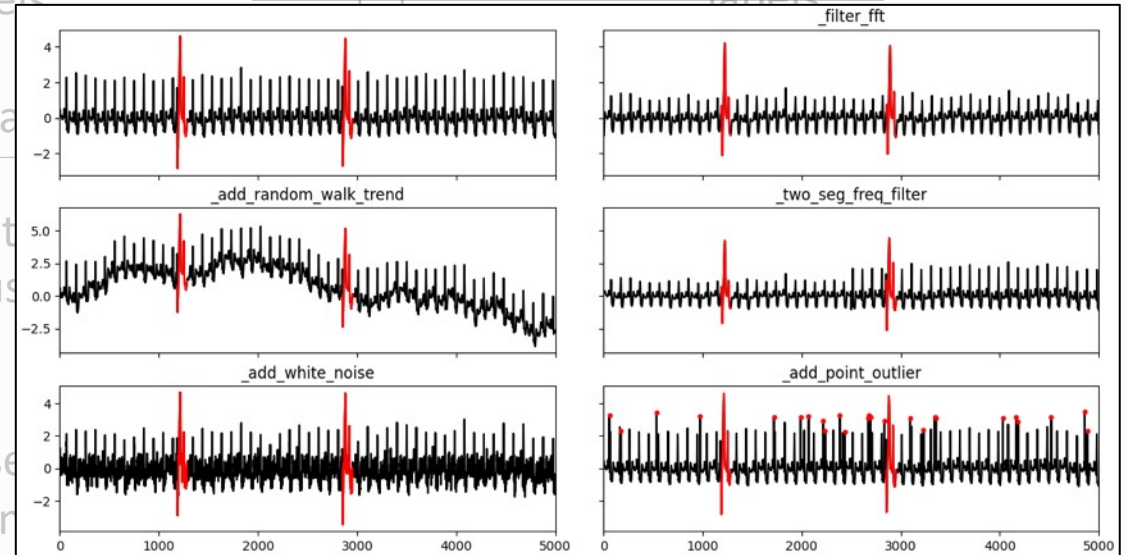
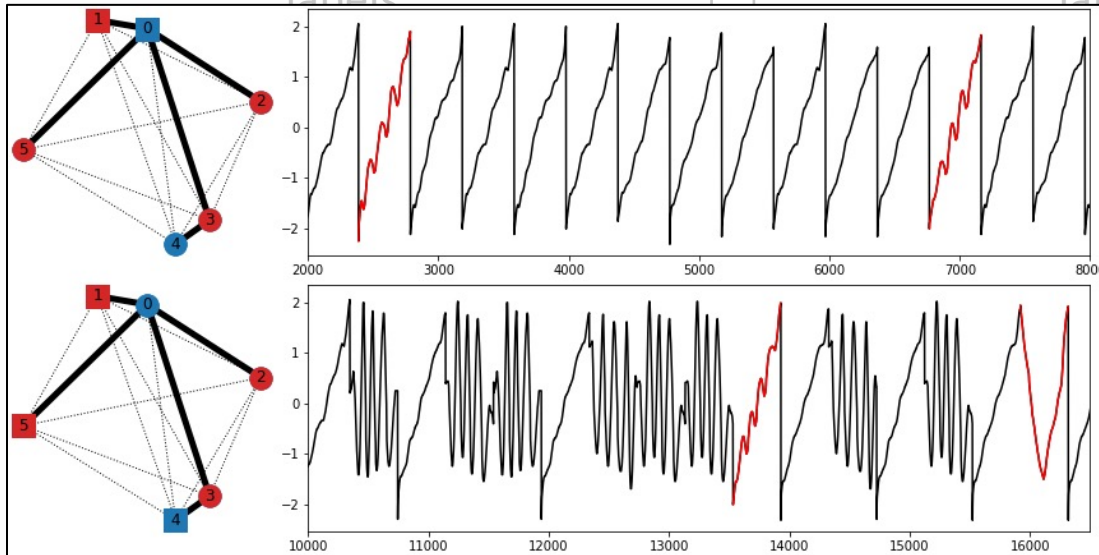
Set of 250

Artificial dataset generation

of 976 time series with labels

Synthetic dataset generation

ies with



- Time series with at least one methods above 0.8 AUC-ROC

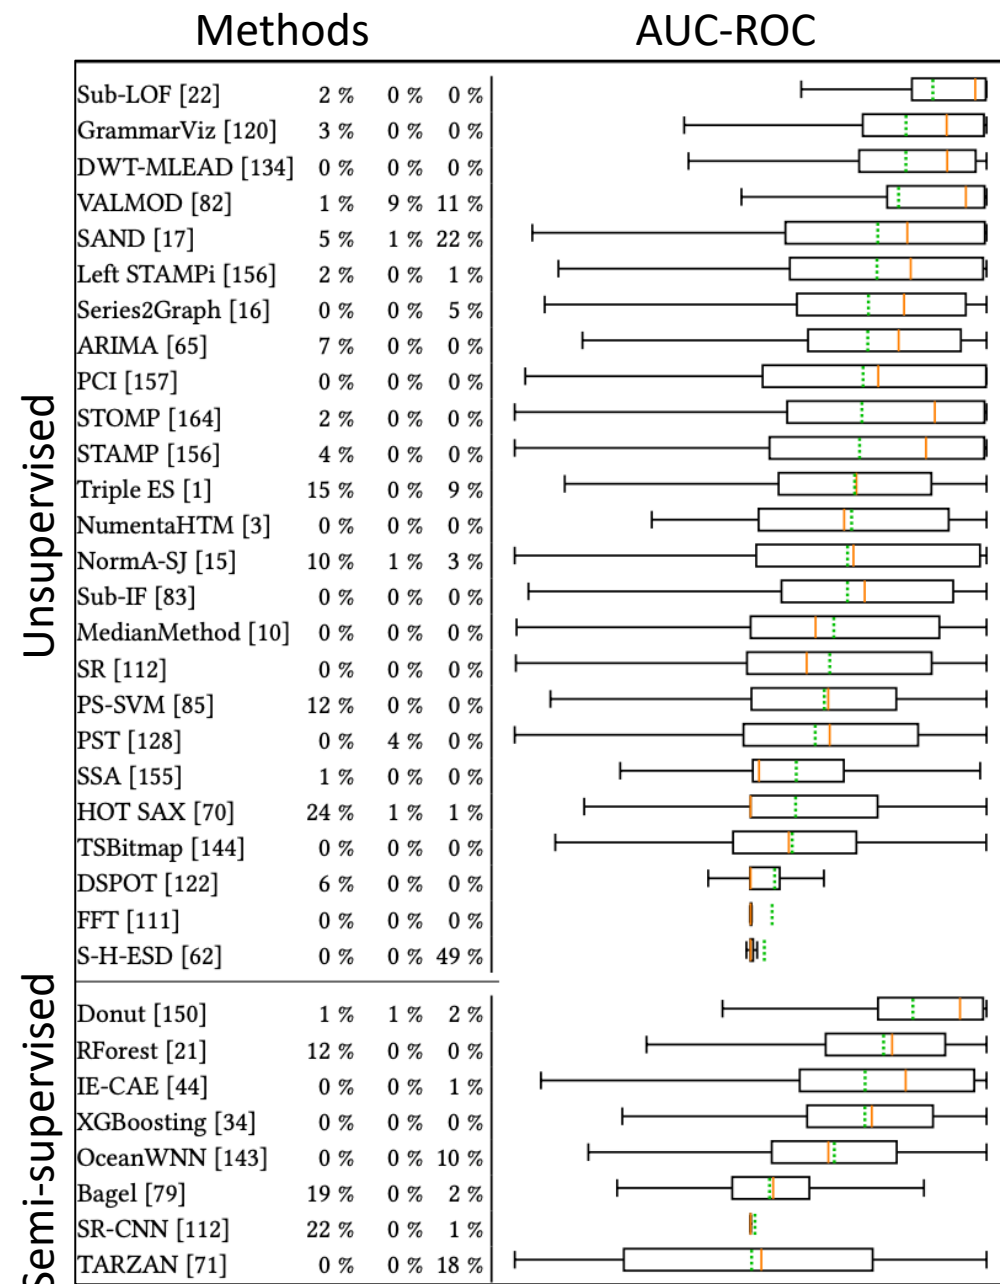
quality.

Anomaly Detection methods: *Experimental evaluation*

Observations on TimeEval [5]:

- Distance-based and Density-based methods have a better accuracy (AUC-ROC) than forecasting and reconstruction-based approaches
- Semi-supervised methods are not outperforming Unsupervised approaches

[5] Sebastian Schmidl, Phillip Wenig, and Thorsten Papenbrock. 2022. Anomaly detection in time series: a comprehensive evaluation. Proc. VLDB Endow. 15, 9 (May 2022), 1779–1797.

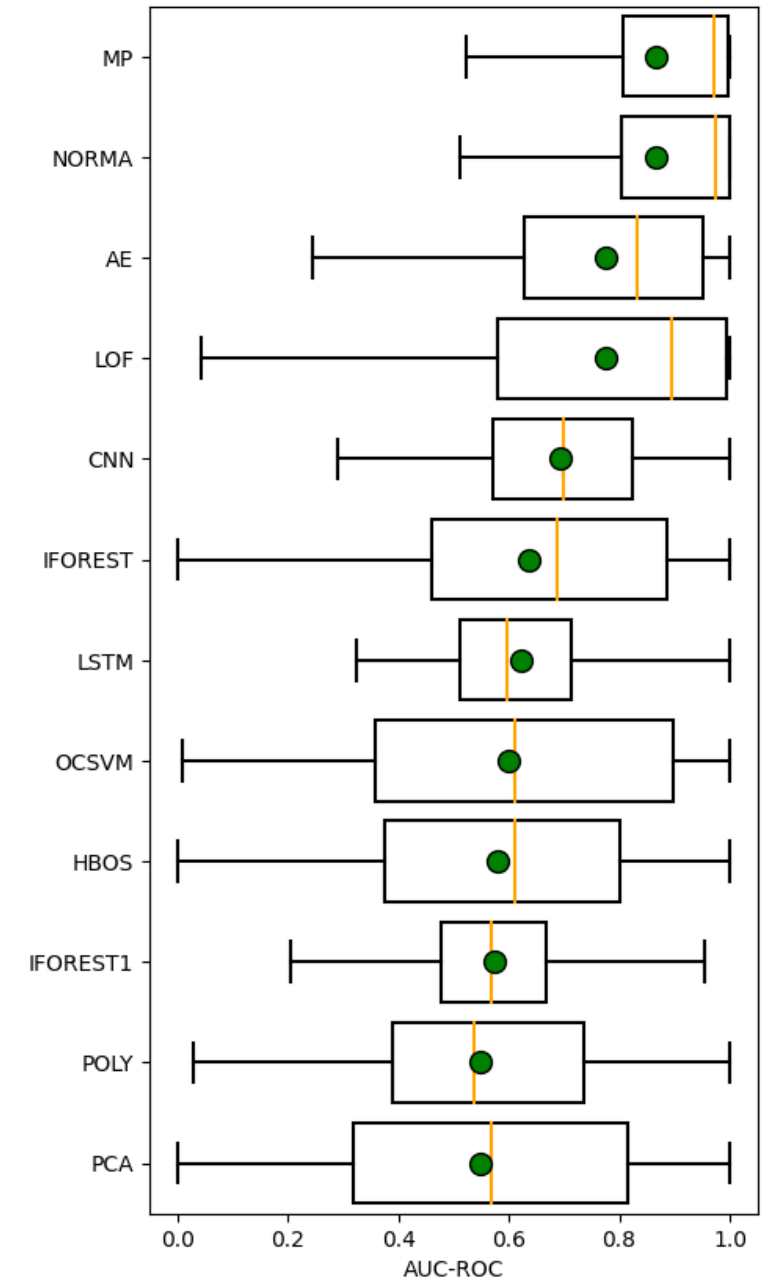


Anomaly Detection methods: *Experimental evaluation*

Observations on HEX/UCR [18]:

- Distance-based methods have a better accuracy (AUC-ROC) than forecasting and distribution-based approaches

[18] R. Wu and E. Keogh, "Current Time Series Anomaly Detection Benchmarks are Flawed and are Creating the Illusion of Progress" in IEEE Transactions on Knowledge & Data Engineering, vol. 35, no. 03, pp. 2421-2429, 2023.

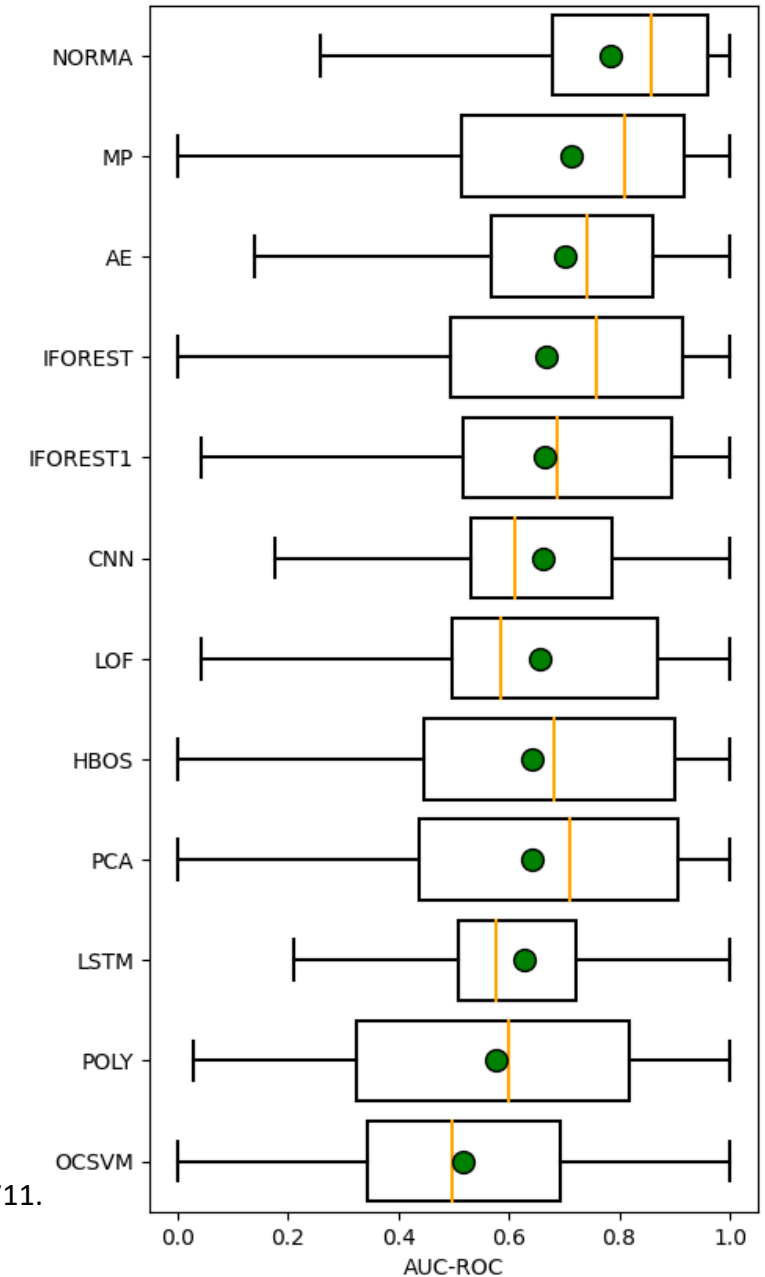


Anomaly Detection methods: *Experimental evaluation*

Observations on TSB-UAD [19]:

- Distance-based methods have a better accuracy (AUC-ROC) than forecasting-based methods.
- Isolation Forest (distribution-based and not proposed for time series) have also a strong accuracy
- AutoEncoder (AE) is also very accurate.

[19] 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. Proc. VLDB Endow. 15, 8 (April 2022), 1697–1711.



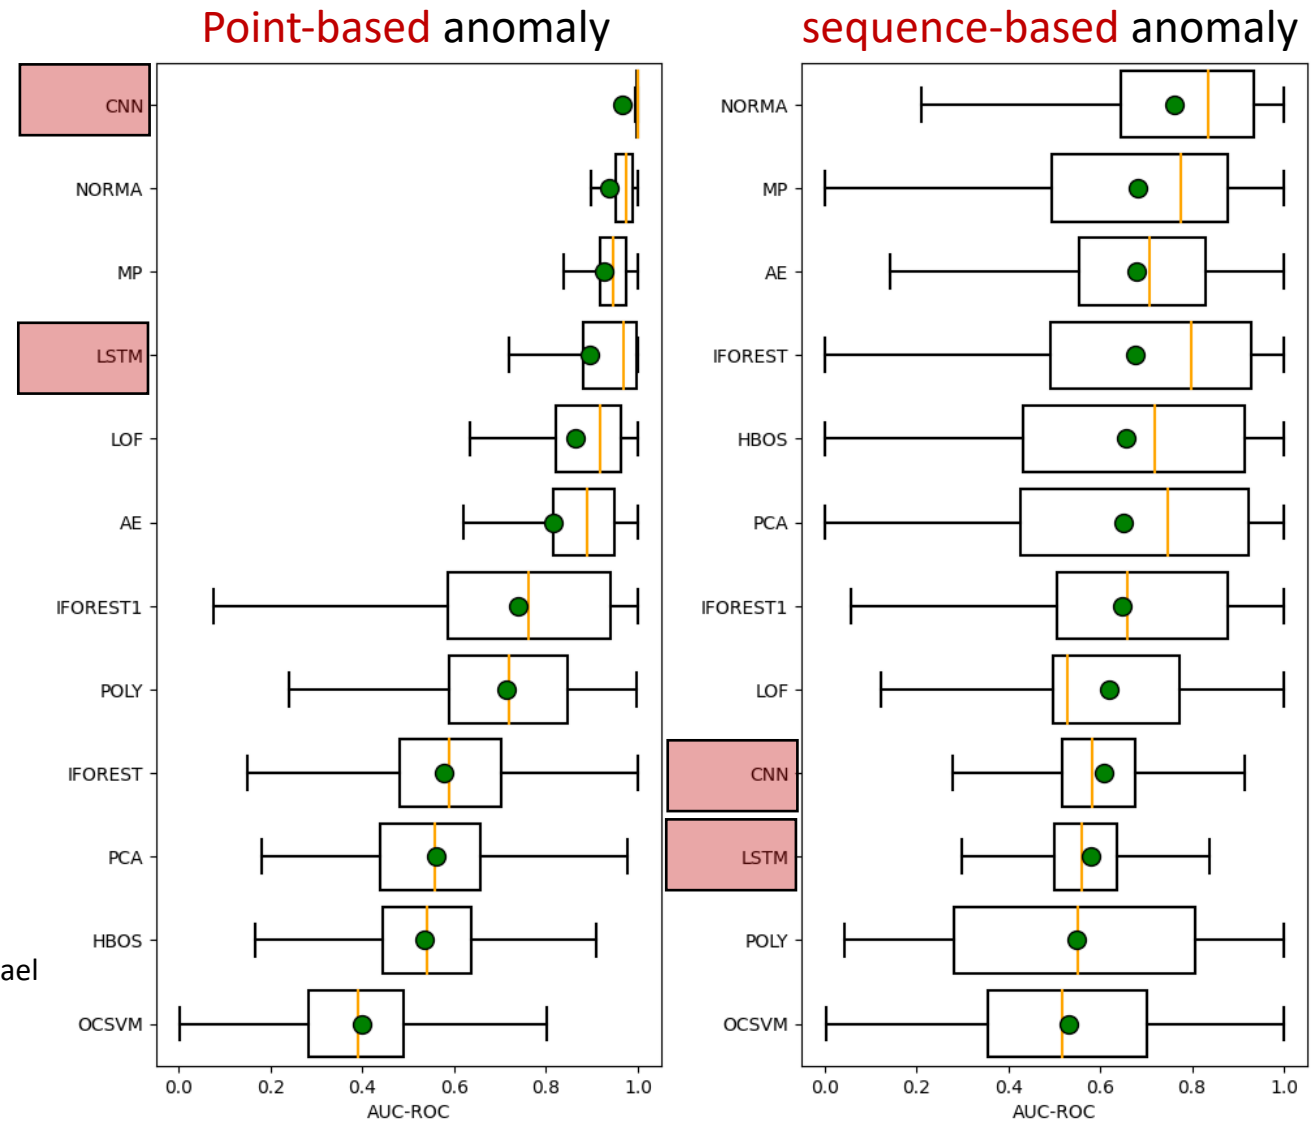
Anomaly Detection methods:

Experimental evaluation

Observations on TSB-UAD [19]:

- Forecasting methods (LSTM and CNN) are very **accurate** for point anomalies
- But have **poor performances** on **sequence-based** anomalies.

[19] 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. Proc. VLDB Endow. 15, 8 (April 2022), 1697–1711.

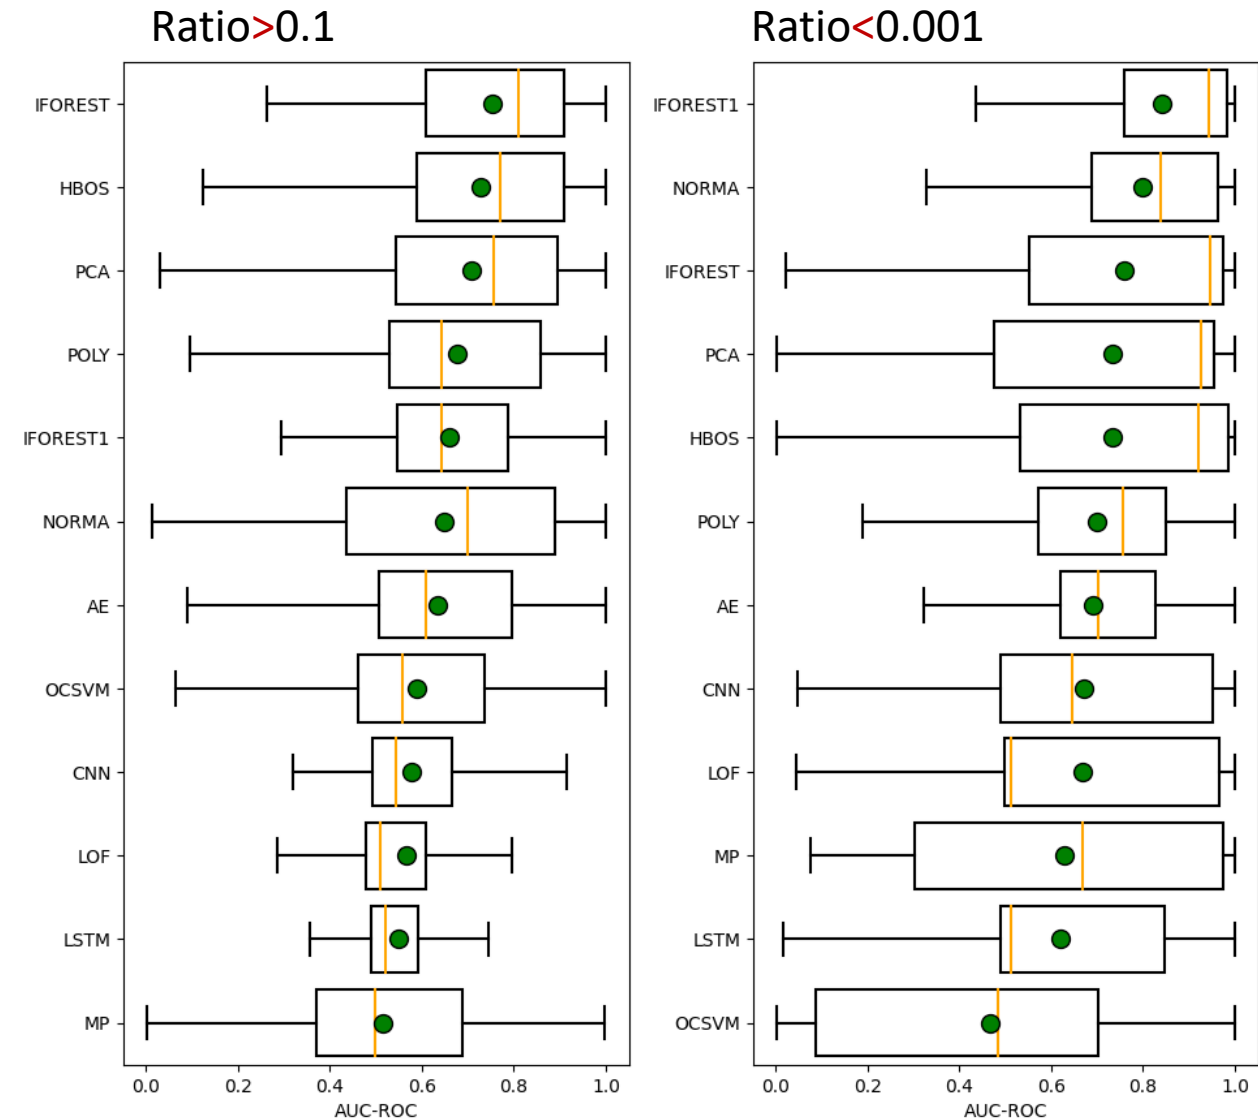


Anomaly Detection methods: *Experimental evaluation*

Observations on TSB-UAD [19]:

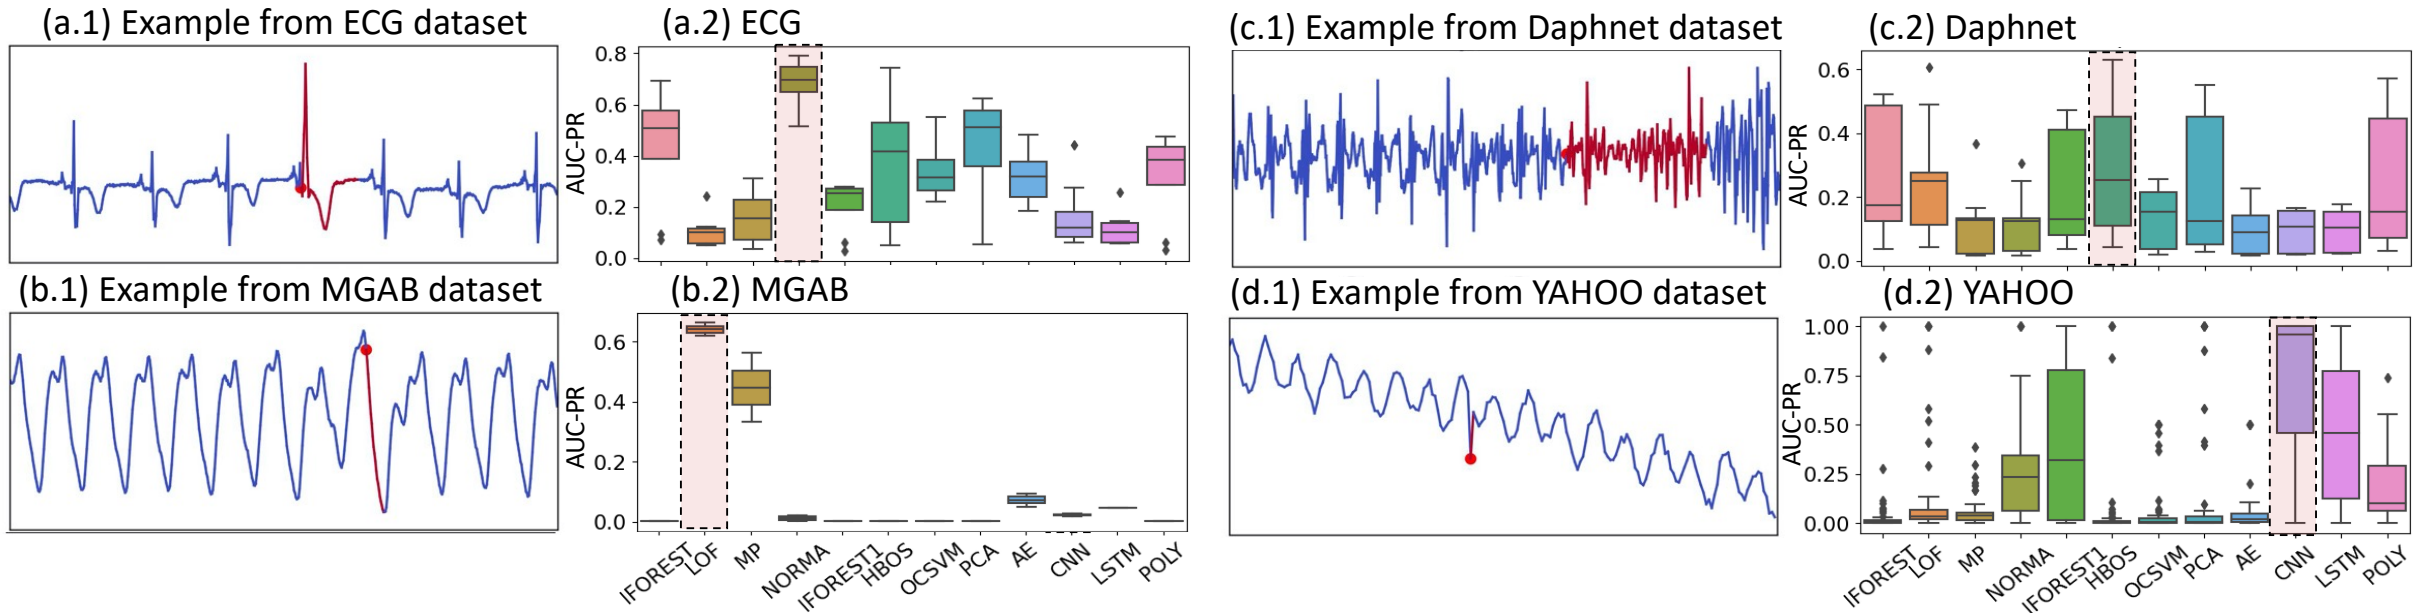
- The ratio of normal/abnormal points has a **strong impact** on the methods ranking.

[19] 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. Proc. VLDB Endow. 15, 8 (April 2022), 1697–1711.



Anomaly Detection methods: *Experimental evaluation*

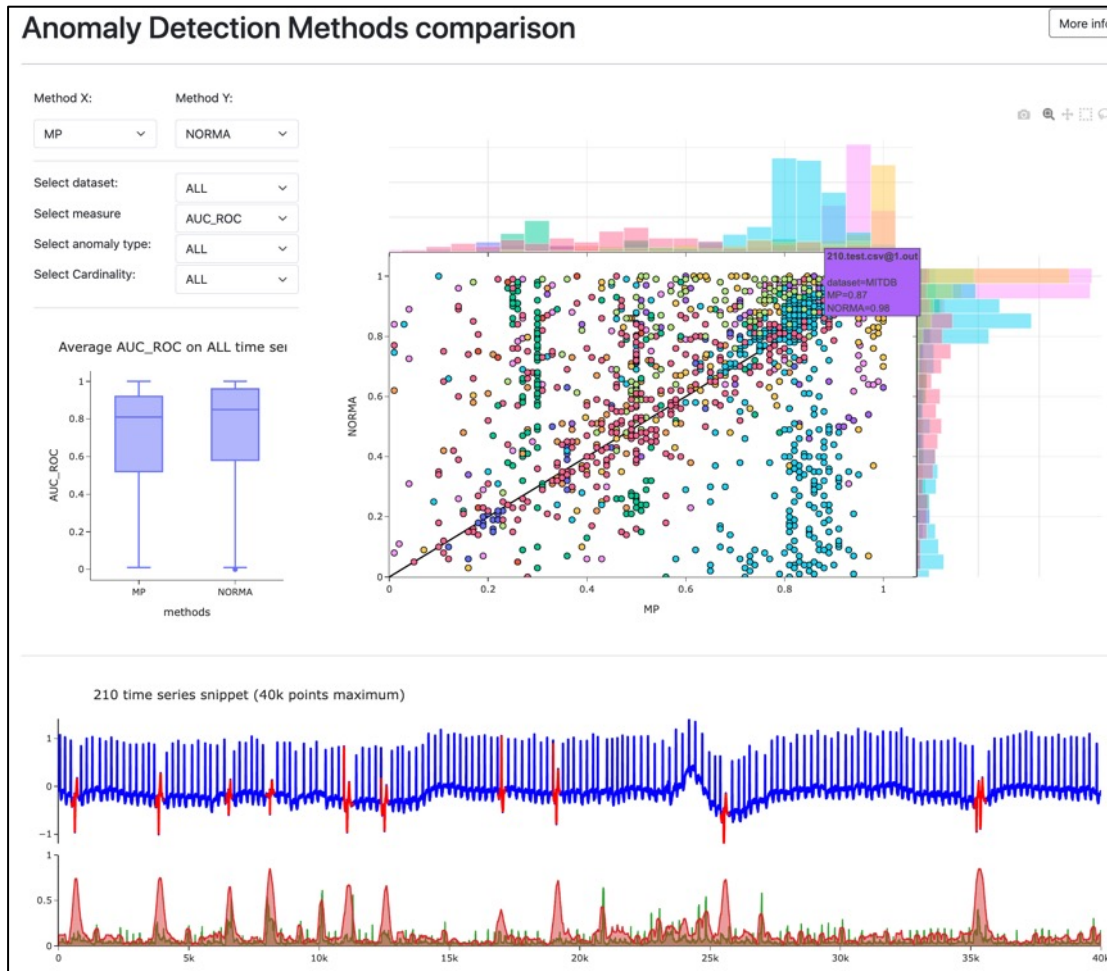
Observation from the results applied on specific datasets (TSB-UAD [19])



There is **no overall winner**.

[19] 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. Proc. VLDB Endow. 15, 8 (April 2022), 1697–1711.

Anomaly Detection methods: *Experimental evaluation*



Theseus [27]

An interactive tool to compare anomaly detection methods



VLDB 2022



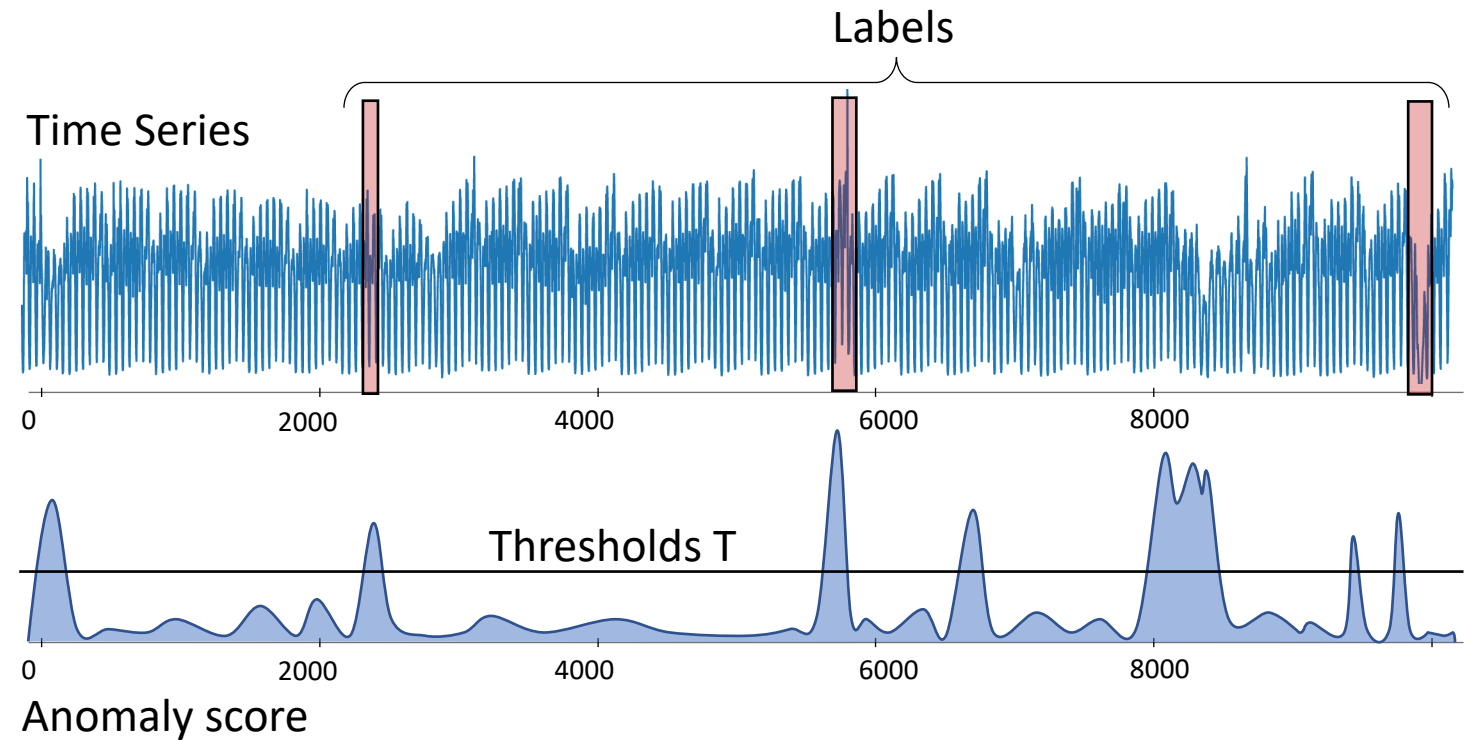
Github repo



Evaluation Measures

Evaluation measures: *Threshold-based*

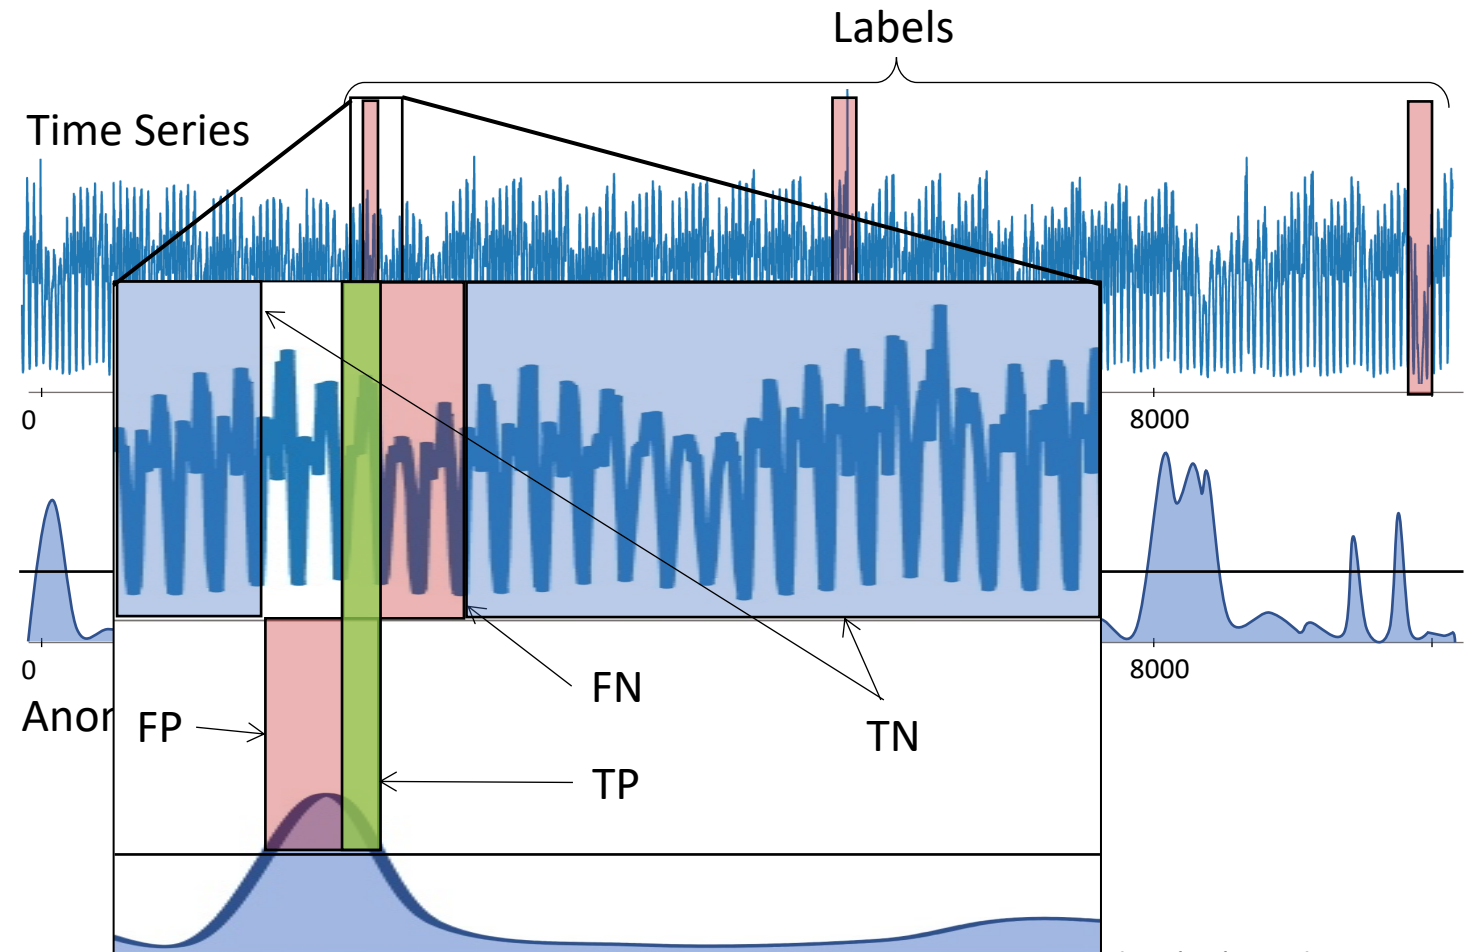
Threshold-based Evaluation Measures:



Evaluation measures: *Threshold-based*

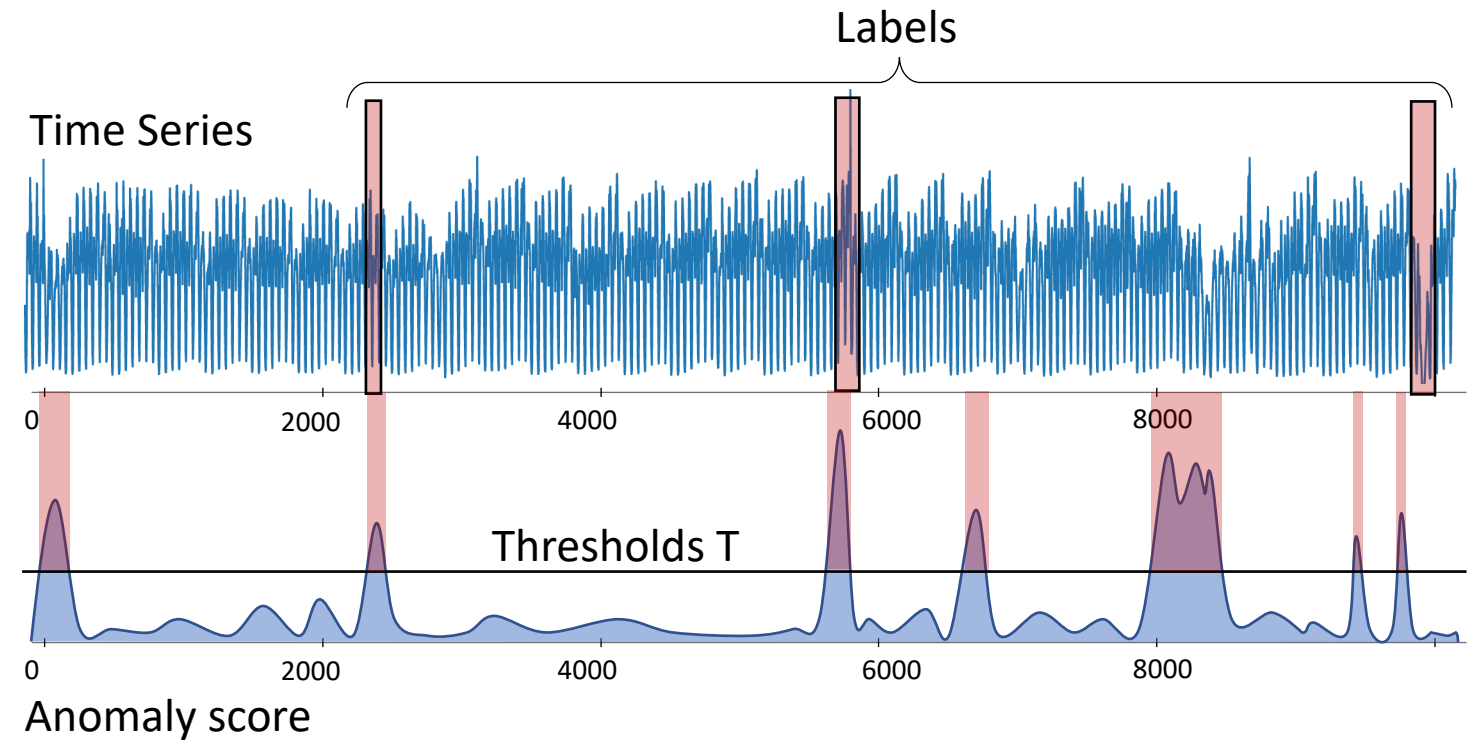
Threshold-based Evaluation Measures:

- Precision: $\frac{TP}{TP+FP}$
- Recall (true positive rate): $\frac{TP}{TP+FN}$
- False positive rate: $\frac{FP}{FP+TN}$
- F-score: $\frac{(1+\beta^2)*Precision}{\beta^2*Precision+Recall}$



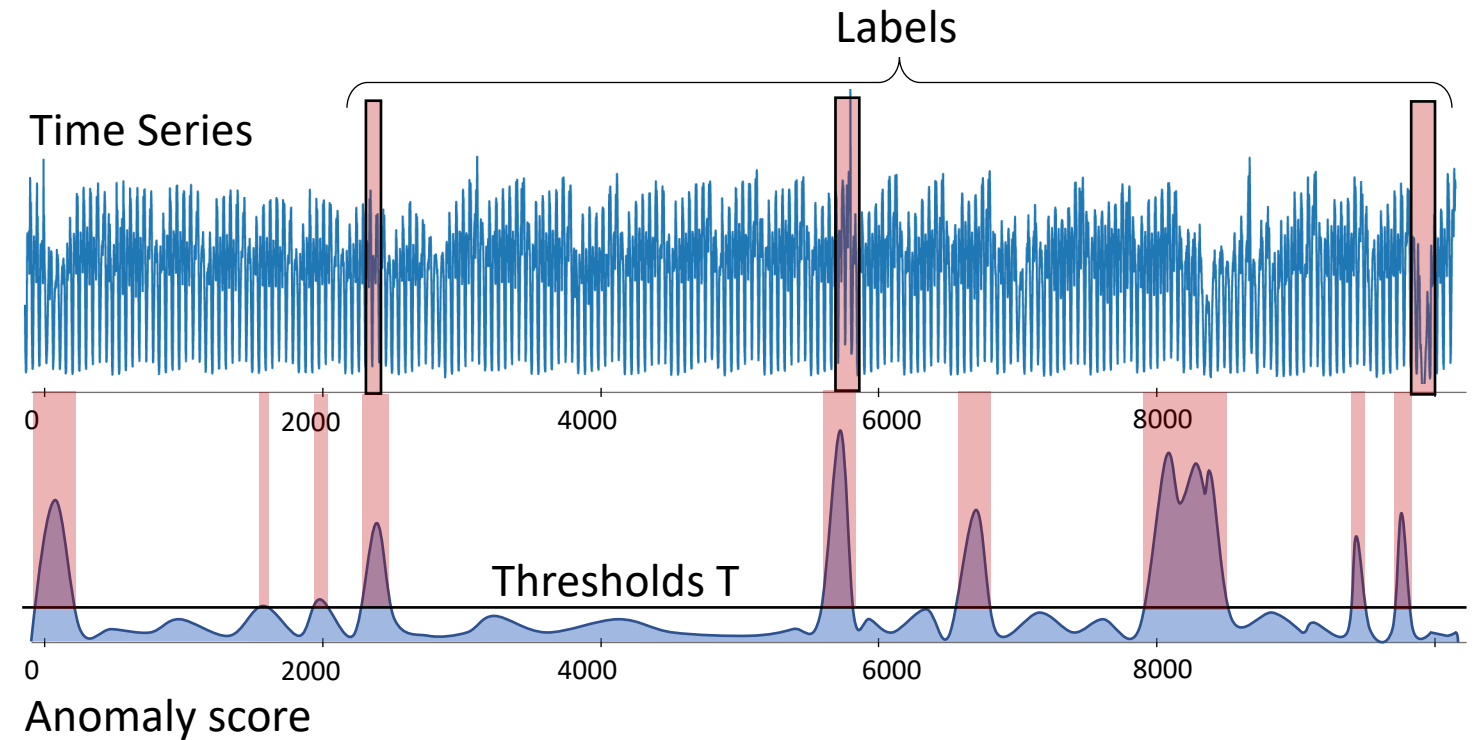
Evaluation measures: *AUC-based*

How do we set the threshold?



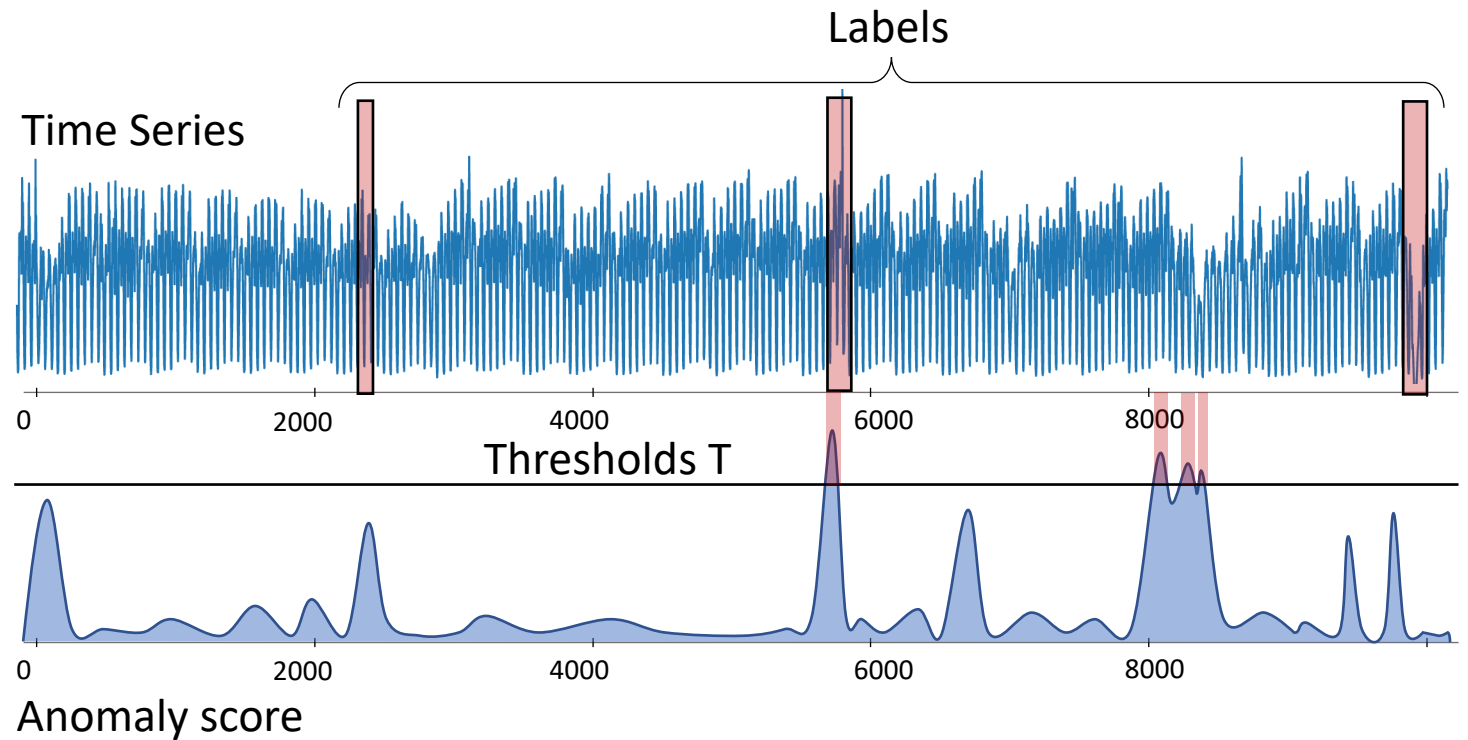
Evaluation measures: *AUC-based*

How do we set the threshold?



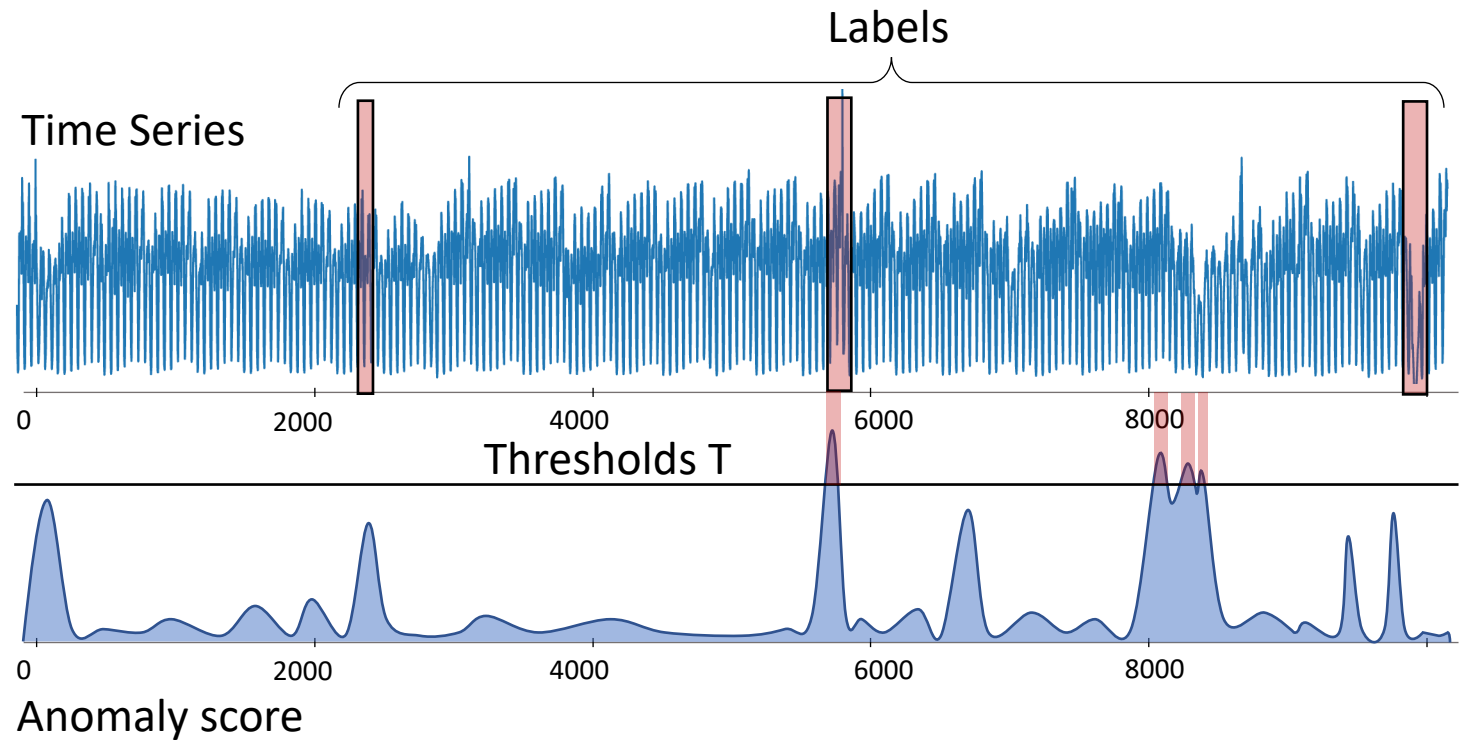
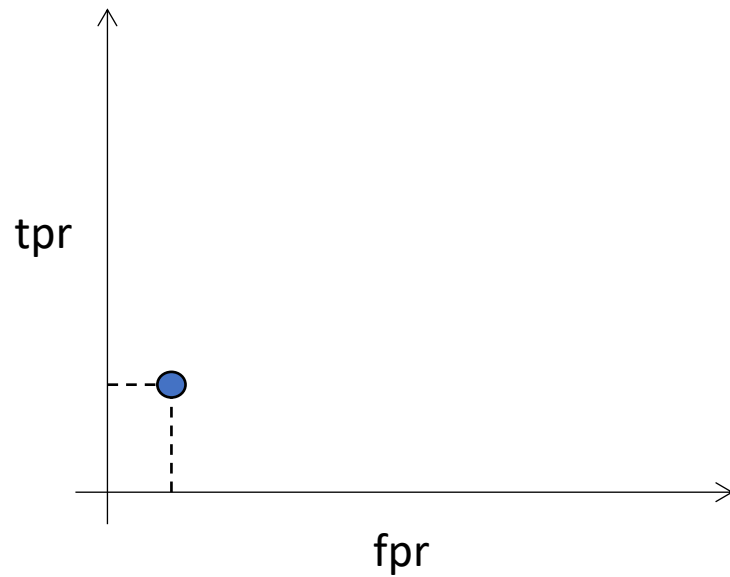
Evaluation measures: *AUC-based*

How do we set the threshold?



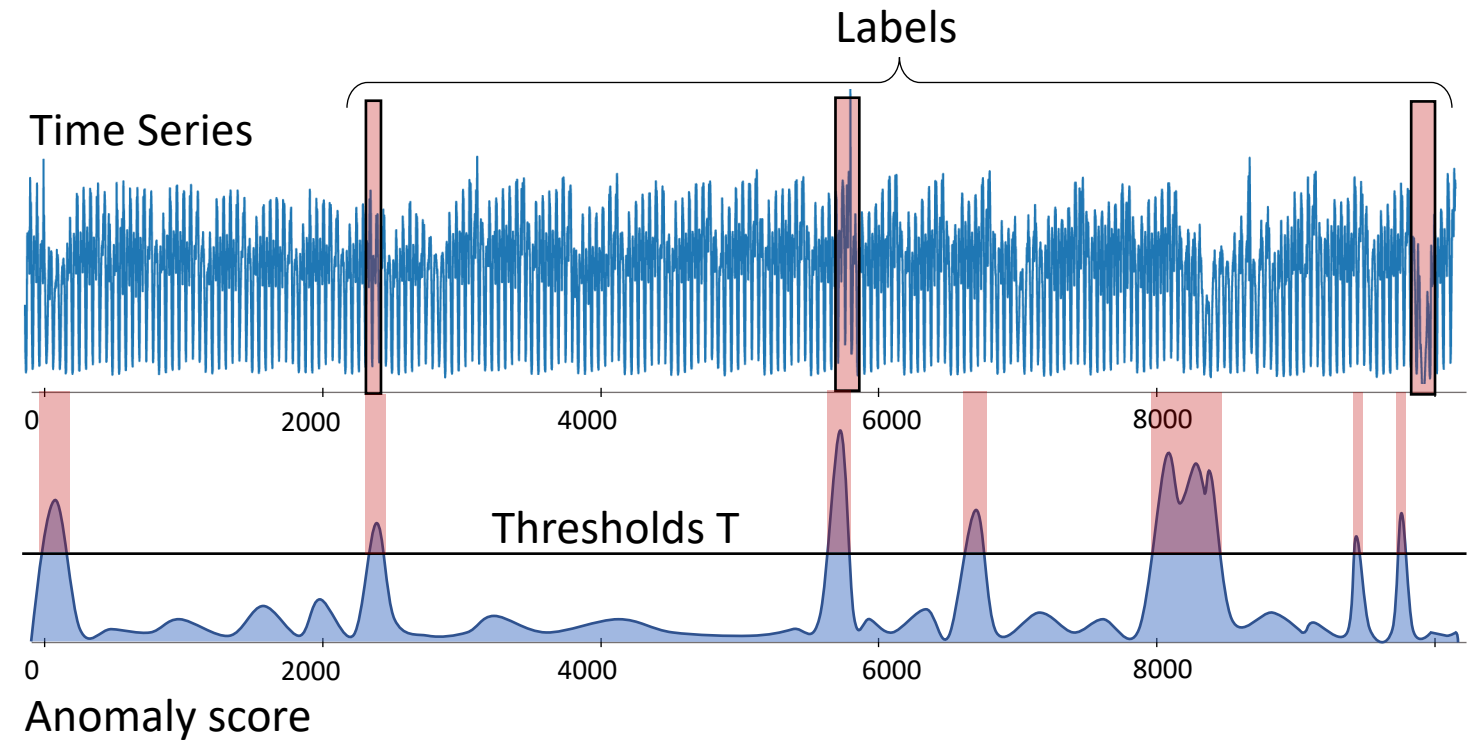
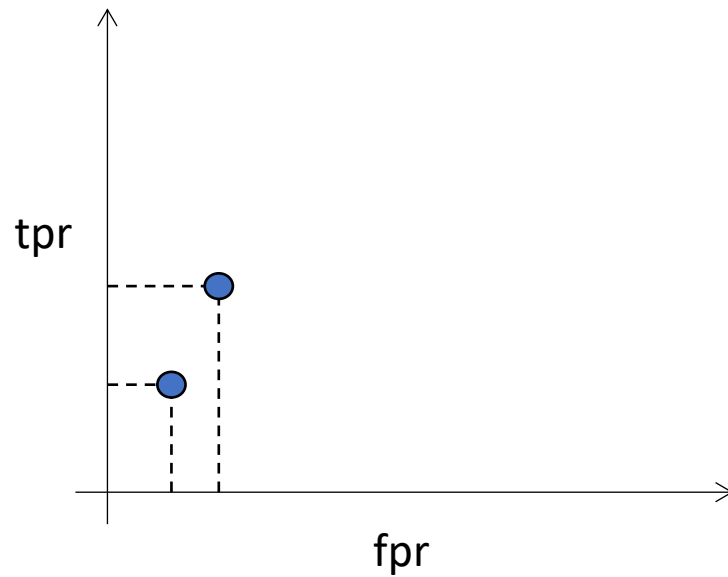
Evaluation measures: *AUC-based*

AUC-based Evaluation Measures:



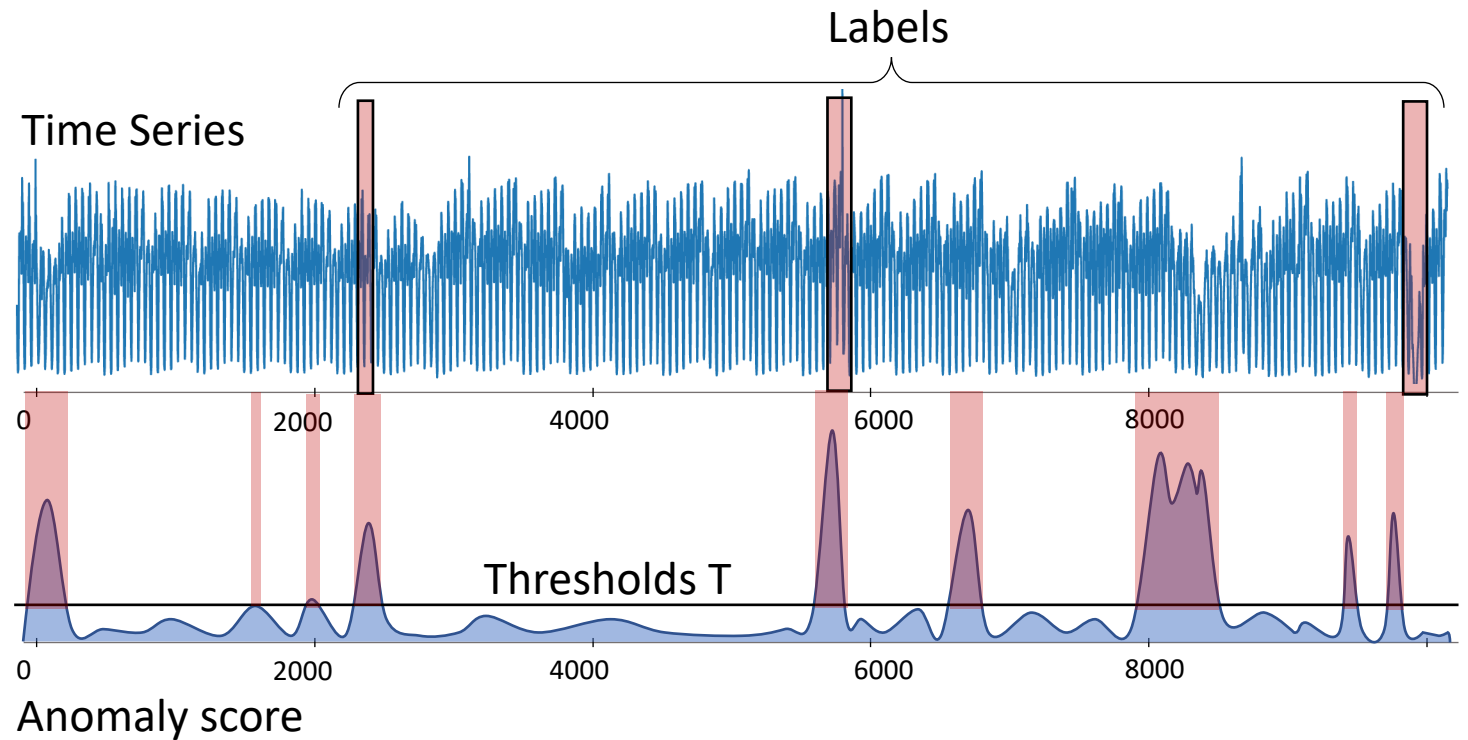
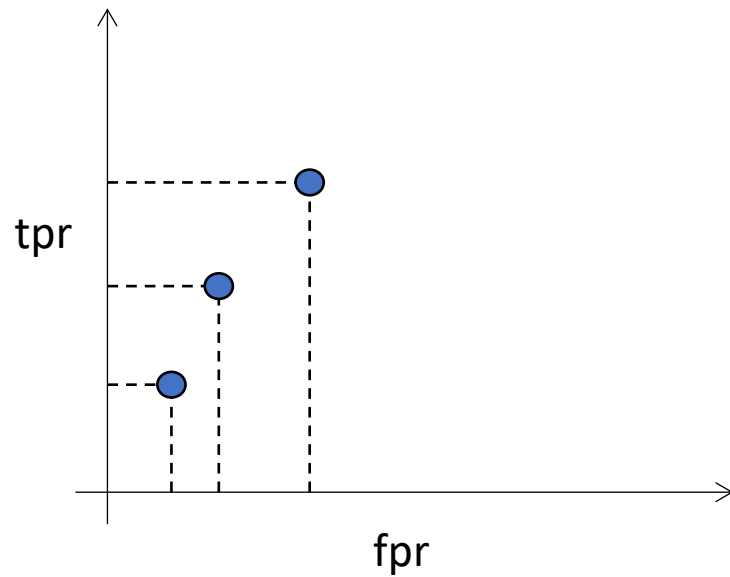
Evaluation measures: *AUC-based*

AUC-based Evaluation Measures:



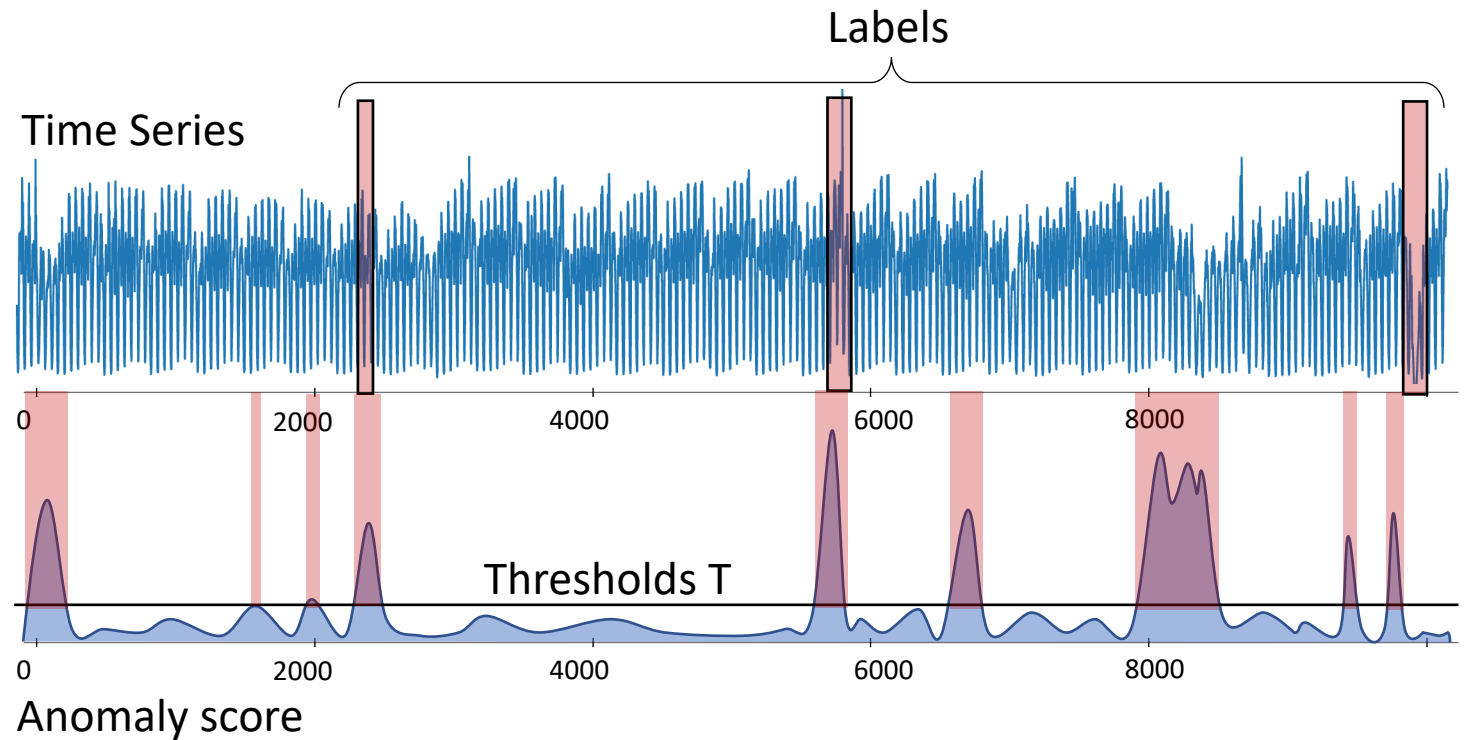
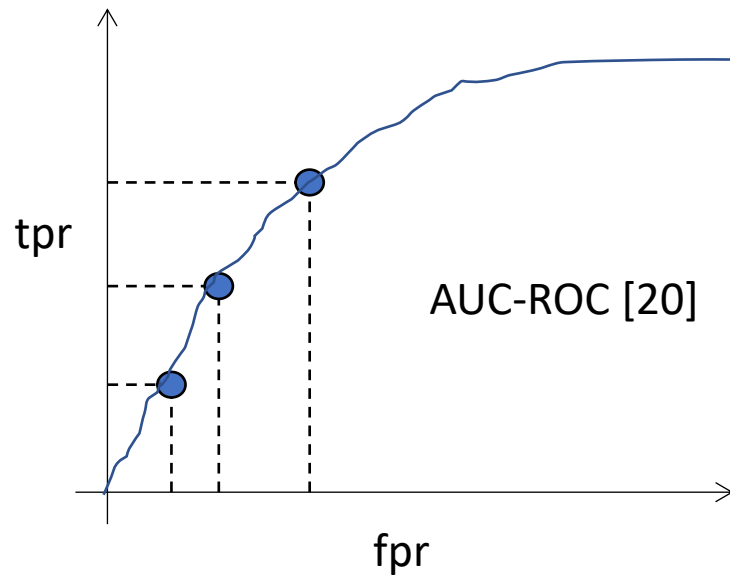
Evaluation measures: *AUC-based*

AUC-based Evaluation Measures:



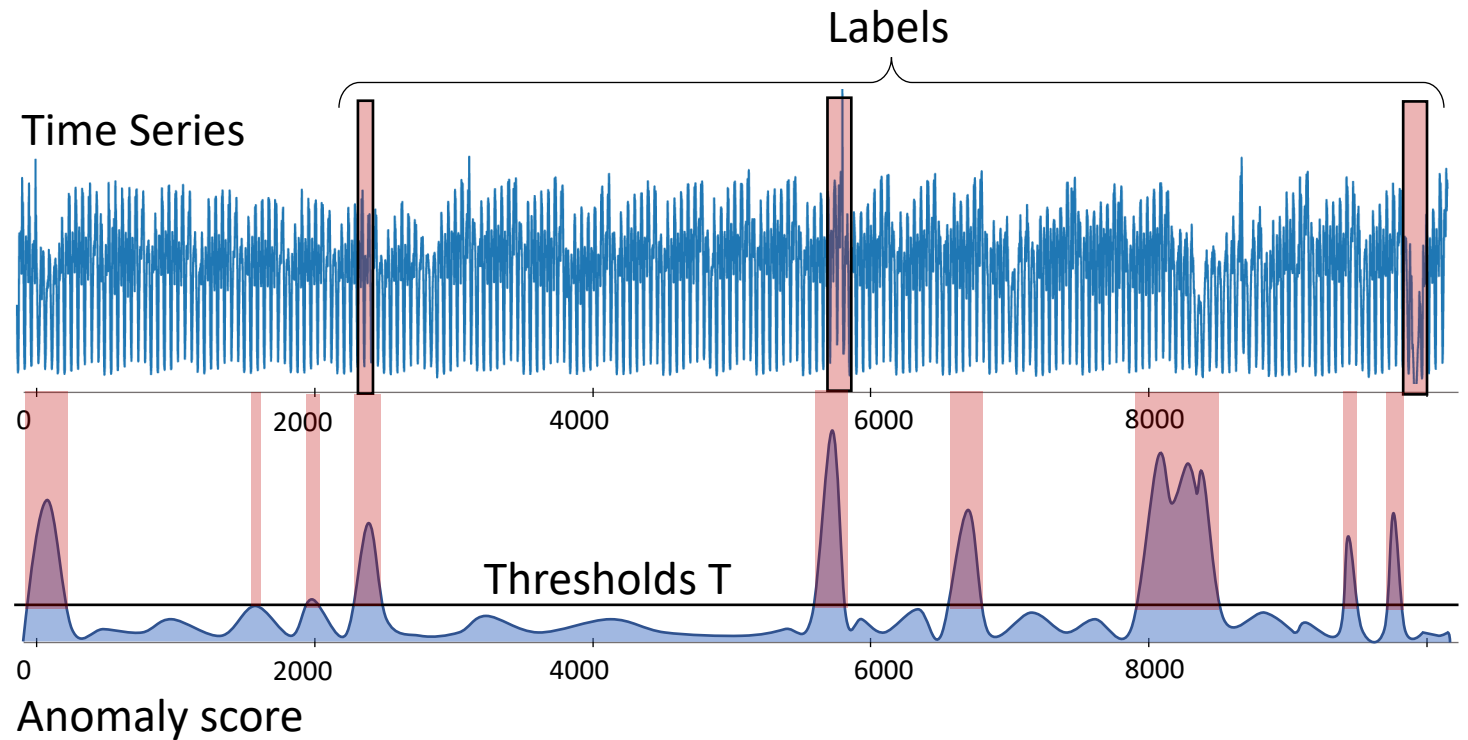
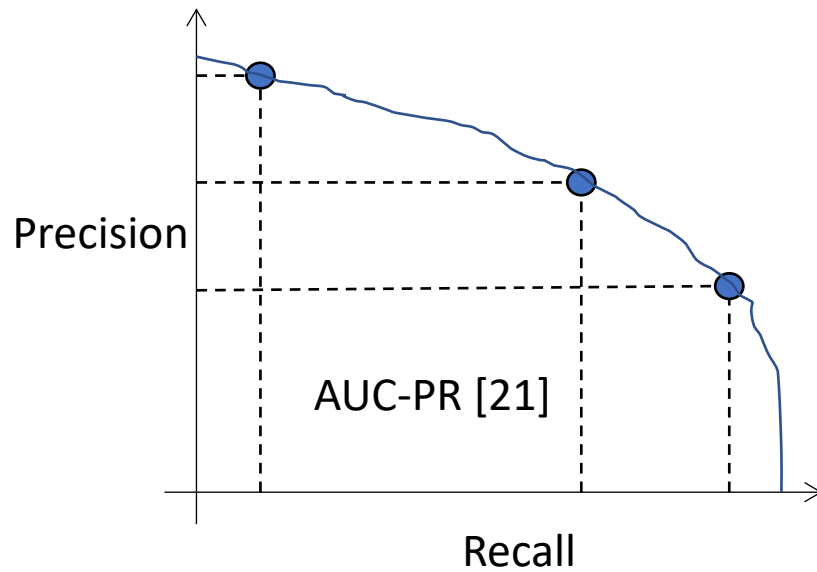
Evaluation measures: *AUC-based*

AUC-based Evaluation Measures:



Evaluation measures: *AUC-based*

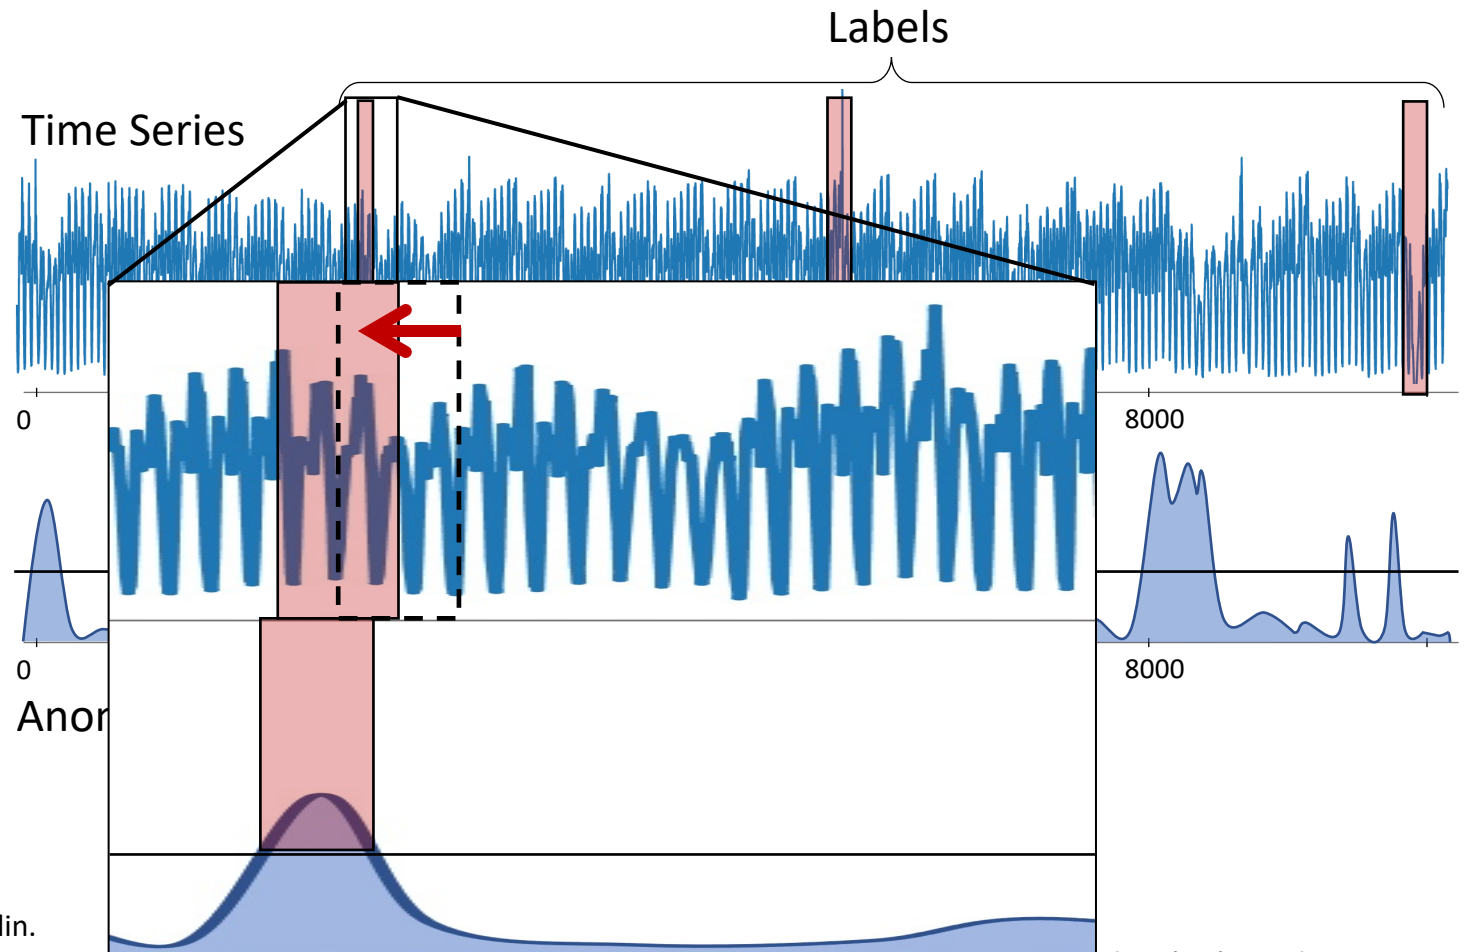
AUC-based Evaluation Measures:



Evaluation measures: *Labeling issue*

Labeling can be an issue for time series [22]:

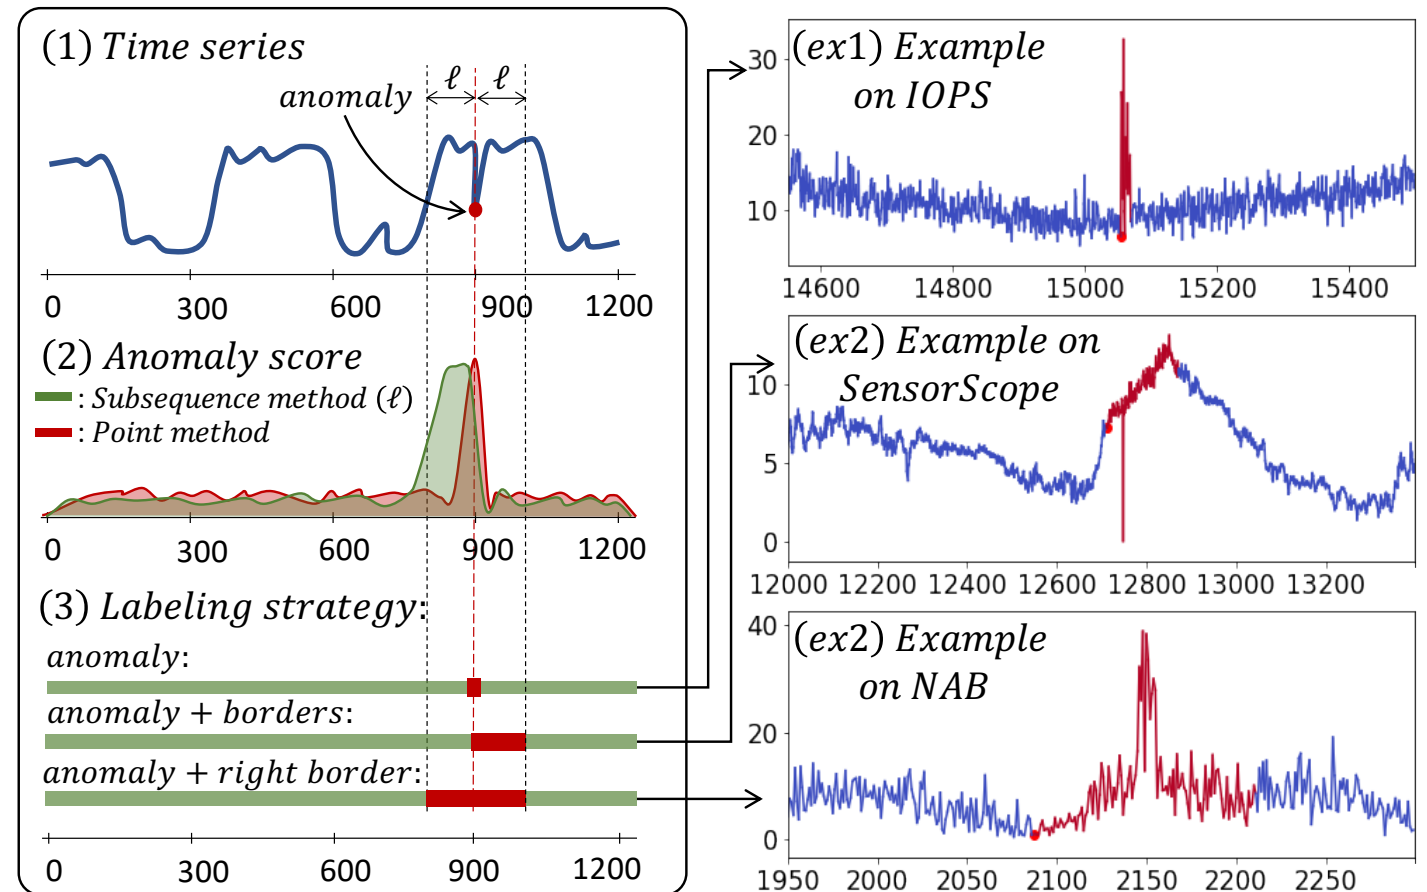
- Misalignment can lead to significant changes of accuracy values.



Evaluation measures: *Labeling issue*

Labeling can be an issue for time series [22]:

- Misalignment can lead to significant changes of accuracy values.
- This is a real issue because of:
 - Different Labeling strategies between domains and applications
 - Methods that produce misaligned anomaly scores.

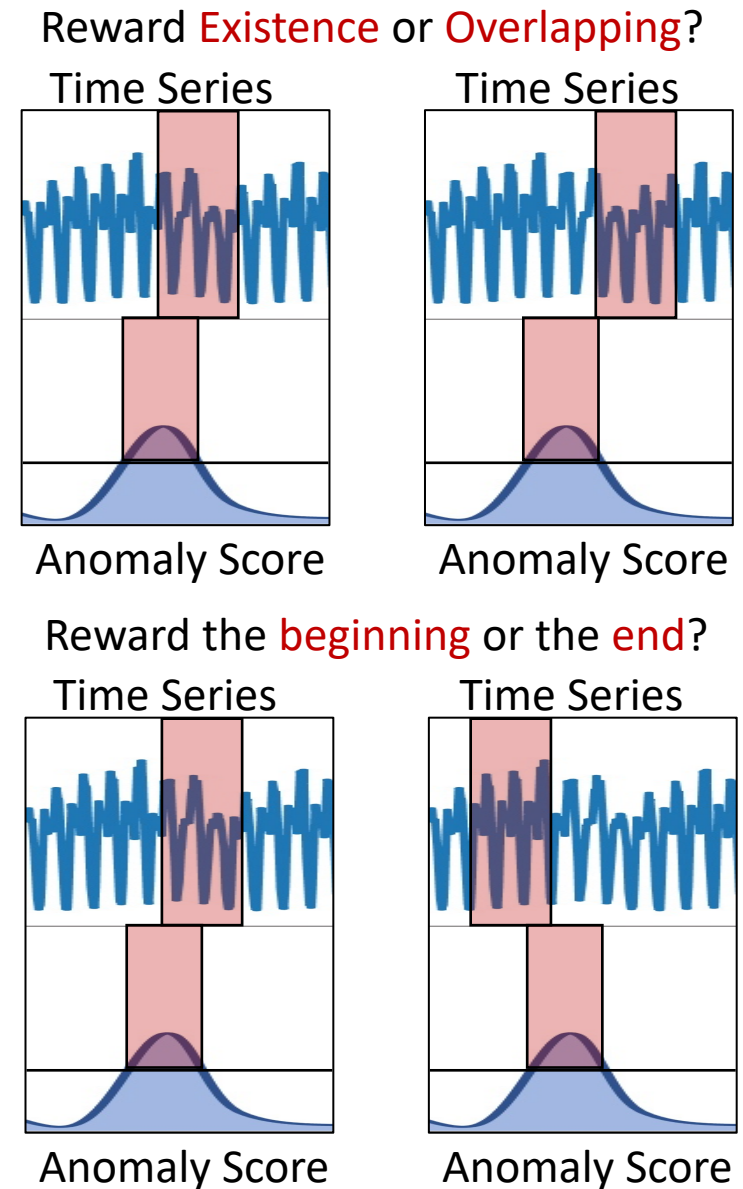


Evaluation measures: *Labeling issue*

Existing solutions:

- Range Precision and Recall [23]:

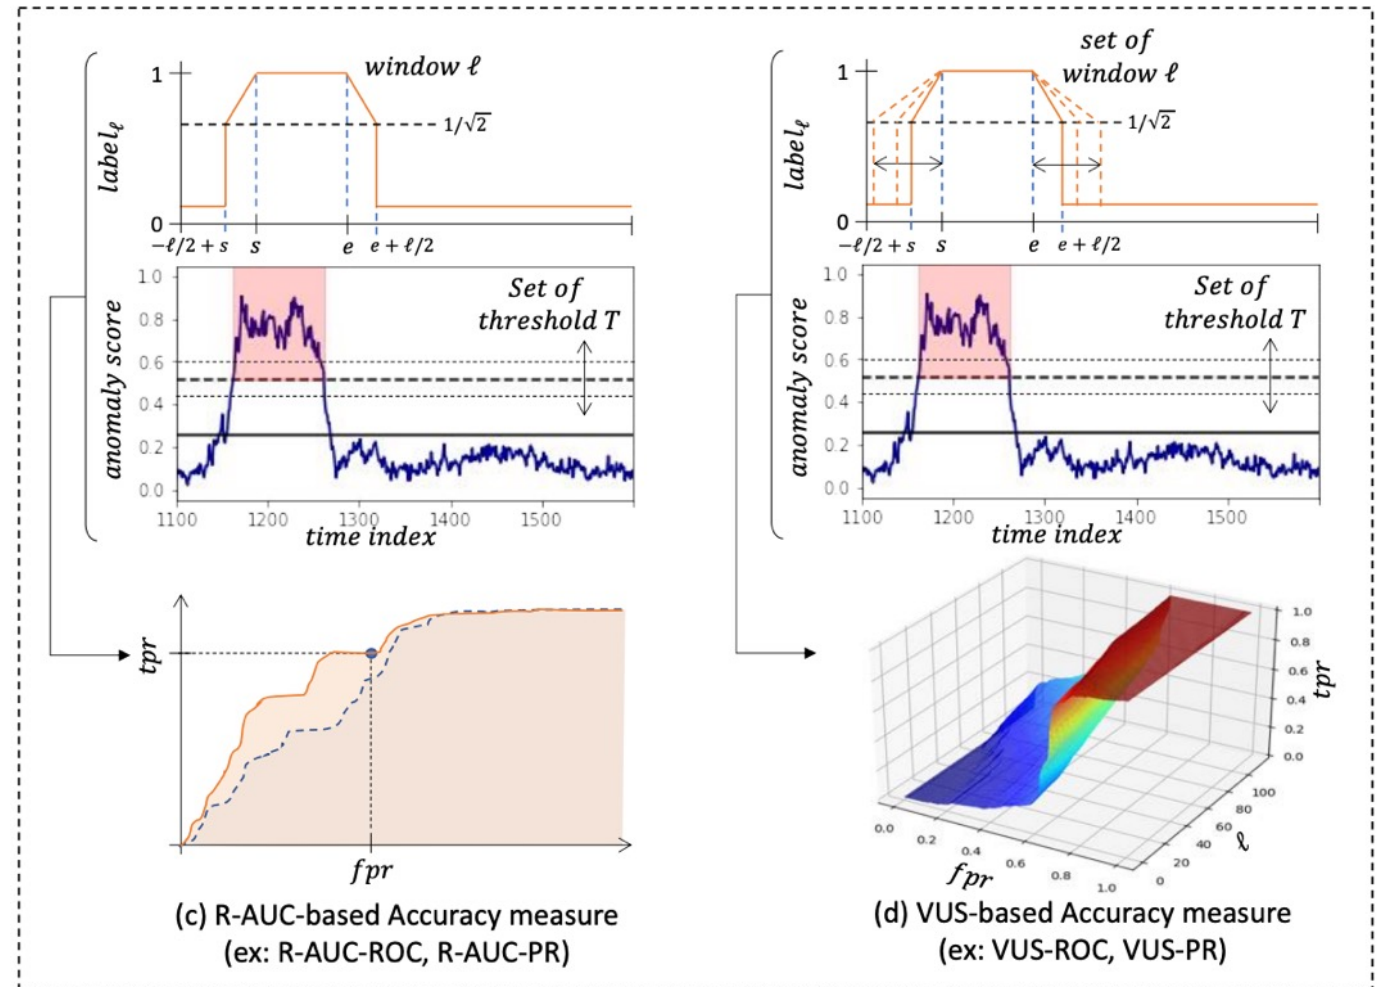
- $Recall_T(R, P) = \frac{\sum_{i=1}^{N_r} Recall_T(R_i, P)}{N_r}$
- $Recall_T(R_i, P) = \alpha * ExistenceR(R_i, P) + (1 - \alpha) * OverlappingR(R_i, P)$
- $Precision_T(R, P) = \frac{\sum_{i=1}^{N_p} Precision_T(R, P_i)}{N_p}$
- $Precision_T(R, P_i) = CardinalityFactor(P_i, R) * \sum_{j=1}^{N_r} w(P_i, P_i \cap R_j, \delta)$
- Functions $w()$, $\delta()$ are tunable functions to represent the overlap size and position respectively.



Evaluation measures: *Labeling issue*

Existing solutions:

- Volume Under the Surface [22] (VUS):
- Modify the labels with buffer regions at the beginning and at the end of an anomaly
- We vary the buffer size (as well as the threshold) and we obtain a surface
- We use the volume under the surface (VUS) as accuracy



Conclusion and Open Problems

Conclusion and Open Problems

If you are interested in anomaly detection in time series...

Anomaly Detection in Time Series: A Comprehensive Evaluation

Sebastian Schmid¹
Hasso Plattner Institute,
University of Potsdam
Potsdam, Germany
sebastian.schmid@hpi.de

Phillip Wenig²
Hasso Plattner Institute,
University of Potsdam
Potsdam, Germany
phillip.wenig@hpi.de

Thorsten Papenbrock
Philips University of Marburg
Marburg, Germany
papenbrock@informatik.uni-marburg.de

ABSTRACT
Detecting anomalous subsequences in time series data is an important task in areas ranging from manufacturing processes over finance applications to health care monitoring. An anomaly can indicate important events, such as production faults, delivery bottlenecks, system defects, or heart flicker, and is therefore of central interest. Because time series are often large and exhibit complex patterns, data scientists have developed various specialized algorithms for the automatic detection of such anomalous patterns. The number and variety of anomaly detection algorithms has grown significantly in the past and, because many of these solutions have been developed independently and by different research communities, there is no comprehensive study that systematically evaluates and compares the different approaches. For this reason, choosing the best detection technique for a given anomaly detection task is a difficult challenge.

This comprehensive, scientific study carefully evaluates most state-of-the-art anomaly detection algorithms. We collected and re-implemented 71 anomaly detection algorithms from different domains and evaluated them on 976 time series datasets. The algorithms have been selected from different algorithm families and detection approaches to represent the entire spectrum of anomaly detection techniques. In the paper, we provide a concise overview of the techniques and their commonalities; we evaluate their individual strengths and weaknesses and, thereby, consider factors, such as effectiveness, efficiency, and robustness. Our experimental results should ease the algorithm selection problem and open up new research directions.

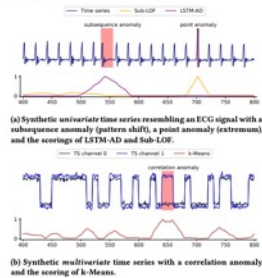


Figure 1: Example time series with anomalies and scorings.

1 ANOMALY DETECTION WILDERNESS

<https://github.com/HPI-Information-Systems/TimeEval>

The data points of a time series record are one or multiple real-valued variables. Each variable models one channel of the time series. If the data points consist of real numbers, the time series

S. Schmidl et al. PVLDB (2022) [5]

TSB-UAD: An End-to-End Benchmark Suite for Univariate Time-Series Anomaly Detection

John Paparrizos
The Ohio State University
paparrizos@osu.edu

Yuhao Kang
University of Chicago
yuhao.kang@chicagogov.edu

Paul Boniol
University of Paris
paul.boniol@univ-paris.fr

Ruey S. Tsay
University of Chicago
rueytsay@chicagobooth.edu

Therms Palpanas
University of Paris & IUF
therms@univ-paris.fr

Michael J. Franklin
University of Chicago
mjfranklin@chicagogov.edu

ABSTRACT
The detection of anomalies in time series has gained ample academic and industrial attention. However, no comprehensive benchmark exists to evaluate time-series anomaly detection methods. It is common to use (i) proprietary or synthetic data, often biased to support particular claims; or (ii) a limited collection of publicly available datasets. Consequently, we often observe methods performing exceptionally well in one dataset but surprisingly poorly in another, creating an illusion of progress. To address the issues above, we thoroughly studied over one hundred papers to identify, collect, process, and systematically format datasets proposed in the past decades. We summarize our effort in TSB-UAD, a new benchmark to ease the evaluation of univariate time-series anomaly detection methods. Overall, TSB-UAD contains 13766 time series with labeled anomalies spanning different domains with high variability of anomaly types, ratios, and sizes. TSB-UAD includes 18 previously proposed datasets containing 1980 time series and we contribute two collections of datasets. Specifically, we generate 958 time series using a principled methodology for transforming 128 time-series classification datasets into time series with labeled anomalies. In addition, we present data transformations with which we introduce new anomalies, resulting in 10828 time series with varying complexity for anomaly detection. Finally, we evaluate 12 representative methods demonstrating that TSB-UAD is a robust resource for assessing anomaly detection methods. TSB-UAD provides a valuable, reproducible, and frequently updated resource to establish a leaderboard of time-series anomaly detection methods.

<https://github.com/TheDatumOrg/TSB-UAD>

A wide range of technological advances in sensing solutions enables collecting enormous amounts of time-varying measurements commonly referred to as time series. In particular, analysts estimate

J. Paparrizos et al. PVLDB (2022) [19]

Current Time Series Anomaly Detection Benchmarks are Flawed and are Creating the Illusion of Progress

Renjie Wu and Eamonn J. Keogh

Abstract—Time series anomaly detection has been a perennially important topic in data science, with papers dating back to the 1950s. However, in recent years there has been an explosion of interest in this topic, much of it driven by the success of deep learning in other domains and for other time series tasks. Most of these papers test on one or more of a handful of popular benchmark datasets, created by Yahoo, Numerix, NASA, etc. In this work we make a surprising claim. The majority of the individual exemplars in these datasets suffer from one or more of four flaws. Because of these four flaws, we believe that many published comparisons of anomaly detection algorithms may be unreliable, and more importantly, much of the apparent progress in recent years may be illusory. In addition to demonstrating these claims, with this paper we introduce the UCR Time Series Anomaly Archive. We believe that this resource will perform a similar role as the UCR Time Series Classification Archive, by providing the community with a benchmark that will allow meaningful comparisons between approaches and a meaningful gauge of overall progress.

Index Terms—Anomaly detection, benchmark datasets, deep learning, time series analysis

1 INTRODUCTION

TIME series anomaly detection has been a perennially important topic in data science, with papers dating back to the dawn of computer science [1]. However, in the last five years there has been an explosion of interest in this topic, with at least one or two papers on the topic appearing each year in virtually every database, data mining, and machine learning conference, including SIGKDD [2], [3], ICDM [4], ICDE, SIGMOD, VLDB, etc.

A large fraction of this increase in interest seems to be largely driven by researchers anxious to transfer the considerable success of deep learning in other domains and

neural networks, and a variational auto-encoder (VAE) over-sampling model.” This description sounds like it has many “moving parts”, and indeed, the dozen or so explicitly listed parameters include: convolution filter, activation, kernel size, strides, padding, LSTM input size, dense input size, softmax loss function, window size, learning rate and batch size. All of this is to demonstrate “accuracy exceeding 0.90 (on a subset of the Yahoo’s anomaly detection benchmark datasets).” However, as we will show, much of the results of this complex approach can be duplicated with a single line of code and a few minutes of effort.

Suppose however that someone downloaded the original

R. Wu et al. TKDE (2021) [18]

A review on outlier/anomaly detection in time series data

ANE BLÁZQUEZ-GARCÍA and ANGEL CONDE, Iberlan Technology Research Centre, Basque Research and Technology Alliance (BRTA), Spain

USUE MORI, Intelligent Systems Group (ISG), Department of Computer Science and Artificial Intelligence, University of the Basque Country (UPV/EHU), Spain

JOSE A. LOZANO, Intelligent Systems Group (ISG), Department of Computer Science and Artificial Intelligence, University of the Basque Country (UPV/EHU), Spain and Basque Center for Applied Mathematics (BCAM), Spain

Recent advances in technology have brought major breakthroughs in data collection, enabling a large amount of data to be gathered over time and thus generating time series. Mining this data has become an important task for researchers and practitioners in the past few years, including the detection of outliers or anomalies that may represent errors or events of interest. This review aims to provide a structured and comprehensive state-of-the-art on outlier detection techniques in the context of time series. To this end, a taxonomy is presented based on the main aspects that characterize an outlier detection technique.

Additional Key Words and Phrases: Outlier detection, anomaly detection, time series, data mining, taxonomy, software

1 INTRODUCTION

Recent advances in technology allow us to collect a large amount of data over time in diverse research areas. Observations that have been recorded in an orderly fashion and which are correlated in time constitute a time series. Time series data mining aims to extract all meaningful knowledge from this data, and several mining tasks (e.g., classification, clustering, forecasting, and outlier detection) have been considered in the literature [Elaing and Agon 2012; Fu 2011; Ratanamahatana et al. 2010].

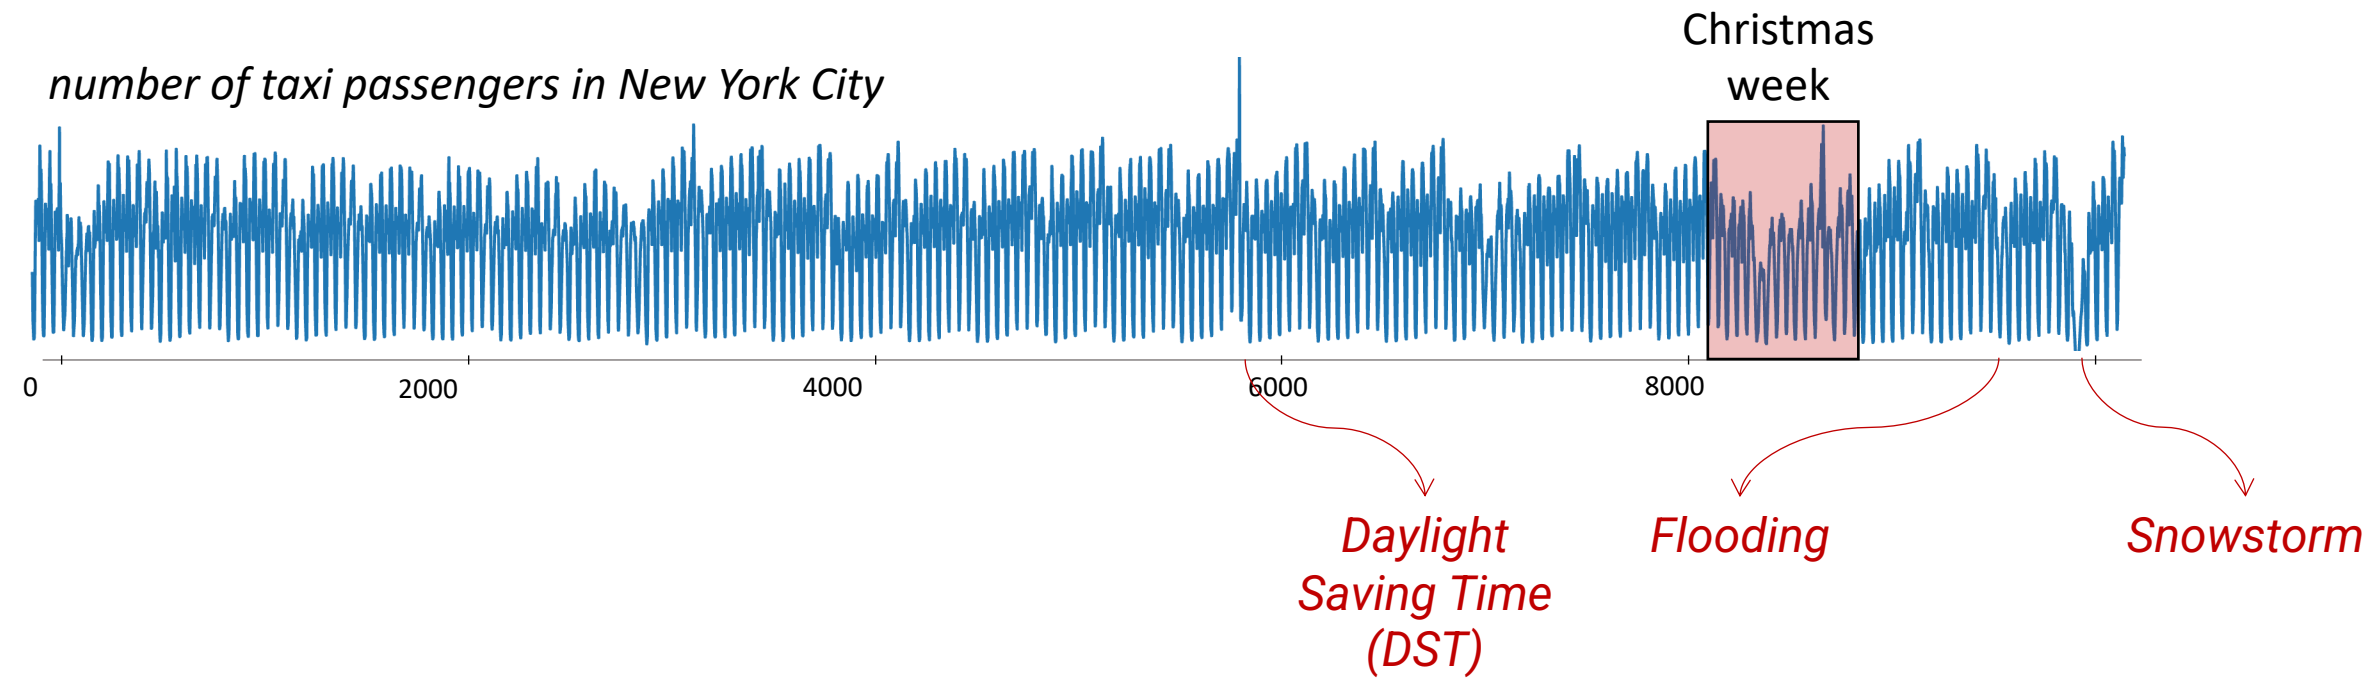
Outlier detection has become a field of interest for many researchers and practitioners and is now one of the main tasks of time series data mining. Outlier detection has been studied in a variety of application domains such as credit card fraud detection, intrusion detection in cybersecurity, or fault diagnosis in industry. In particular, the analysis of outliers in time series data examines anomalous behaviors across time [Gupta et al. 2014a]. In the first study on this topic, which was conducted by Fox [1972], two types of outliers in univariate time series were defined: type I, which affects a single observation; and type II, which affects both a particular observation and the subsequent observations. This work was first extended to four outlier types [Tsay 1988], and then to the case of multivariate time series [Tsay et al. 2000]. Since then, many definitions of the term *outlier* and numerous detection methods have been proposed in the literature. However, to this day, there is still no consensus on the terms used [Caretto et al. 2019]; for example, outlier observations are often referred to as anomalies, discordant observations, discords, exceptions, aberrations, surprises, peculiarities or contaminants.

Authors’ addresses: Ane Blázquez-García, ablaquez@iberlan.es, Iberlan Technology Research Centre, Basque Research and Technology Alliance (BRTA), P.TAM, Arzobispado, 2, Arzobispado, 48940, Spain; Usue Mori, usue.mori@ehu.es, Intelligent Systems Group (ISG), Department of Computer Science and Artificial Intelligence, University of the Basque Country (UPV/EHU), 48940, Spain; Jose A. Lozano, jalozano@bcam.es, Basque Center for Applied Mathematics (BCAM), 48940, Spain.

A. Blazquez-Garcia et al. ACM Computing Survey (2021) [24]

Conclusion and Open Problems

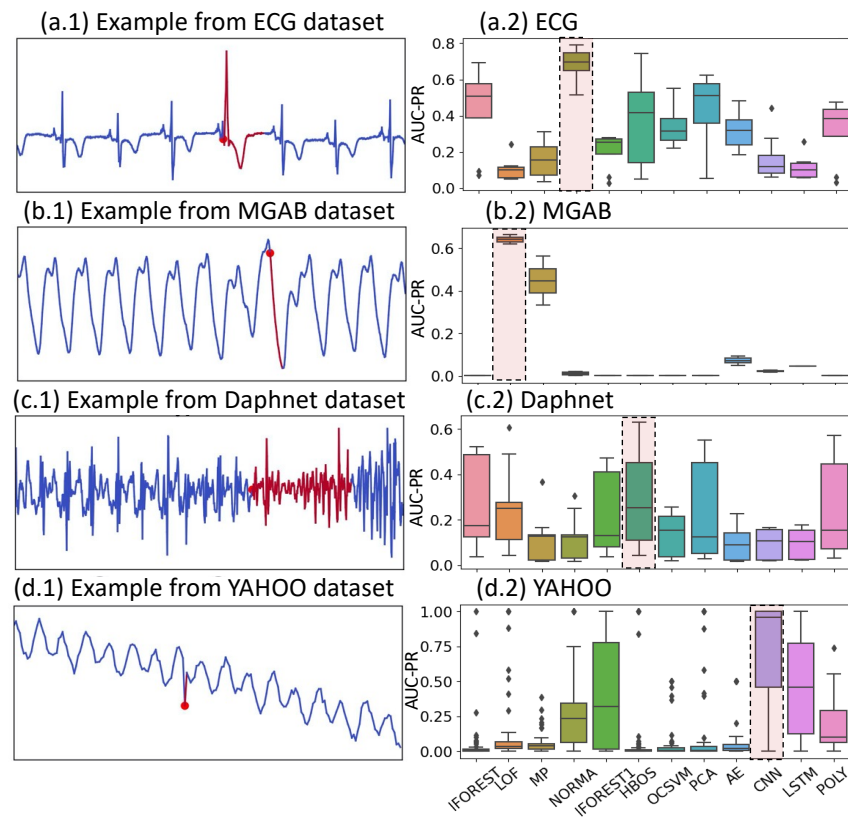
Context-aware Unsupervised Anomaly Detection



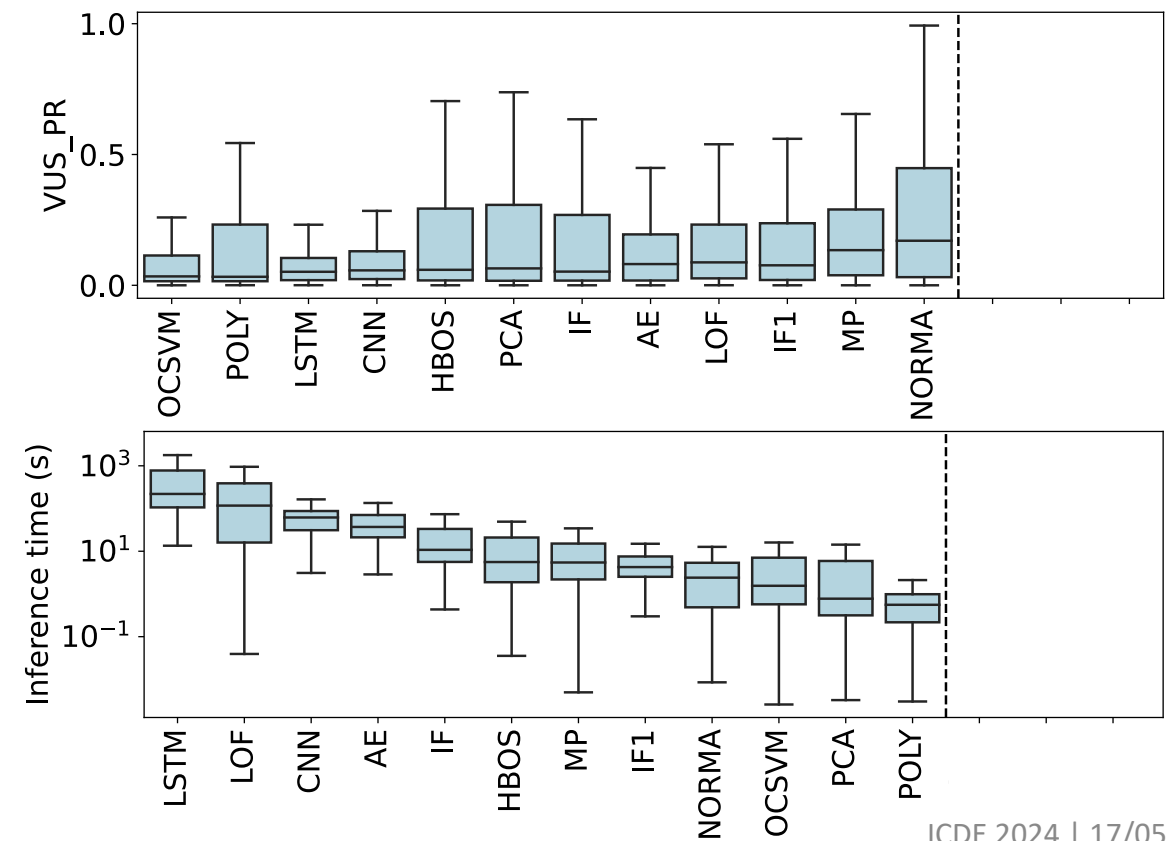
Conclusion and Open Problems

Model selection for anomaly detection

Methods ranking changes
significantly between datasets [19]



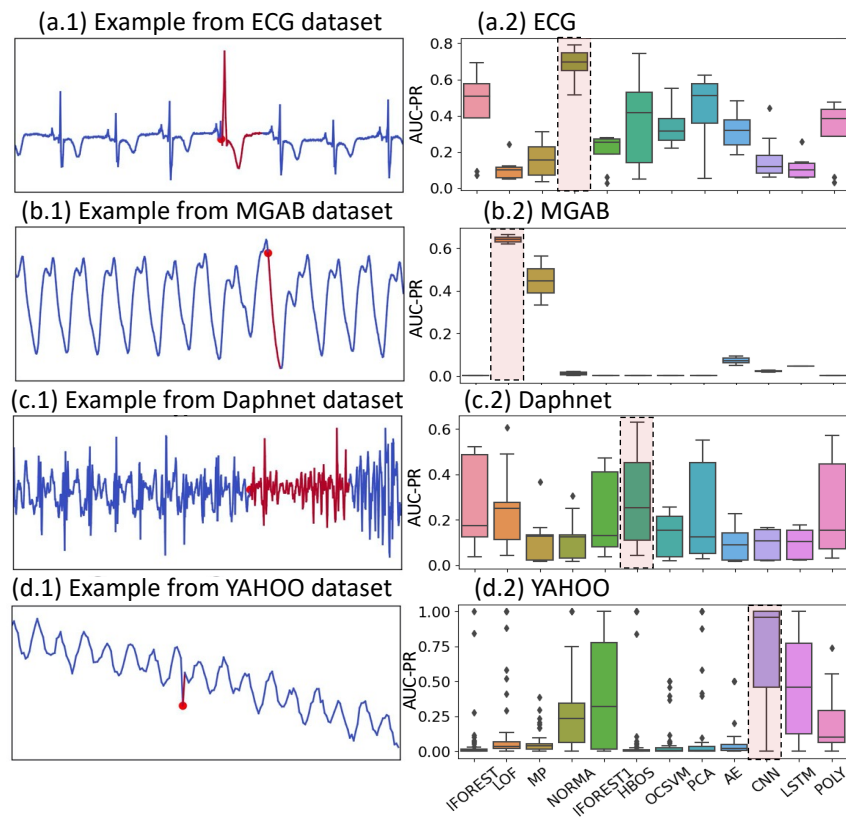
Results over TSB-UAD



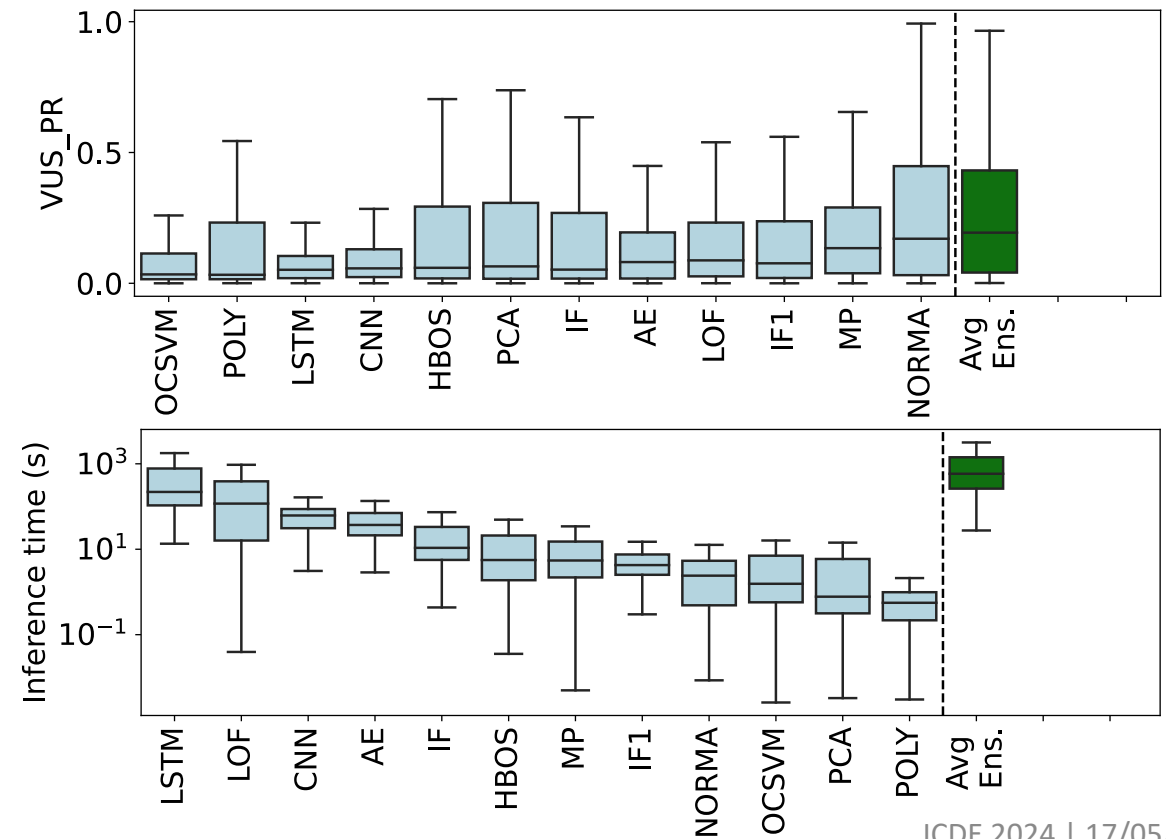
Conclusion and Open Problems

Model selection for anomaly detection

Methods ranking changes significantly between datasets [19]



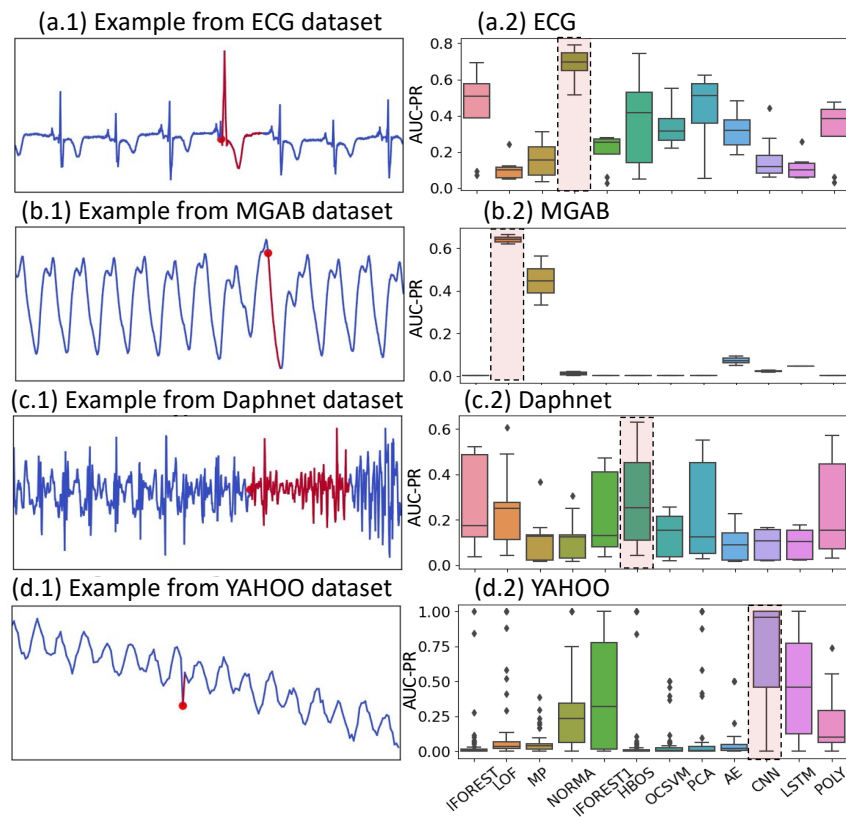
Can *Ensembling* methods solve the problem?



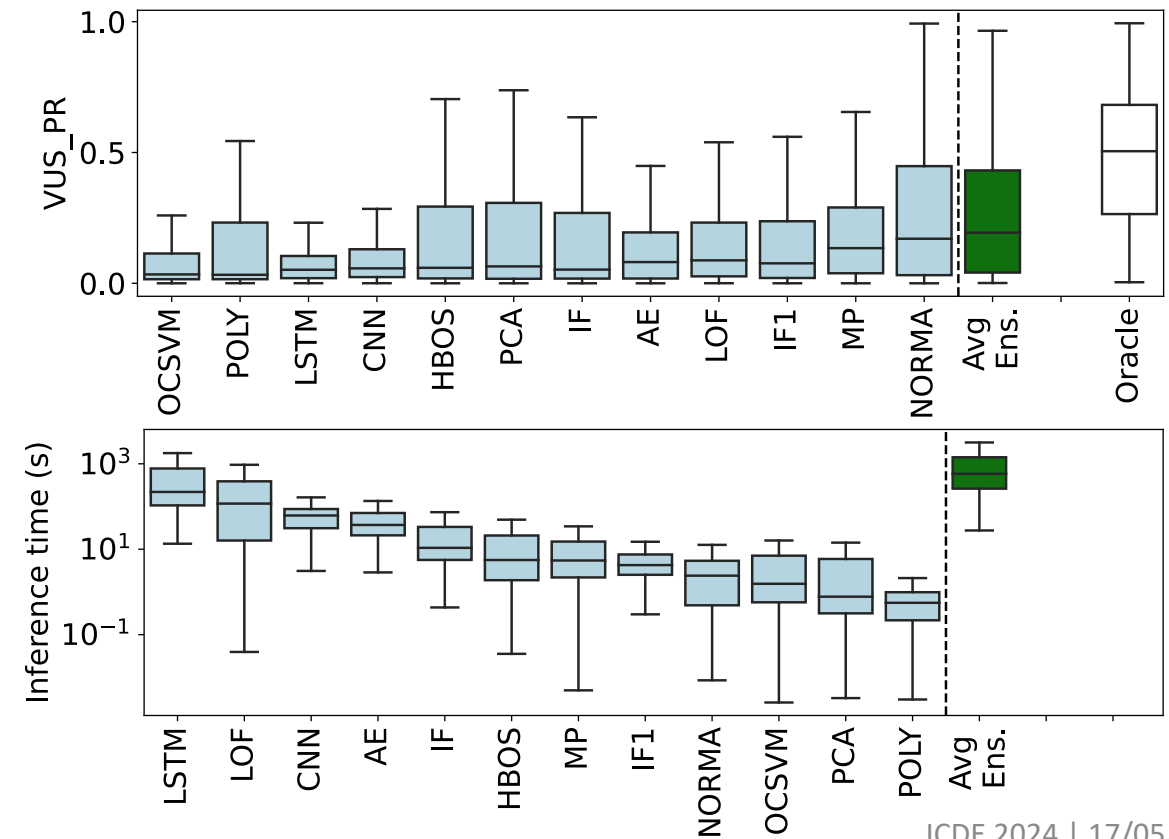
Conclusion and Open Problems

Model selection for anomaly detection

Methods ranking changes significantly between datasets [19]



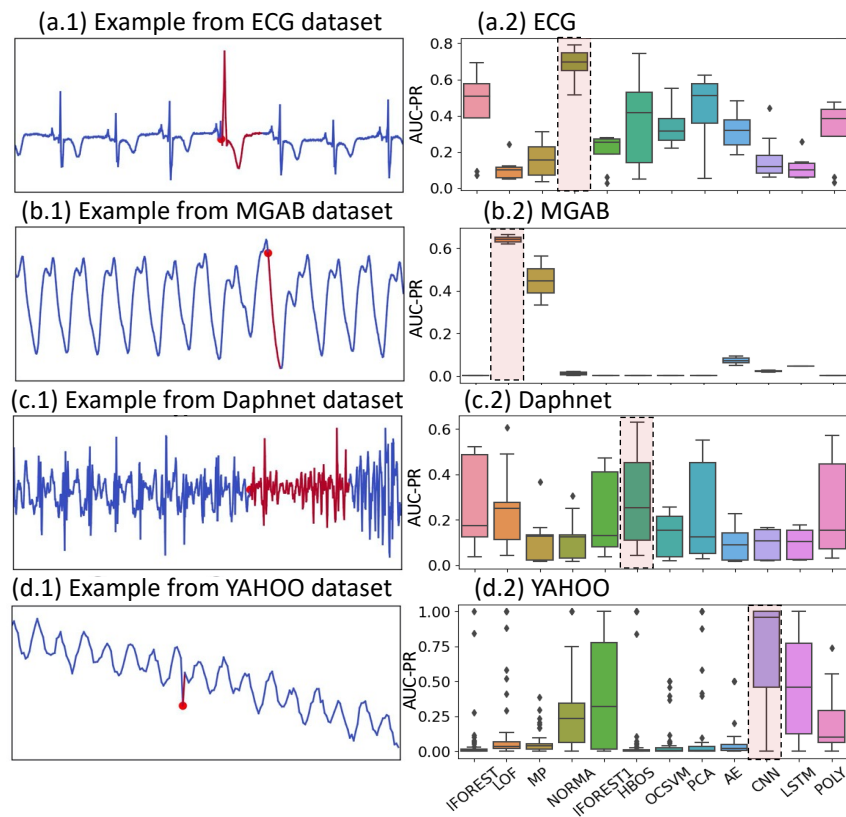
Can *automatic model selection* solve the problem?



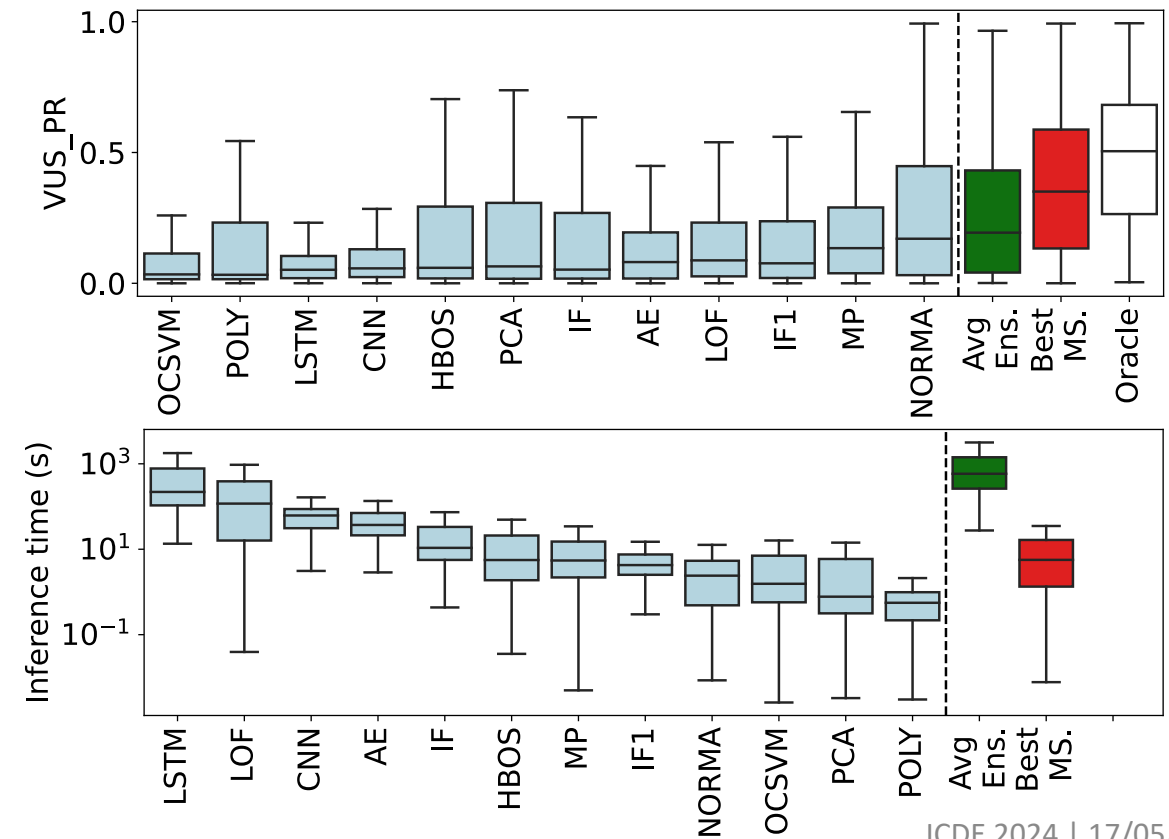
Conclusion and Open Problems

Model selection for anomaly detection

Methods ranking changes significantly between datasets [19]



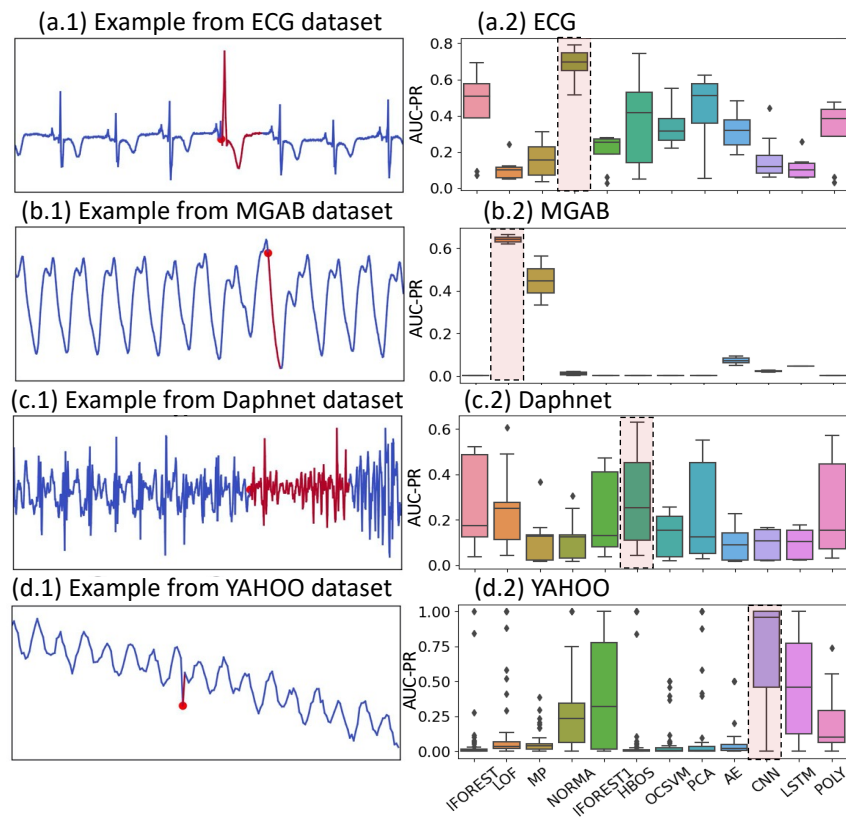
Can *automatic model selection* solve the problem?



Conclusion and Open Problems

Model selection for anomaly detection

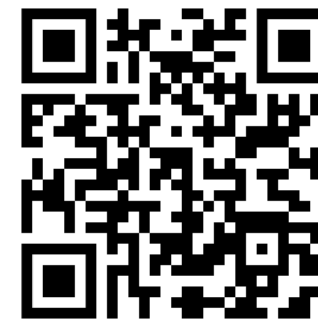
Methods ranking changes significantly between datasets [19]



Can *automatic model selection* solve the problem?

Choose Wisely [29]

An experimental evaluation of model selection for time series anomaly detection



VLDB 2023



ICDE 2024

[29] 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. Proc. VLDB Endow. 16, 11 (July 2023), 3418–3432.

References

- [1] N. Laptev, S. Amizadeh, and Y. Billawala. 2015. **S5 - A Labeled Anomaly Detection Dataset**, version 1.0(16M).
- [2] Markus Thill, Wolfgang Konen, and Thomas Bäck. 2020. **MGAB: The Mackey-Glass Anomaly Benchmark**.
- [3] Pawel Benecki, Szymon Piechaczek, Daniel Kostrzewa, and Jakub Nalepa. 2021. **Detecting Anomalies in Spacecraft Telemetry Using Evolutionary Thresholding and LSTMs**. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (Lille, France) (GECCO '21)
- [4] Scott David Greenwald. 1990. **Improved detection and classification of arrhythmias in noise-corrupted electrocardiograms using contextual information**. Thesis. Massachusetts Institute of Technology.
- [5] Sebastian Schmidl, Phillip Wenig, and Thorsten Papenbrock. 2022. **Anomaly detection in time series: a comprehensive evaluation**. Proc. VLDB Endow. 15, 9 (May 2022), 1779–1797.
- [6] Chin-Chia Michael Yeh, Yan Zhu, Liudmila Ulanova, Nurjahan Begum, Yifei Ding, Hoang Anh Dau, Diego Furtado Silva, Abdullah Mueen, and Eamonn J. Keogh. 2016. **Matrix Profile I: All Pairs Similarity Joins for Time Series**. In ICDM.
- [7] Yan Zhu, Zachary Zimmerman, Nader Shakibay Senobari, Chin-Chia Michael Yeh, Gareth Funning, Abdullah Mueen, Philip Brisk, and Eamonn Keogh. 2016. **Matrix Profile II: Exploiting a Novel Algorithm and GPUs to Break the One Hundred Million Barrier for Time Series Motifs and Joins**. In Proceedings of the International Conference on Data Mining (ICDM), 739–748.
- [8] Yue Lu, Renjie Wu, Abdullah Mueen, Maria A. Zuluaga, and Eamonn Keogh. 2022. **Matrix Profile XXIV: Scaling Time Series Anomaly Detection to Trillions of Datapoints and Ultra-fast Arriving Data Streams**. In Proceedings of the 28th ACM SIGKDD.
- [9] C. -C. M. Yeh, N. Kavantzaz and E. Keogh, **Matrix Profile VI: Meaningful Multidimensional Motif Discovery**, 2017 IEEE International Conference on Data Mining (ICDM), New Orleans, LA, USA, 2017, pp. 565-574, doi: 10.1109/ICDM.2017.66. Data Mining (KDD '22).
- [10] Paul Boniol, Michele Linardi, Federico Roncallo, Themis Palpanas, Mohammed Meftah, and Emmanuel Remy. 2021. **Unsupervised and scalable subsequence anomaly detection in large data series**. The VLDB Journal 30, 6 (Nov 2021), 909–931.
- [11] F. T. Liu, K. M. Ting and Z. -H. Zhou, **Isolation Forest**, 2008 Eighth IEEE International Conference on Data Mining, Pisa, Italy, 2008, pp. 413-422
- [12] Markus Goldstein and Andreas Dengel. 2012. **Histogram-based outlier score (hbos): A fast unsupervised anomaly detection algorithm**. KI-2012: poster and demo track 9 (2012).
- [13] Paul Boniol and Themis Palpanas. 2020. **Series2Graph: graph-based subsequence anomaly detection for time series**. Proc. VLDB Endow. 13, 12 (August 2020), 1821–1834.
- [14] Ali Abdul-Aziz, Mark R Woike, Nikunj C Oza, Bryan L Matthews, and John D Iekki. 2012. **Rotor health monitoring combining spin tests and data-driven anomaly detection methods**. Structural Health Monitoring (2012).
- [15] Pankaj Malhotra, Lovekesh Vig, Gautam Shro, and Puneet Agarwal. 2015. **Long Short Term Memory Networks for Anomaly Detection in Time Series**. (2015).
- [16] 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.
- [17] Mayu Sakurada and Takehisa Yairi. 2014. **Anomaly Detection Using Autoencoders with Nonlinear Dimensionality Reduction**. In Proceedings of the MLSDA 2014 2nd Workshop on Machine Learning for Sensory Data Analysis (Gold Coast, Australia QLD, Australia) (MLSDA'14).
- [18] R. Wu and E. Keogh, **Current Time Series Anomaly Detection Benchmarks are Flawed and are Creating the Illusion of Progress** in IEEE Transactions on Knowledge & Data Engineering, vol. 35, no. 03, pp. 2421-2429, 2023.
- [19] 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**. Proc. VLDB Endow. 15, 8 (April 2022), 1697–1711.
- [20] Tom Fawcett. 2006. **An introduction to ROC analysis**. Pattern Recognition Letters 27, 8 (2006), 861–874.
- [21] Jesse Davis and Mark Goadrich. 2006. **The Relationship between Precision-Recall and ROC Curves**. In Proceedings of the 23rd International Conference on Machine Learning (ICML '06).
- [22] 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 (July 2022), 2774–2787.
- [23] Nesime Tatbul, Tae Jun Lee, Stan Zdonik, Mejbah Alam, and Justin Gottschlich. 2018. **Precision and Recall for Time Series**. In Advances in Neural Information Processing Systems, Vol. 31.
- [24] Ane Blázquez-García, Angel Conde, Usue Mori, and Jose A. Lozano. 2021. **A Review on Outlier/Anomaly Detection in Time Series Data**. ACM Comput. Surv. 54, 3, Article 56 (April 2022), 33 pages.
- [25] Paul Boniol, John Paparrizos, Themis Palpanas, and Michael J. Franklin. 2021. **SAND: streaming subsequence anomaly detection**. Proc. VLDB Endow. 14, 10 (June 2021), 1717–1729.
- [26] Schneider, J., Wenig, P. & Papenbrock, T. **Distributed detection of sequential anomalies in univariate time series**. The VLDB Journal 30, 579–602 (2021).

Thank you for attending!

Any Questions?