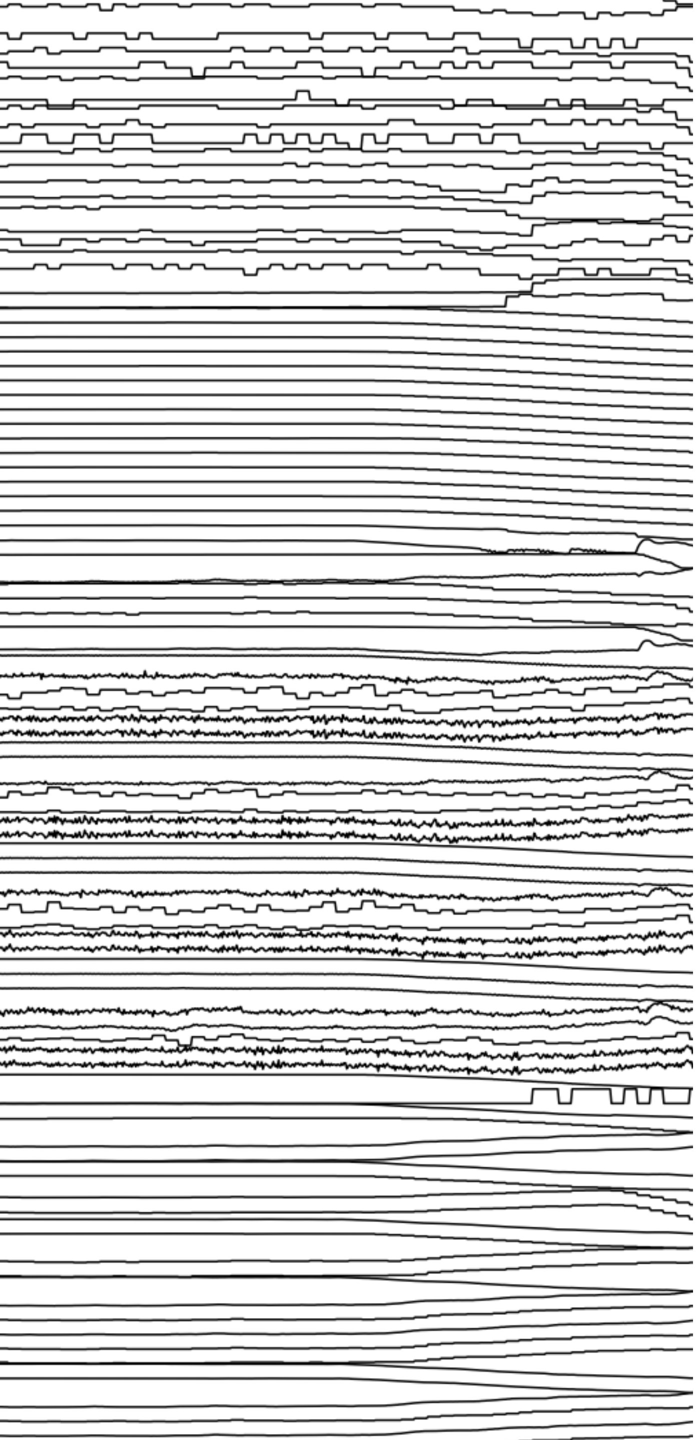

Time Series Anomaly Detection: An Overview

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

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

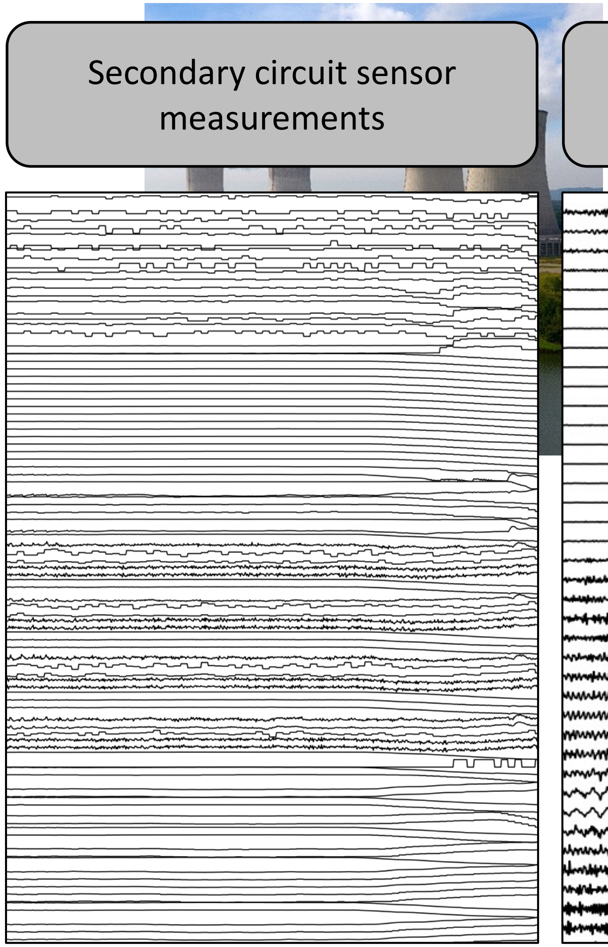

Astrophysics

Fiber-acoustic sensors in the VIRGO north building

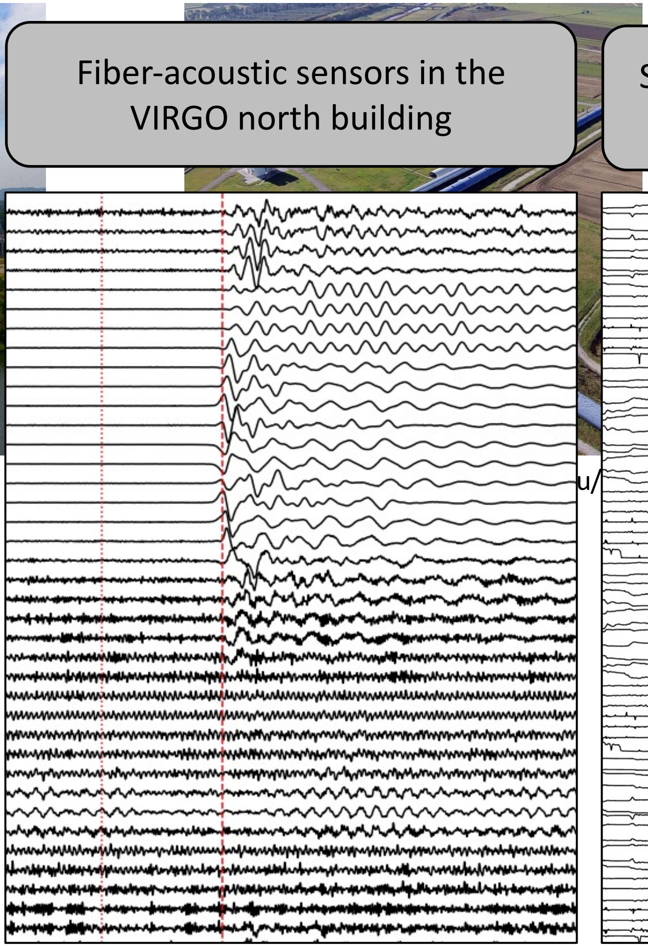

Medicine

Sensor measurements of the Da-Vinci surgery robot

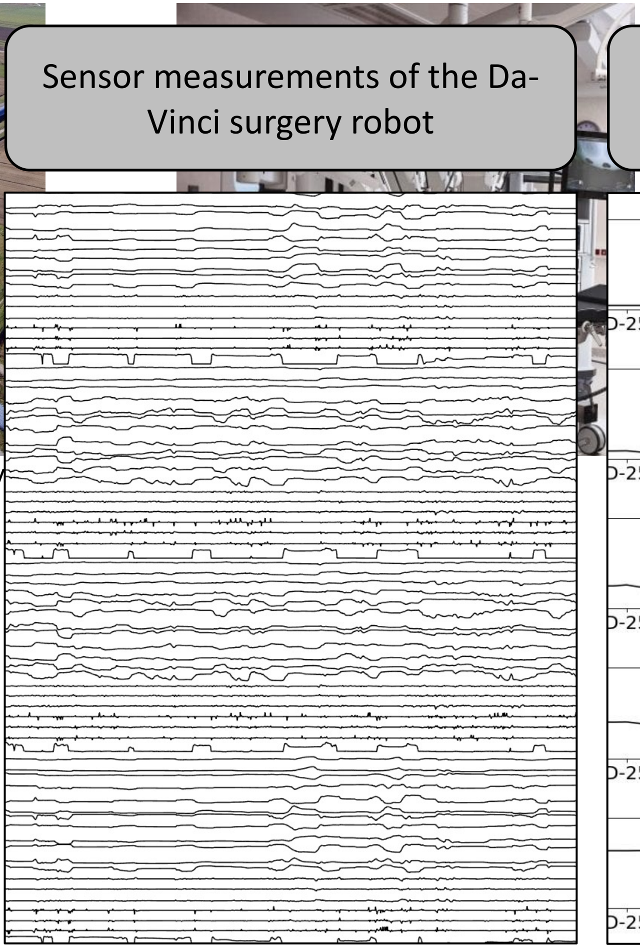

Volcanology

Sensor measurements on le Piton de la Fournaise

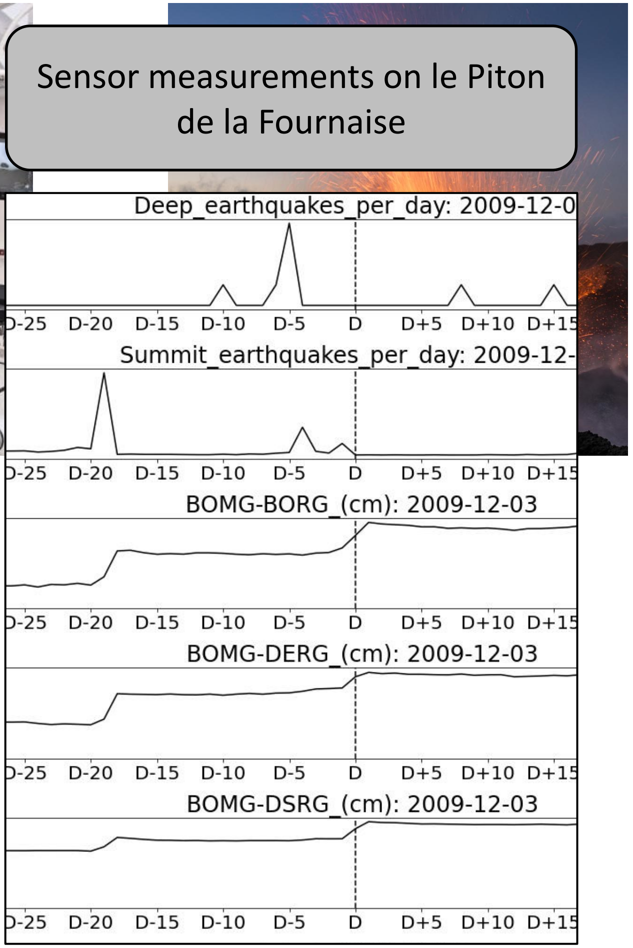

Introduction: *with Important Challenges*

Energy Production

Secondary circuit sensor measurements

Identification of precursors of feed-water pumps vibrations

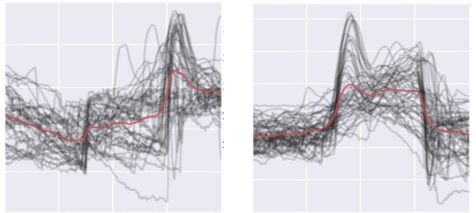

Astrophysics

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

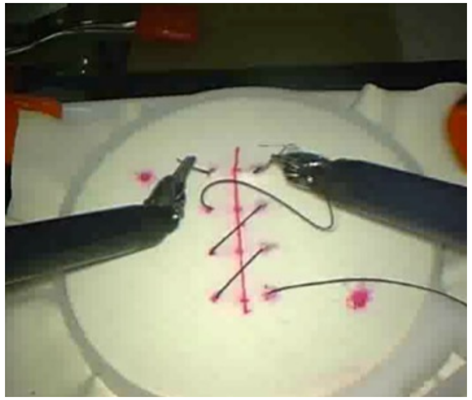

Volcanology

Sensor measurements on le Piton de la Fournaise

Detection of abnormal events on the volcano for predicting eruptions

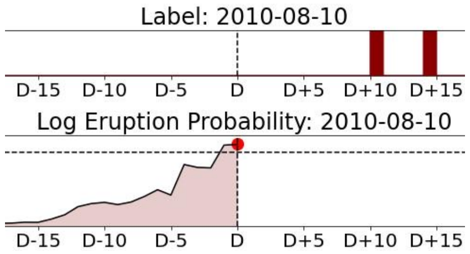

Introduction: *Anomaly Detection in Time Series*

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

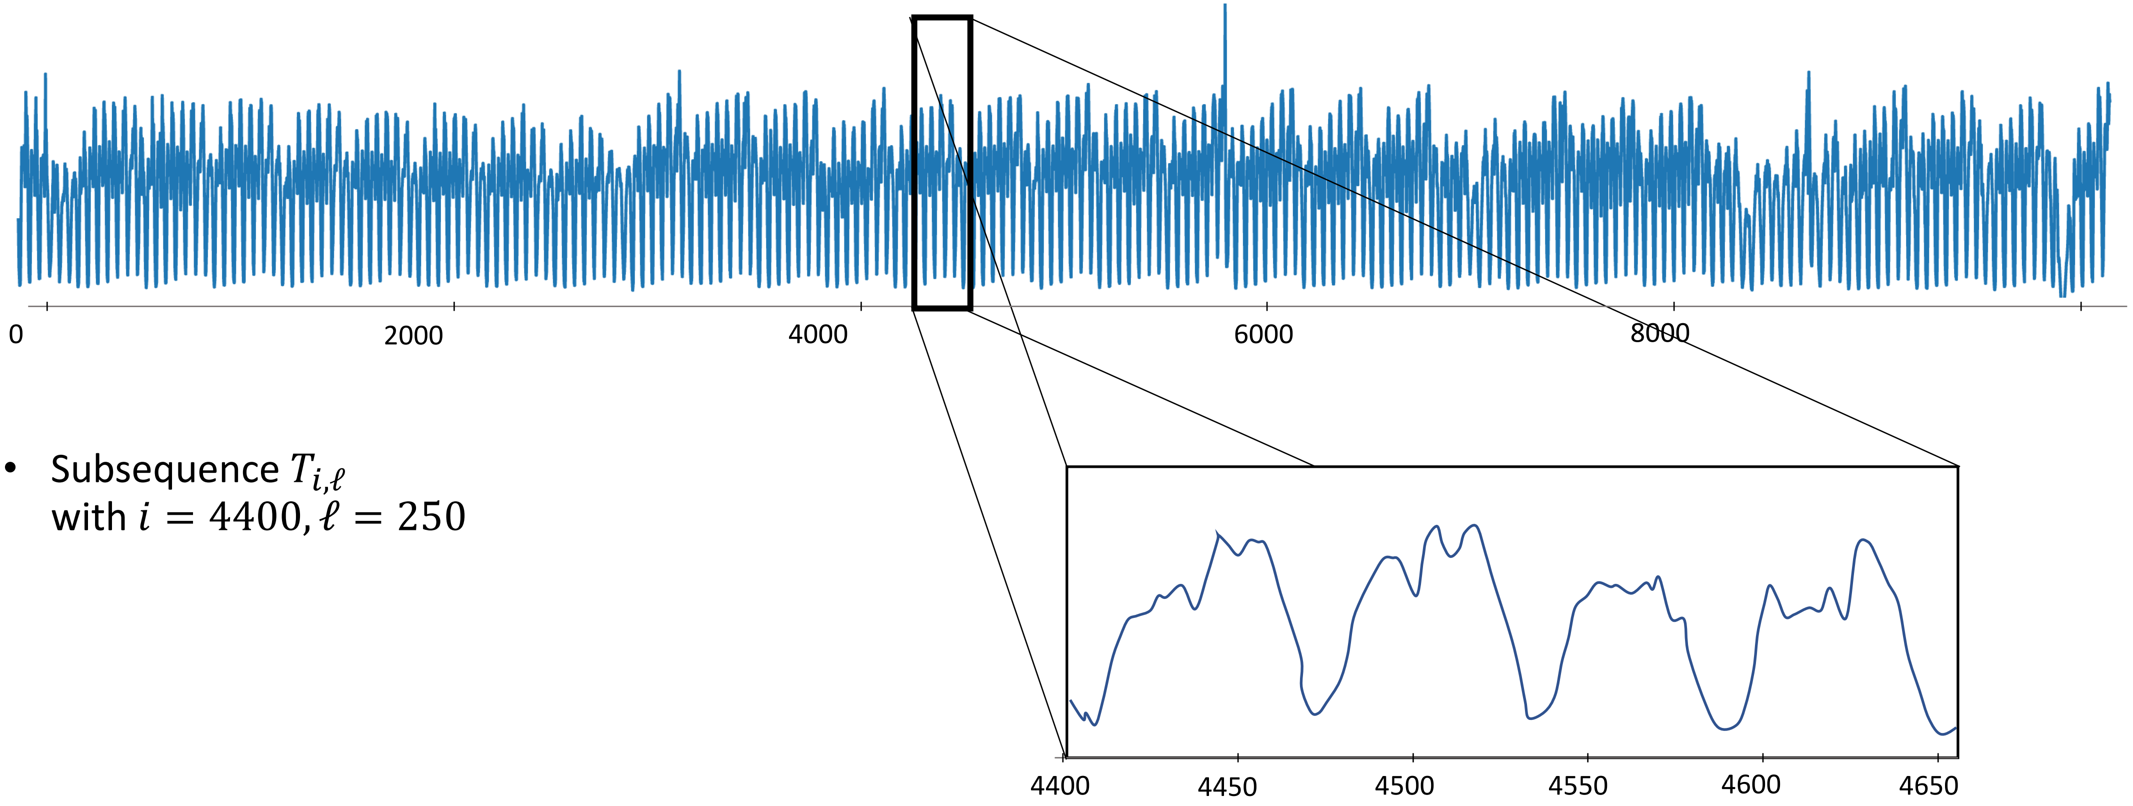

- Subsequence $T_{i,\ell}$
  with $i = 4400, \ell = 250$

Introduction: *Anomaly Detection in Time Series*

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

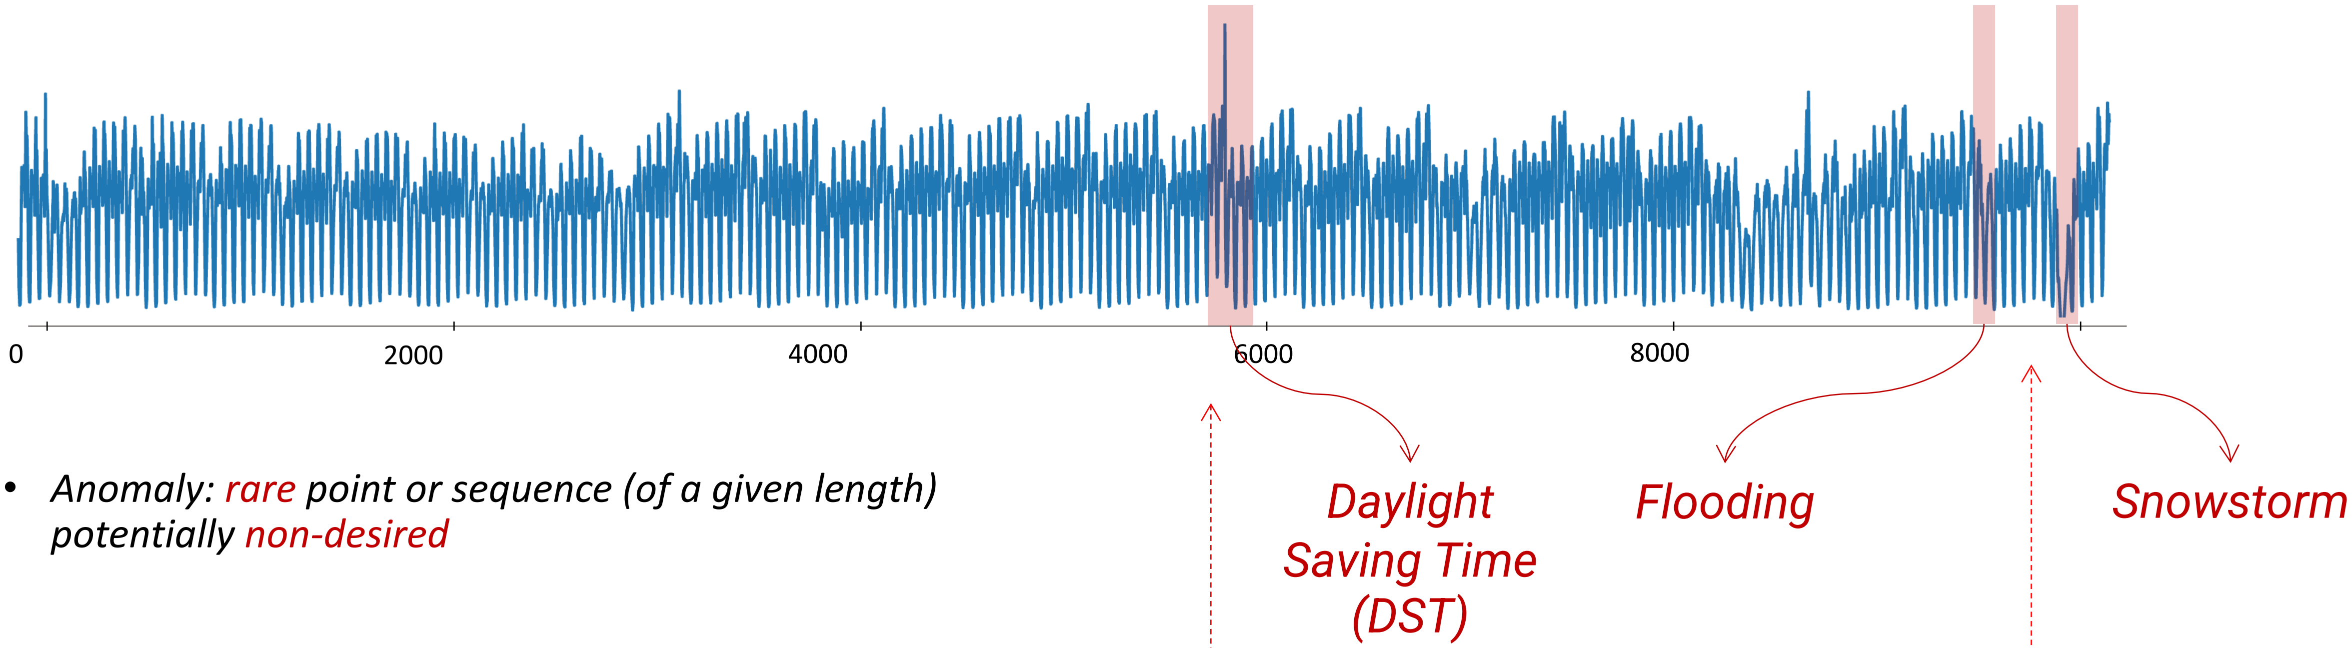

- *Anomaly: rare point or sequence (of a given length) potentially non-desired*

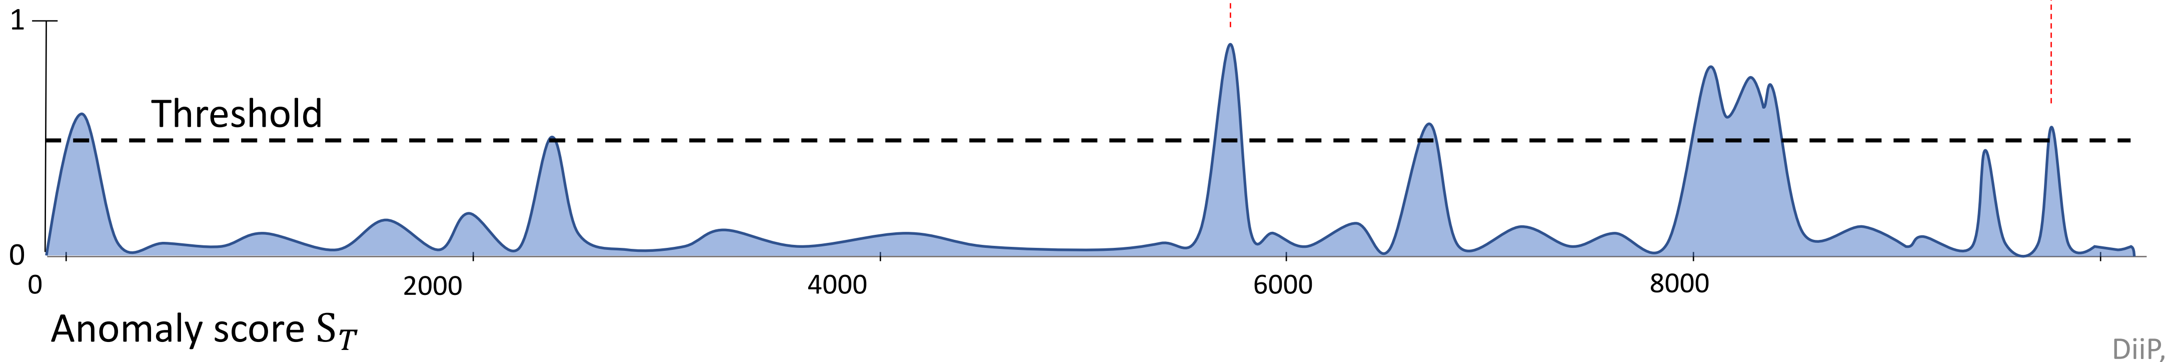

Introduction: *Outline*

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

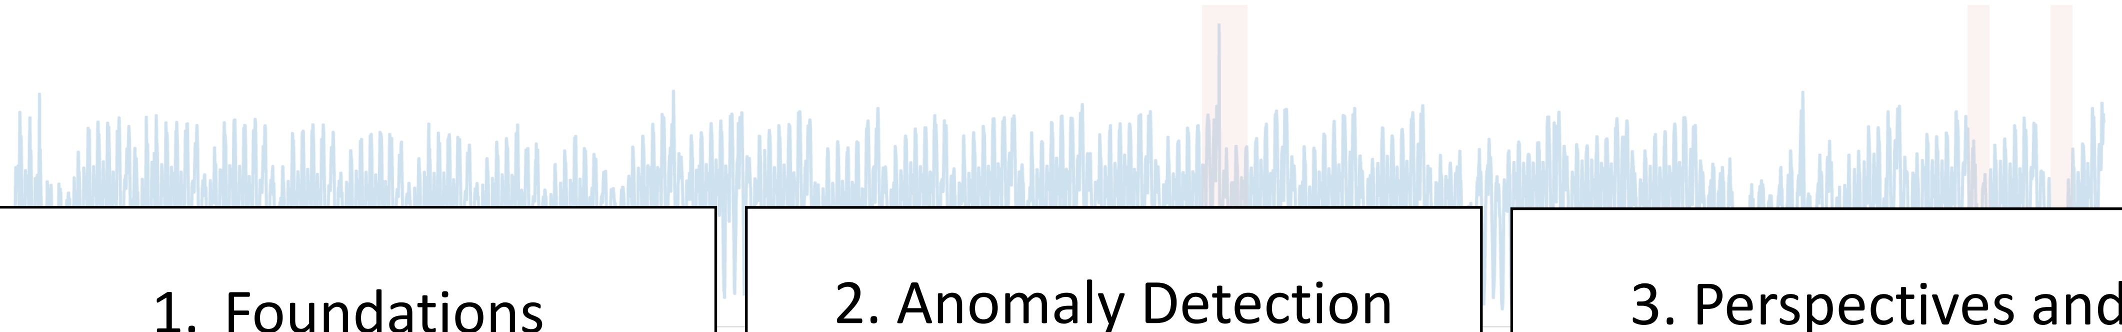

1. Foundations

1.1. Type of Time Series
1.2. Type of Anomalies

2. Anomaly Detection Methods

2.1. A Taxonomy of Methods
2.2. Existing Benchmarks

3. Perspectives and challenges

3.1. Time series labeling issues
3.2. Ensembling

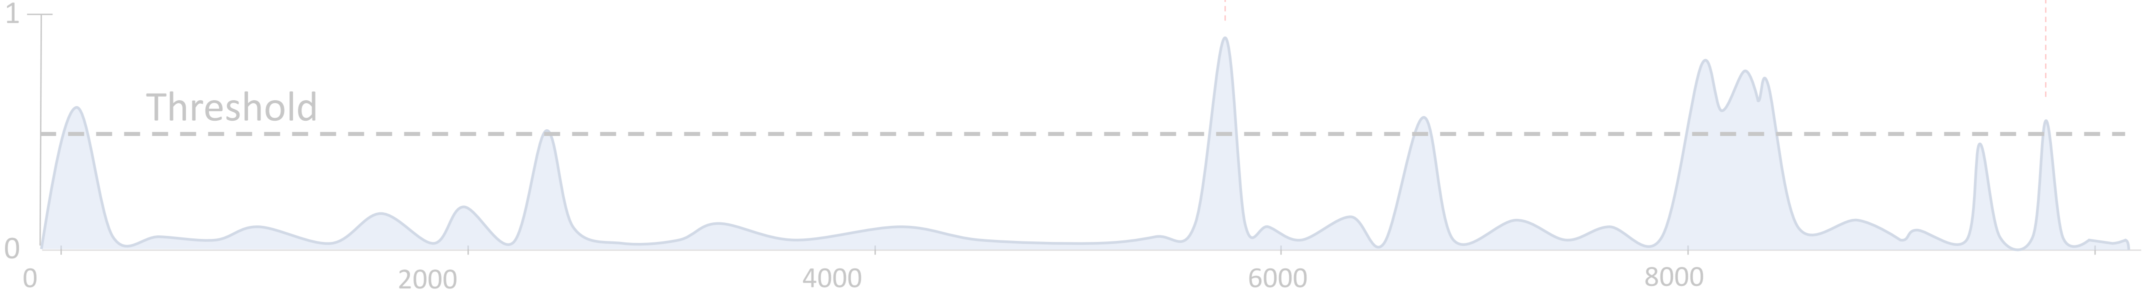

Anomaly score

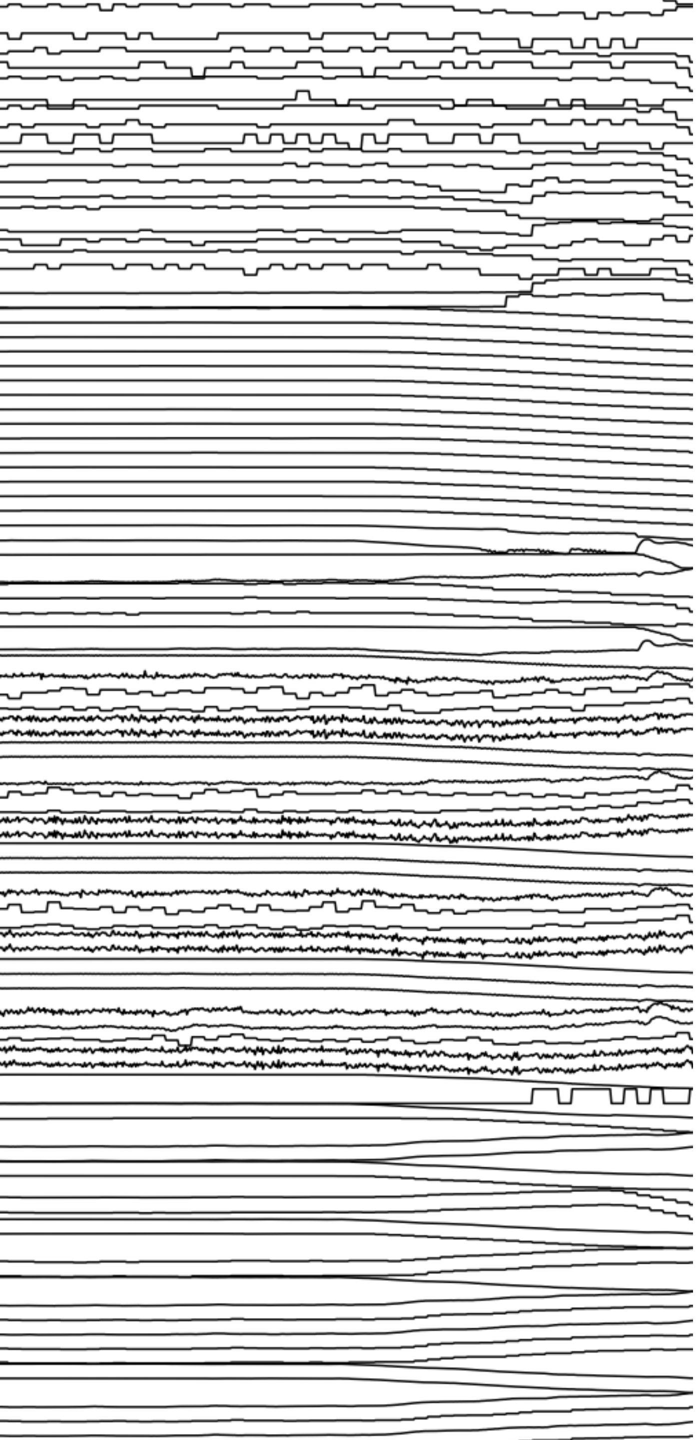

Foundations

Foundations: *Type of time series*

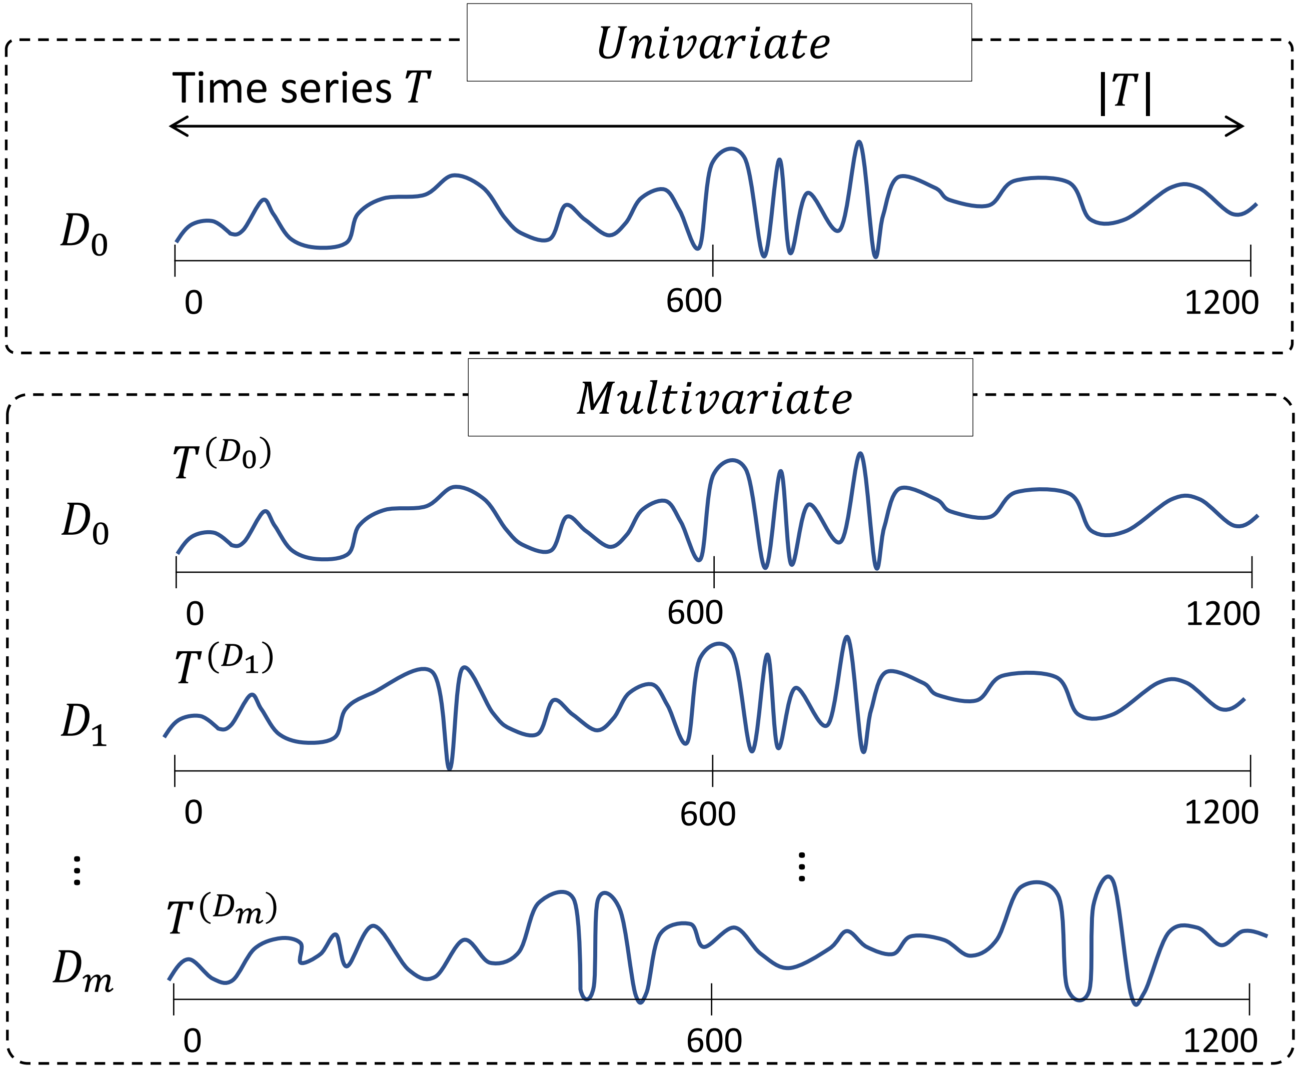

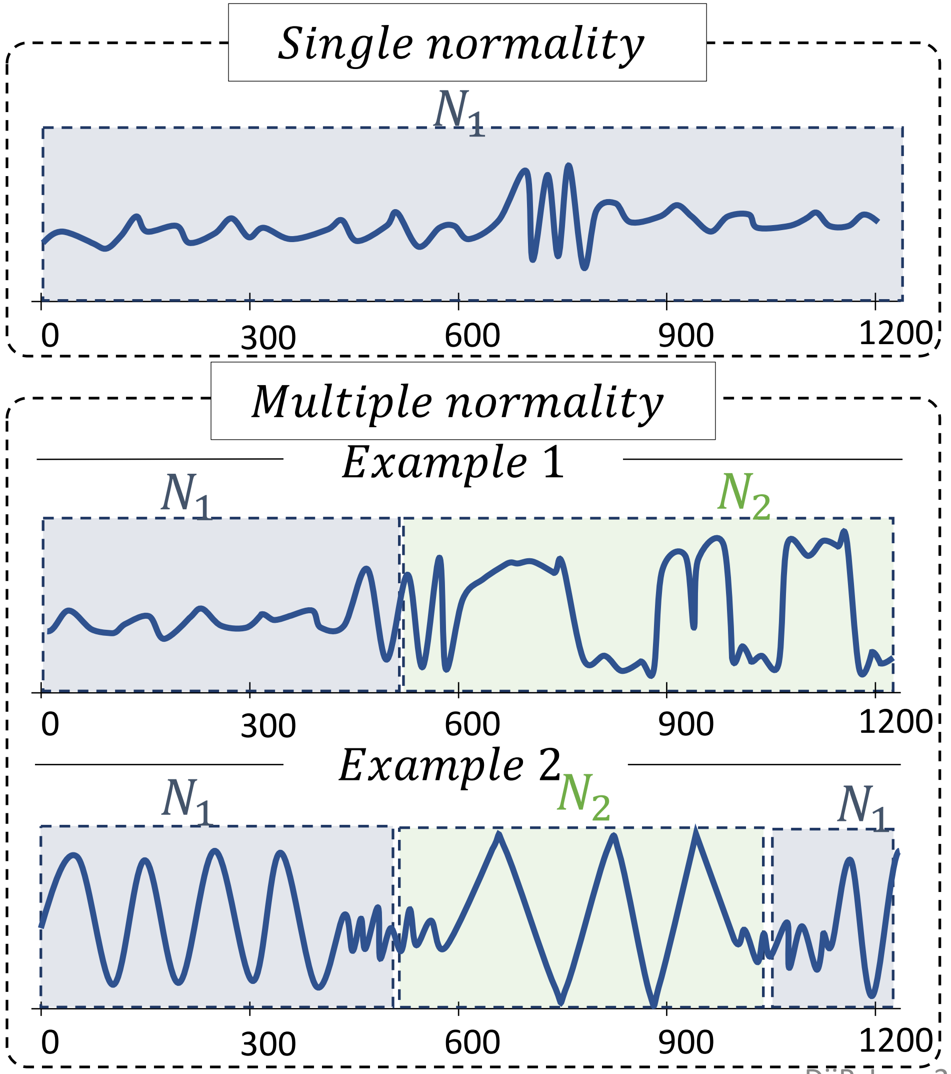

Foundations: *Type of anomalies*

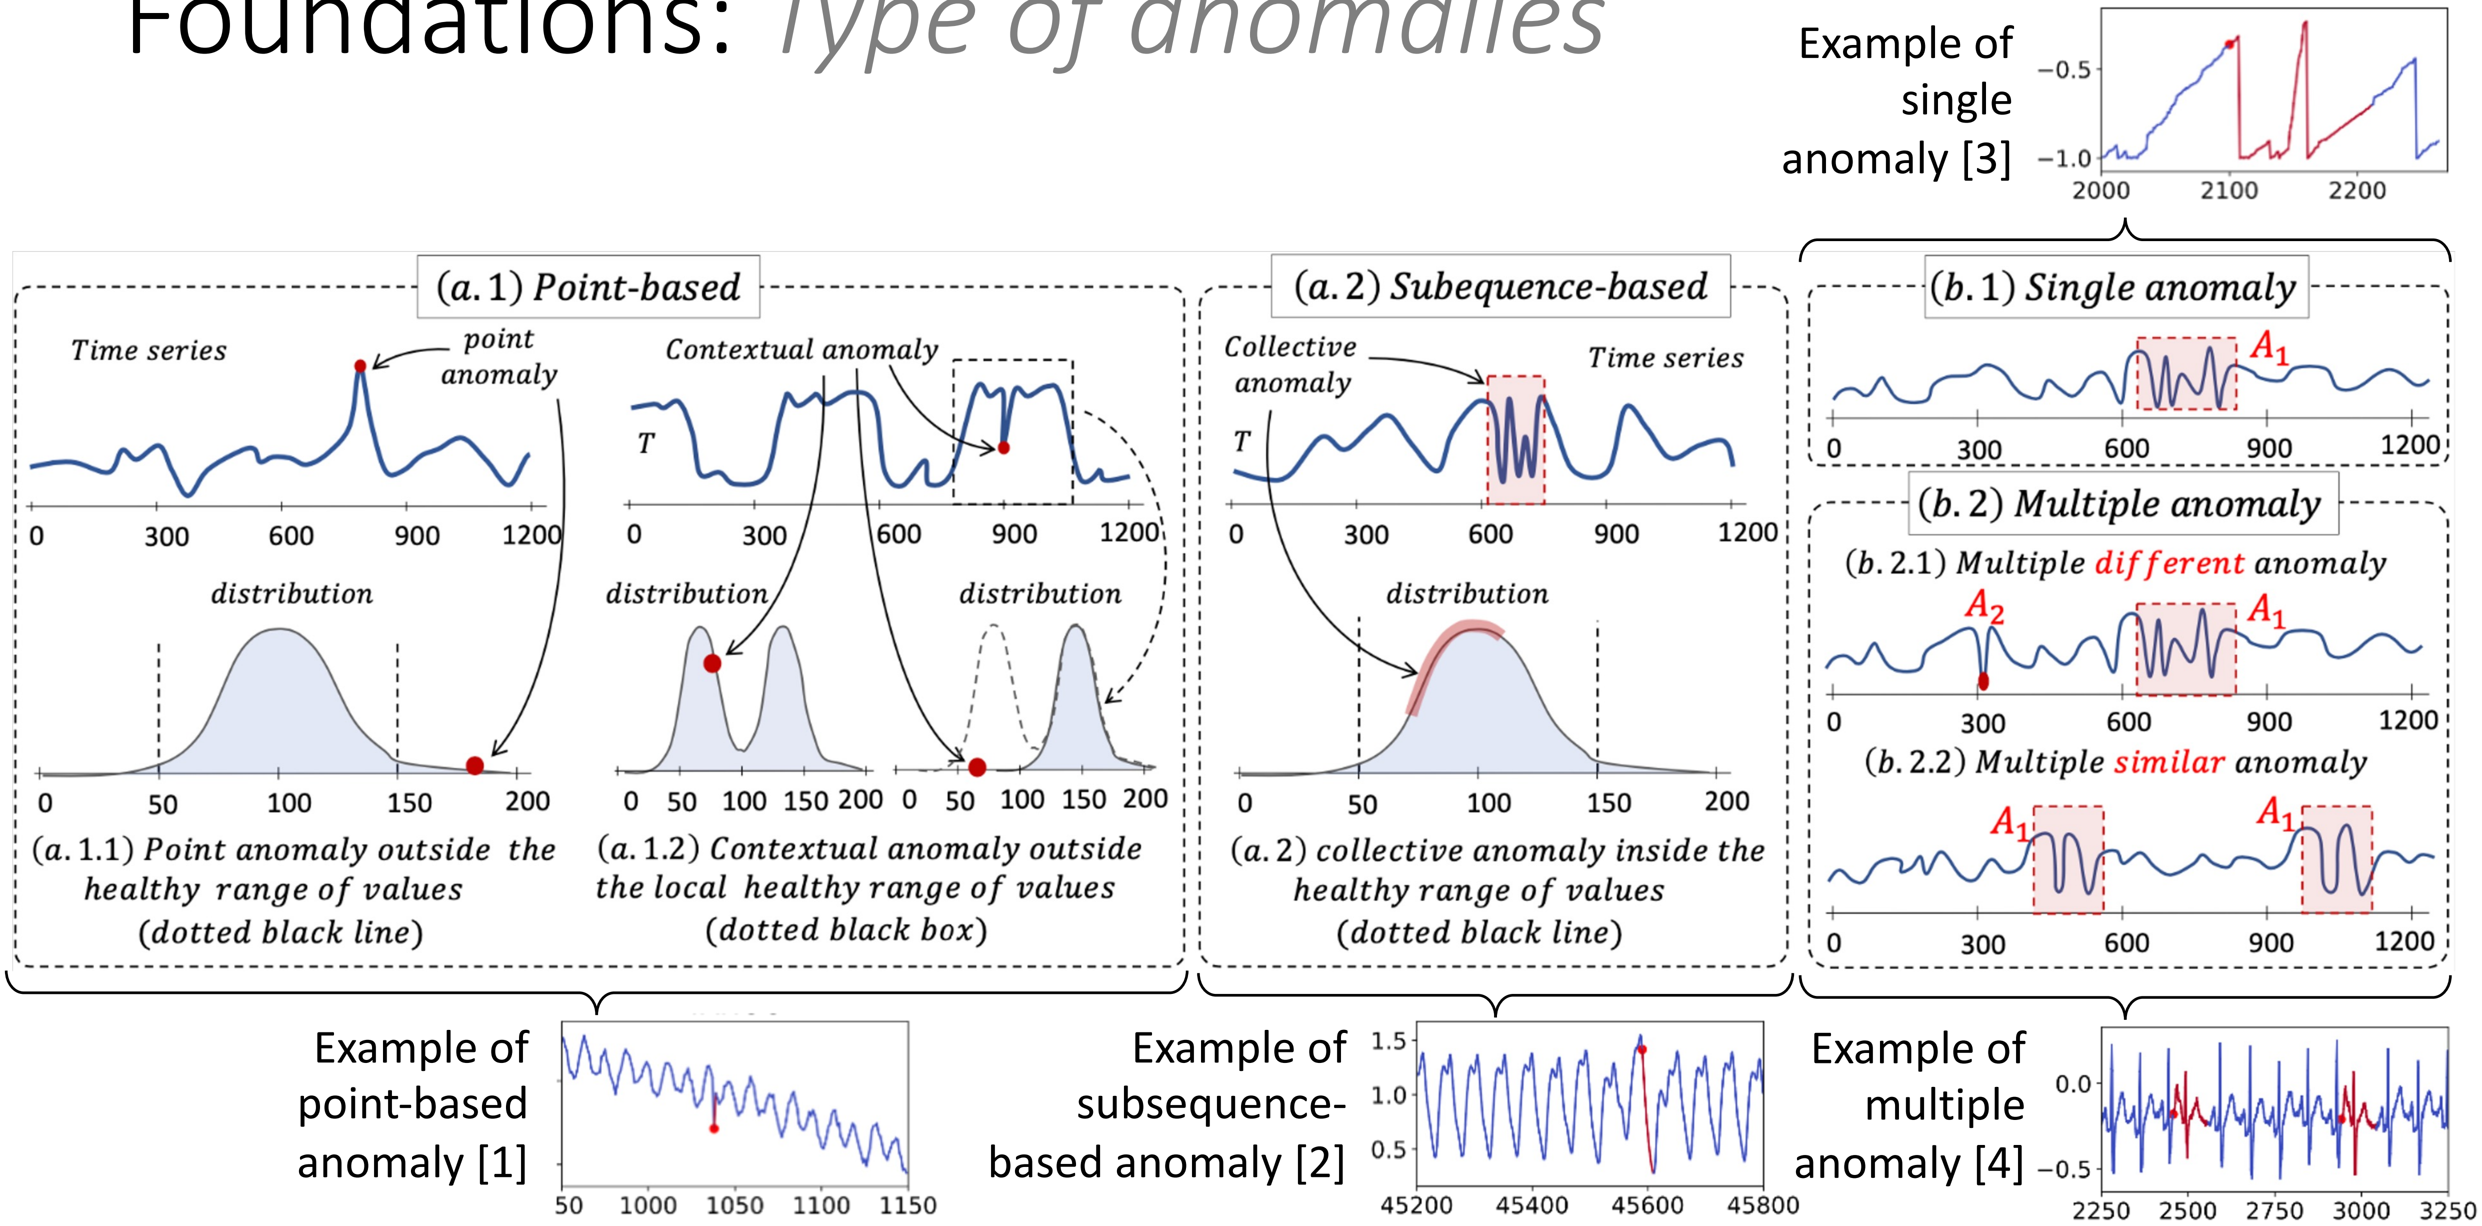

Foundations: *Type of anomalies*

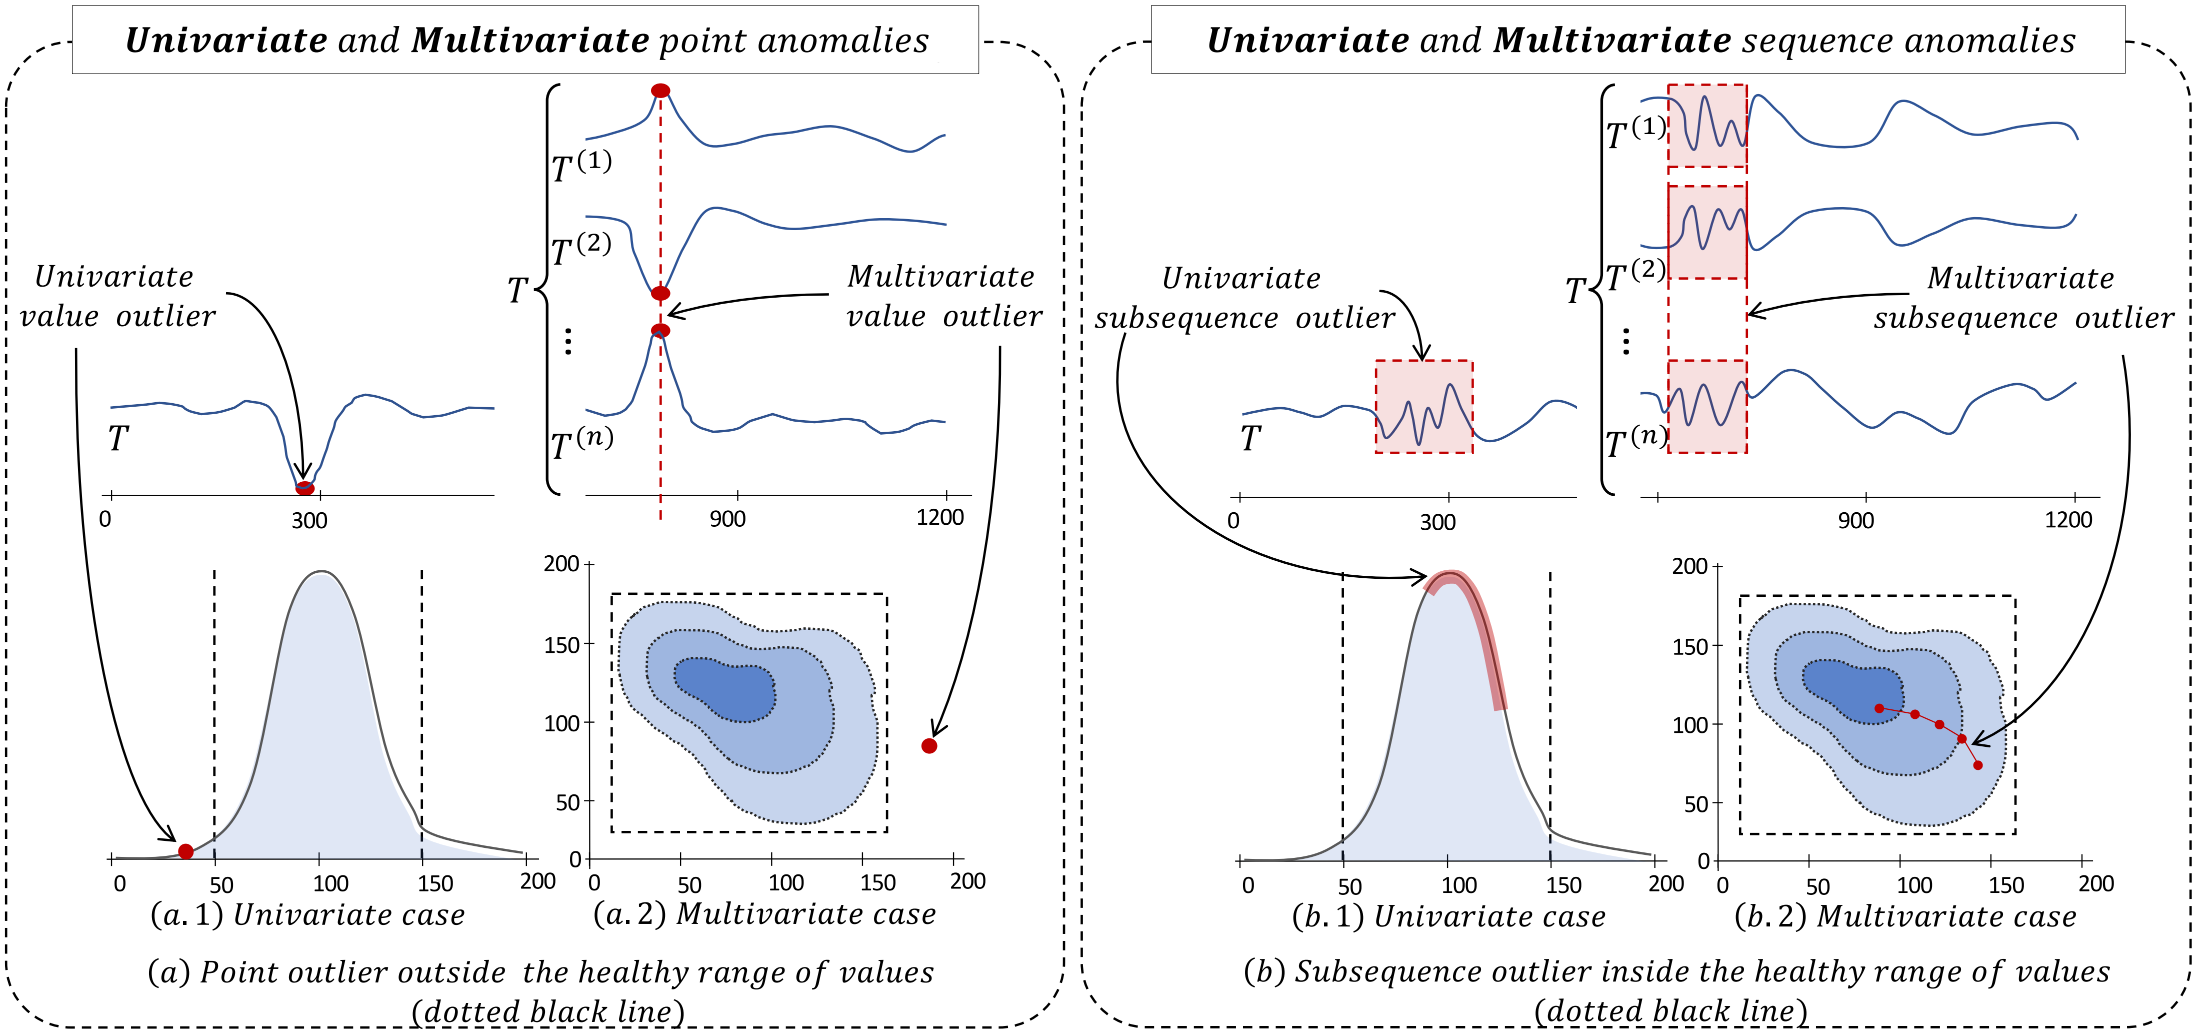

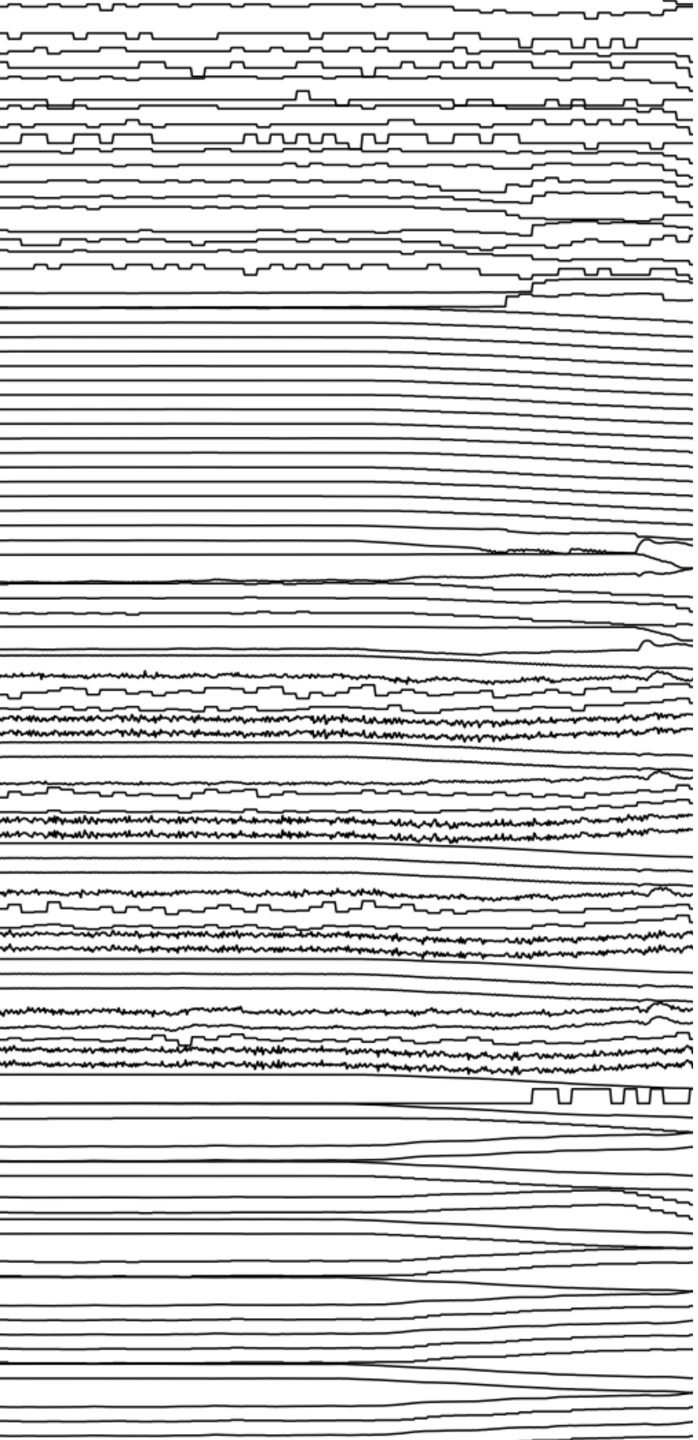

Anomaly Detection Methods

Anomaly Detection methods: *A taxonomy*

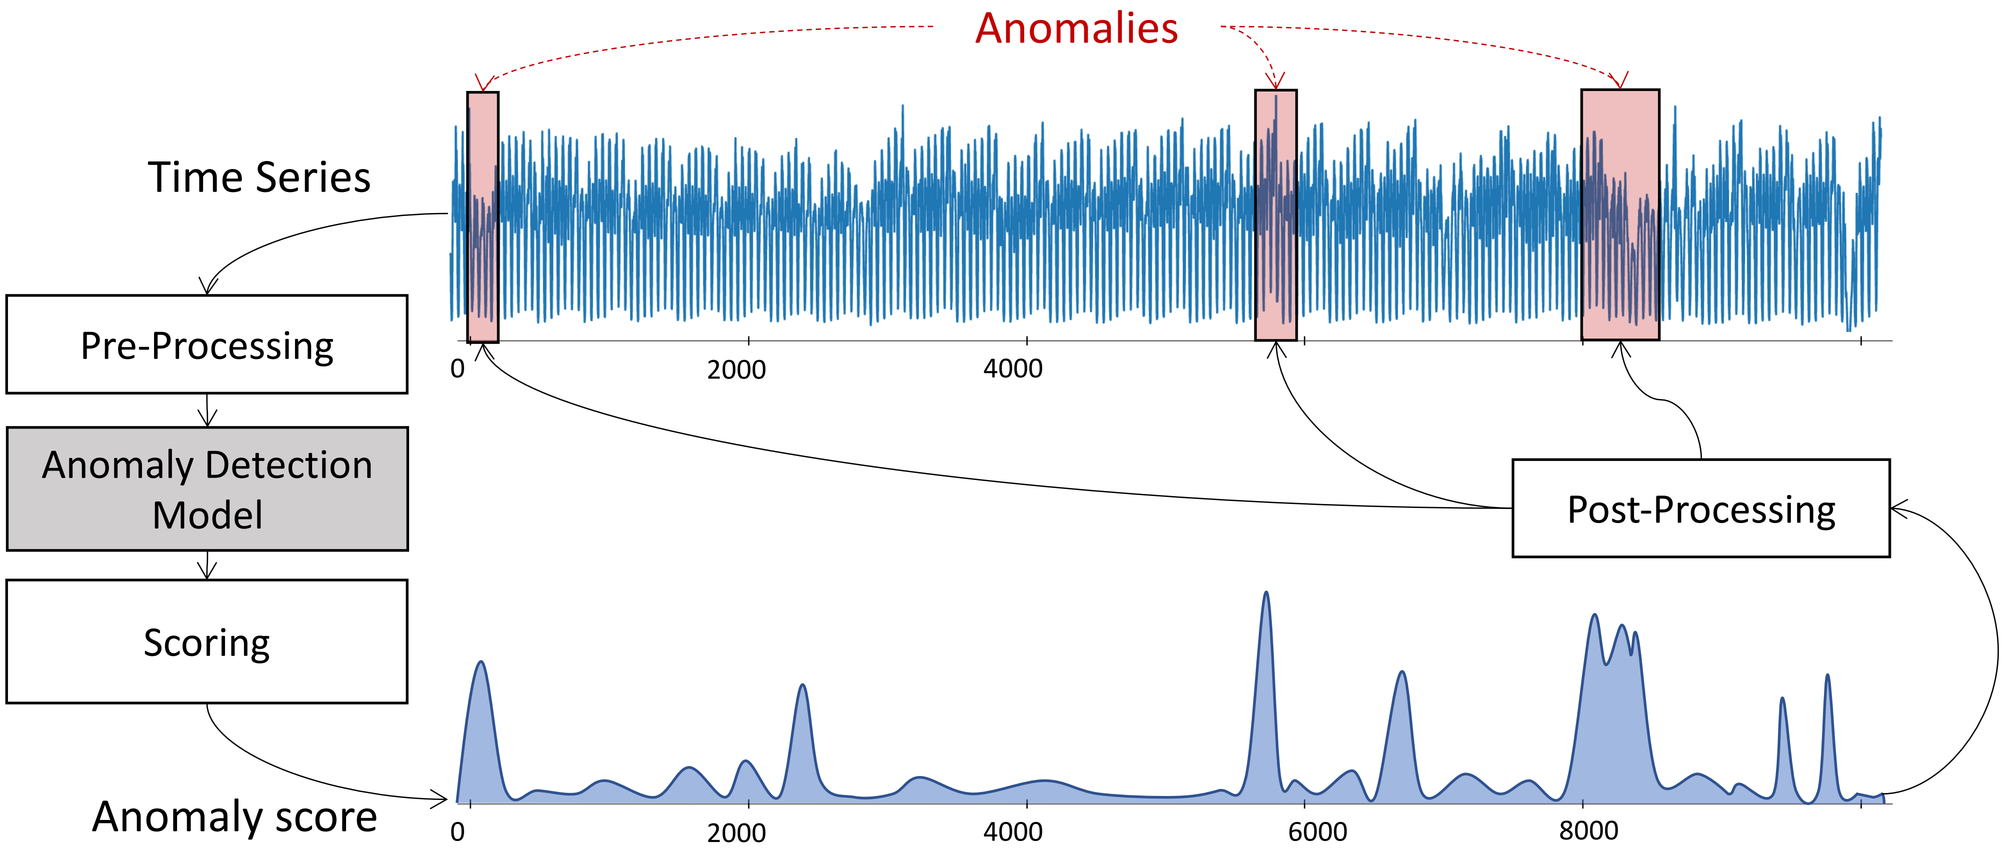

Anomaly Detection methods: *A taxonomy*

By domains [5] …

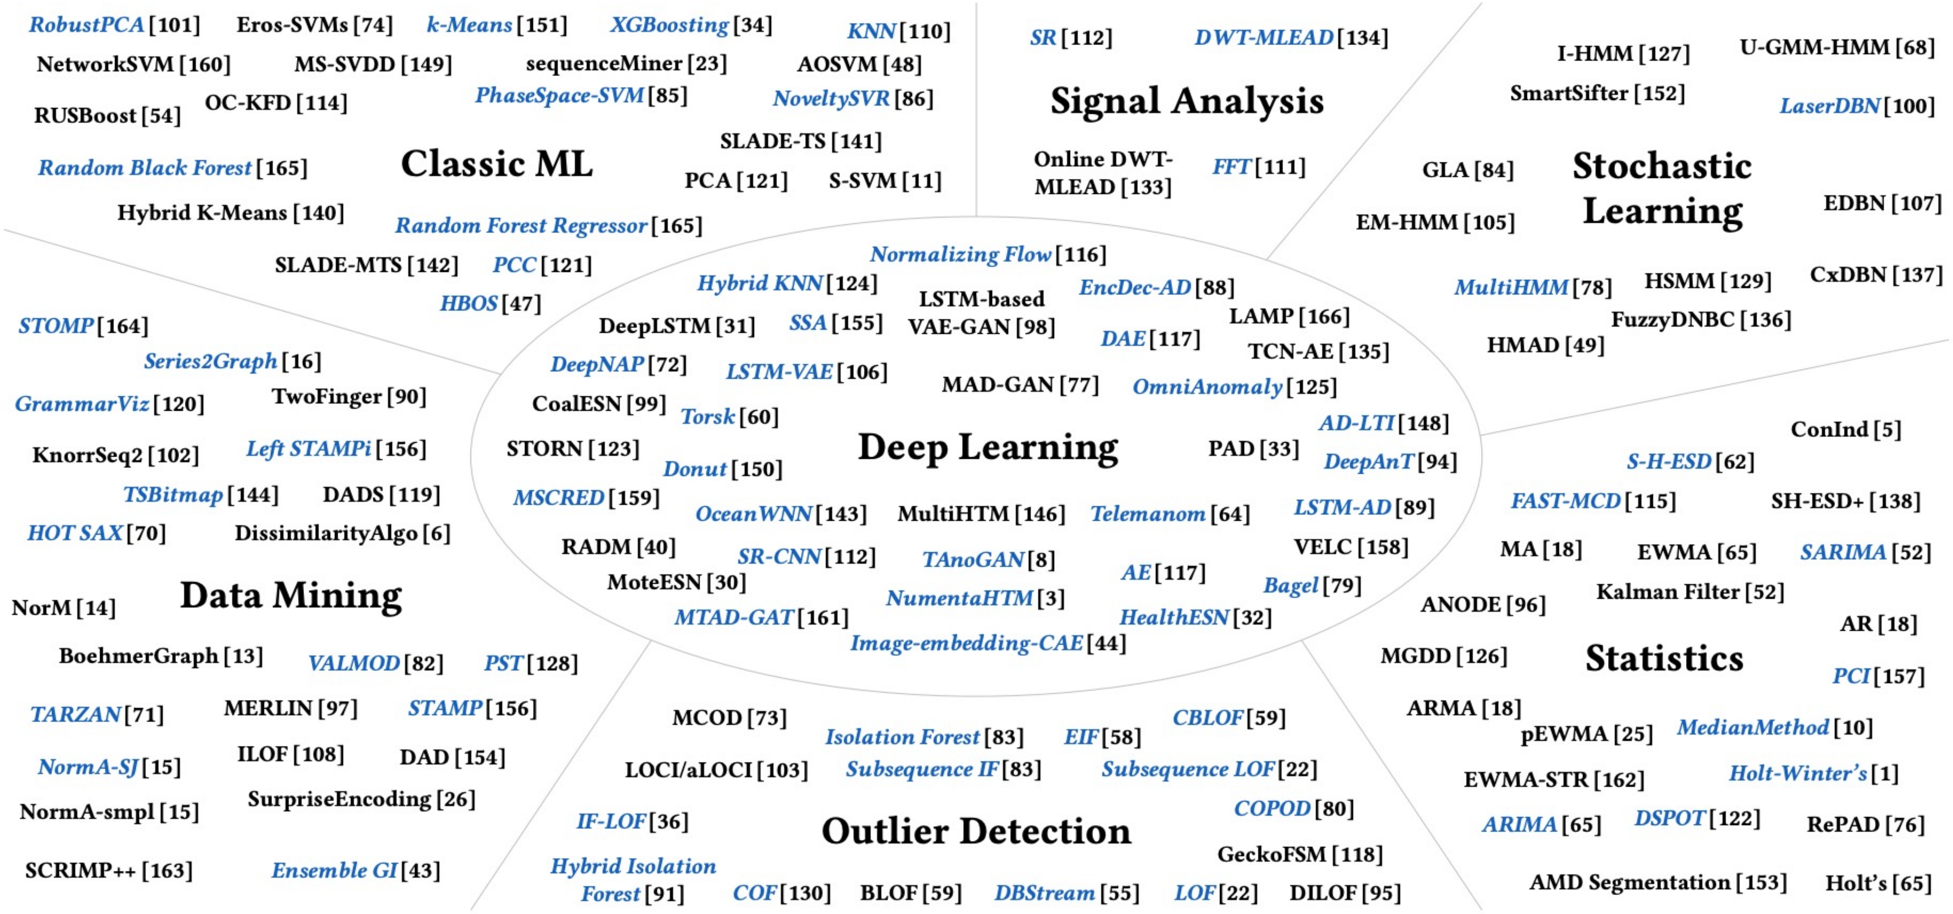

[5] Sebastian Schmidl, Phillip Wenig, and Thorsten Papenbrock. 2022. Anomaly detection in time series: a comprehensive evaluation. Proc. VLDB Endow. 15, 9 (May 2022), 1779–1797.

Anomaly Detection methods: *A taxonomy*

By inputs…

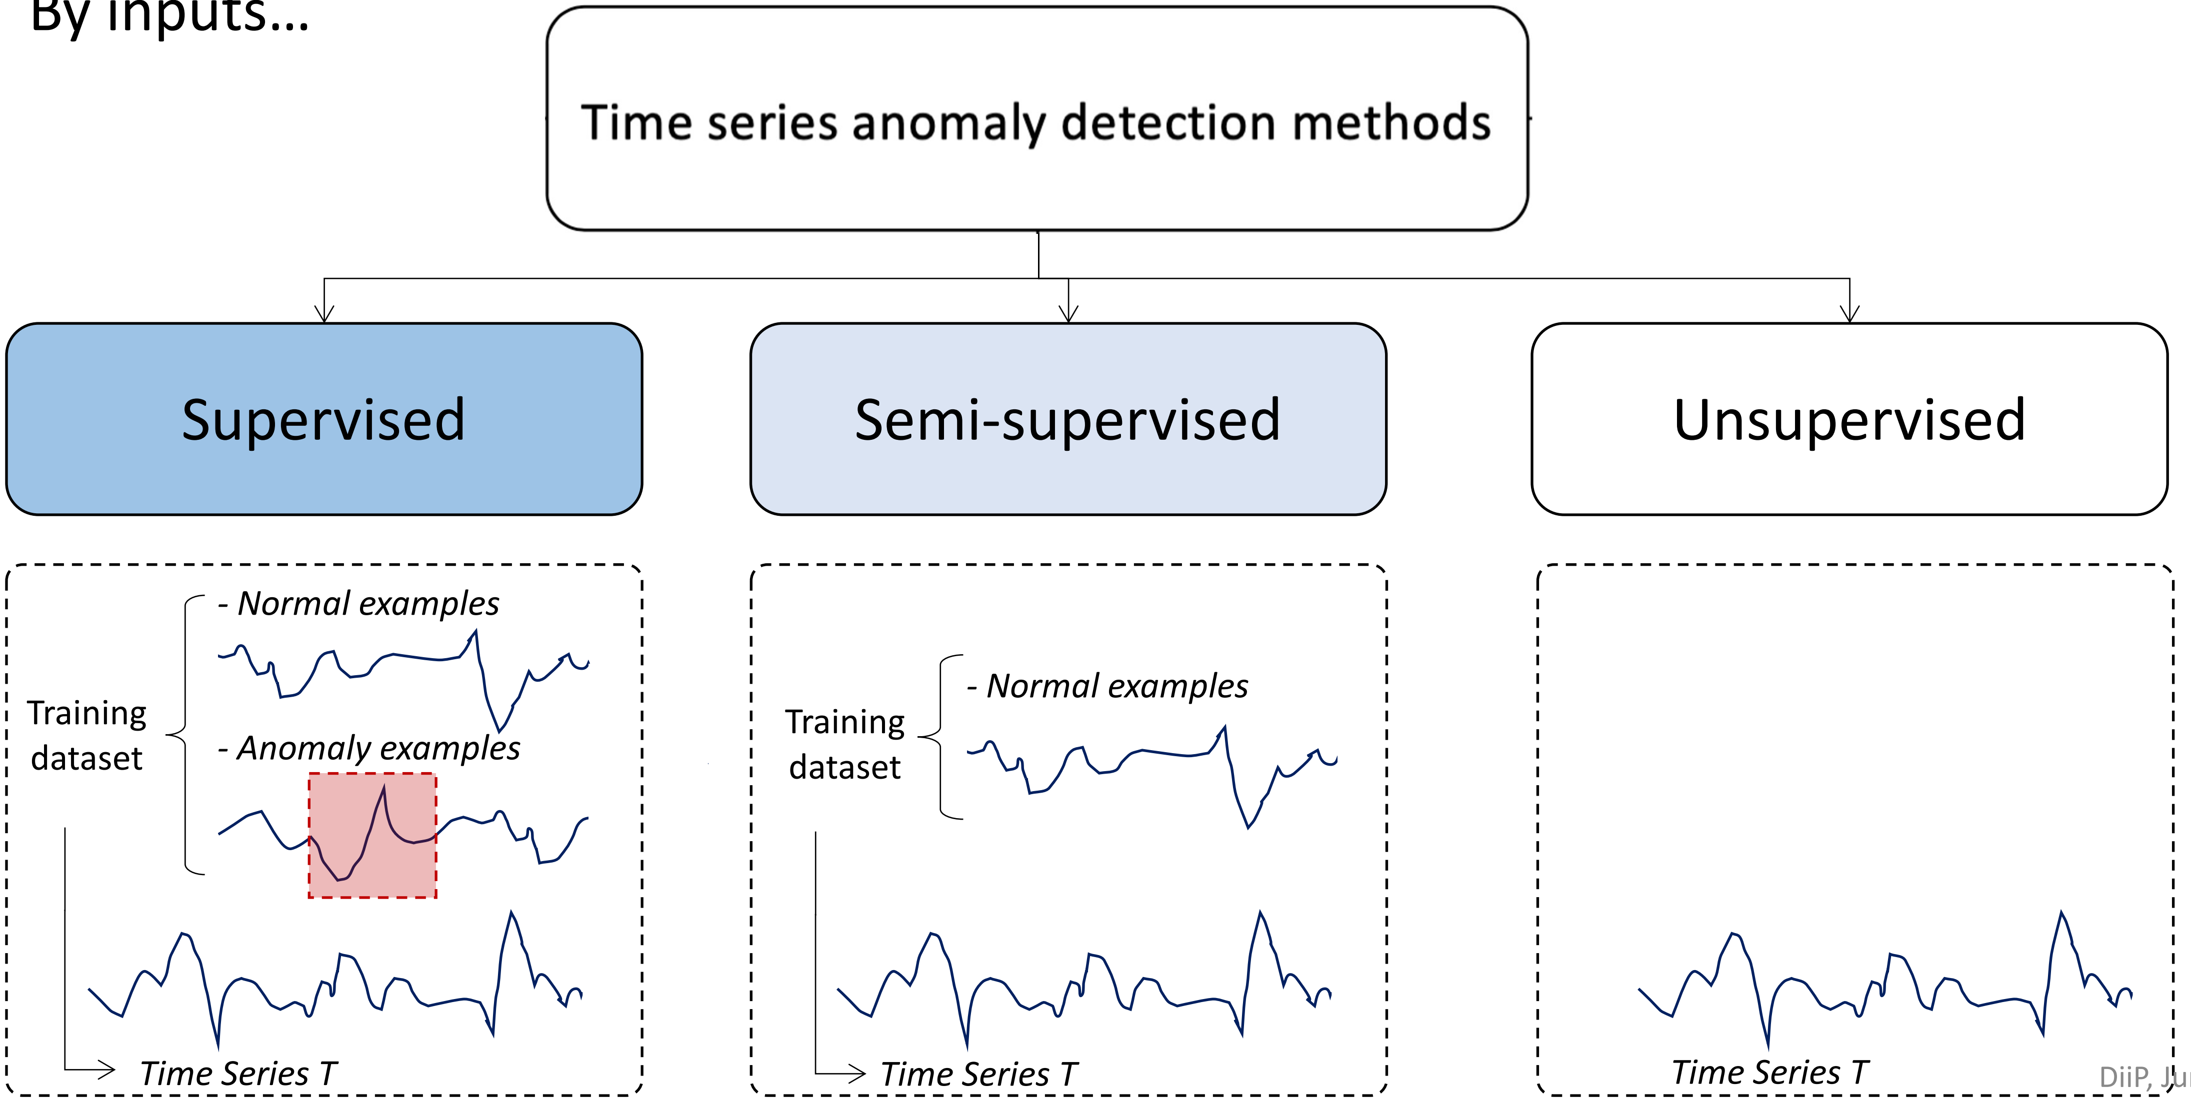

Anomaly Dete

By inputs…

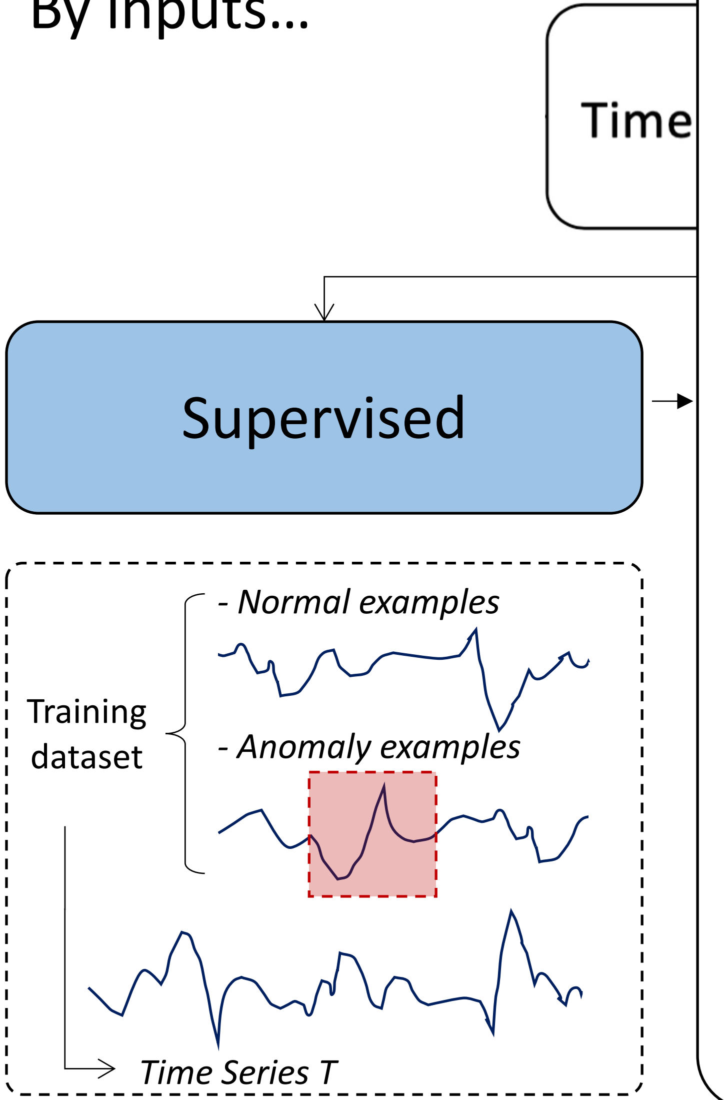

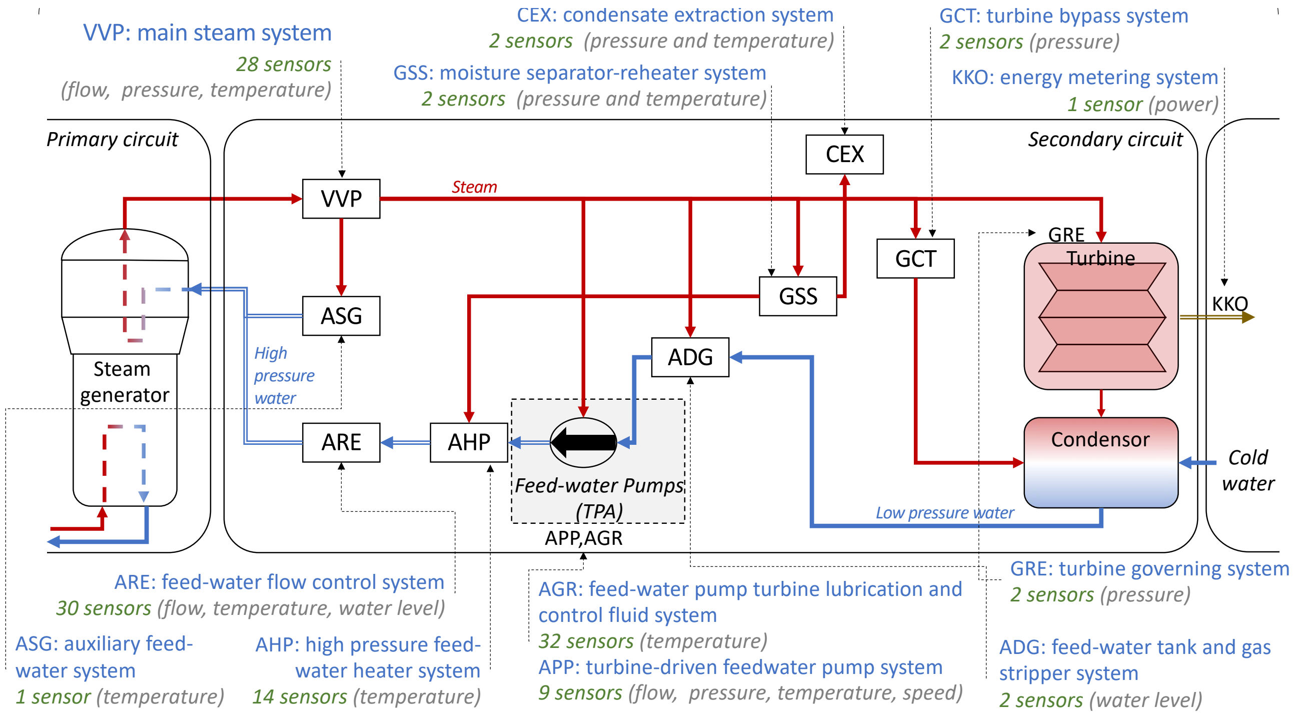

Anomaly Dete

By inputs…

Time

Supervised

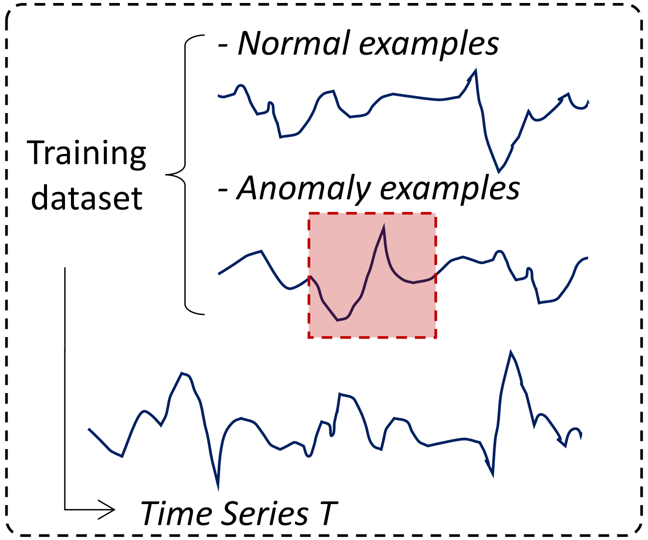

Supervised anomaly detection (e.g., classification)

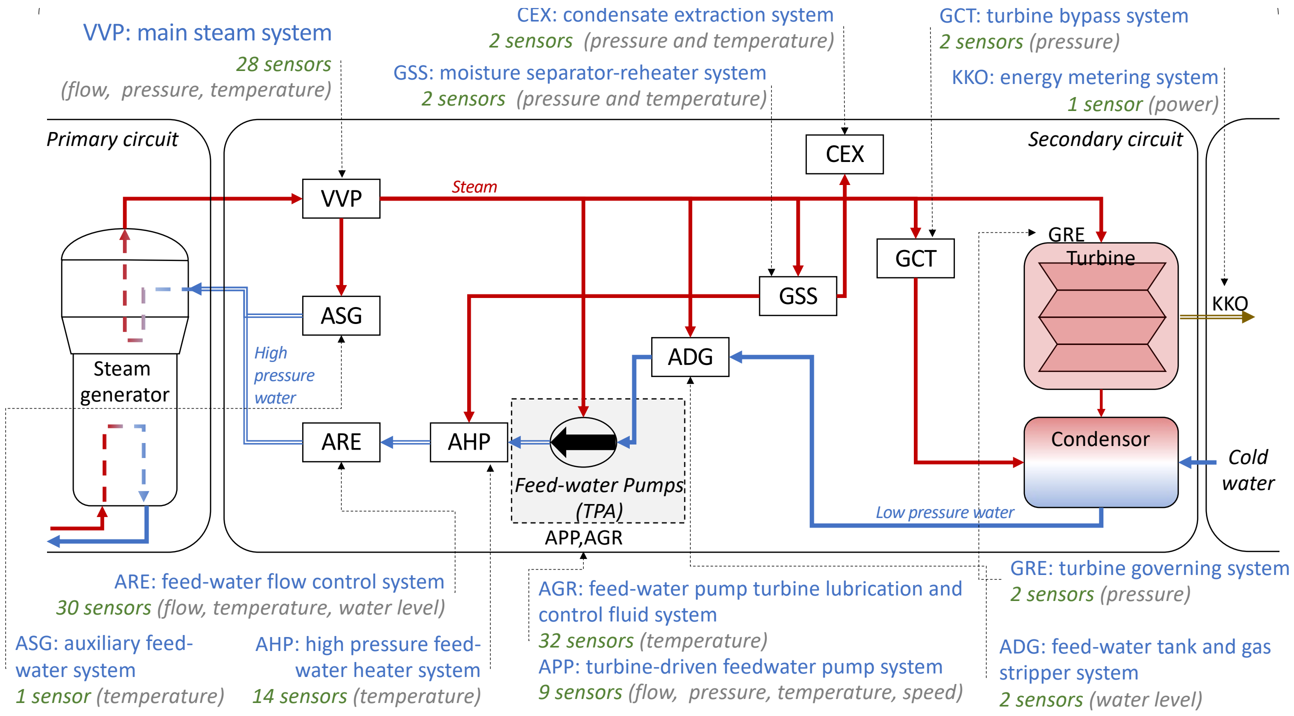

Class 1: Time series without any vibrations

Class 2: Time series with a vibrations

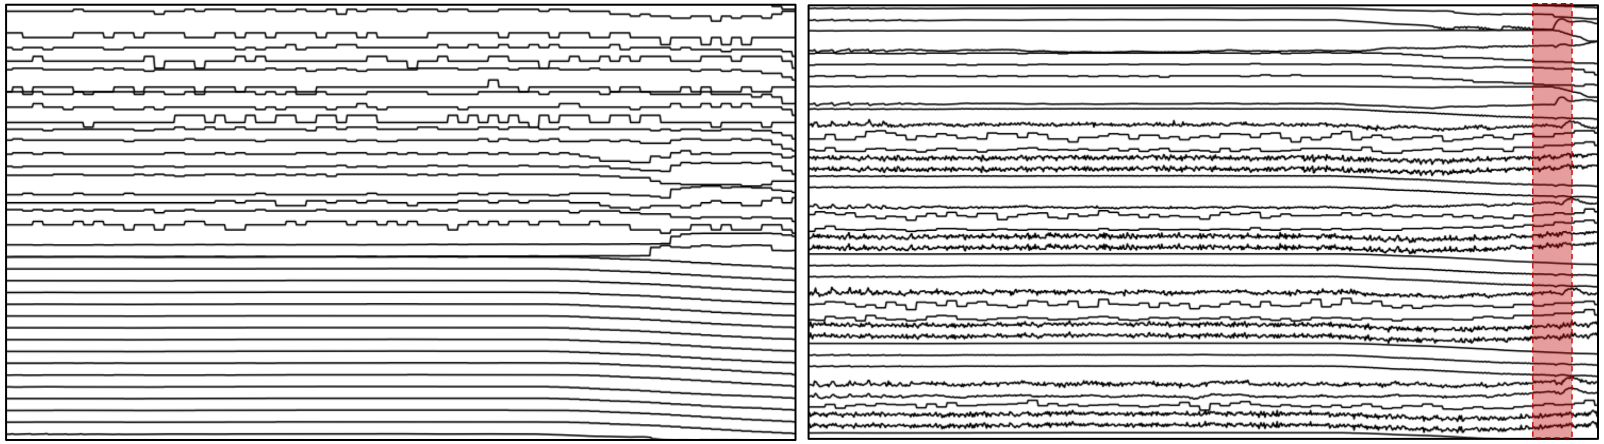

Vibration

Anomaly Dete

By inputs…

Time

Supervised anomaly detection (e.g., classification)

Explanation of the detection

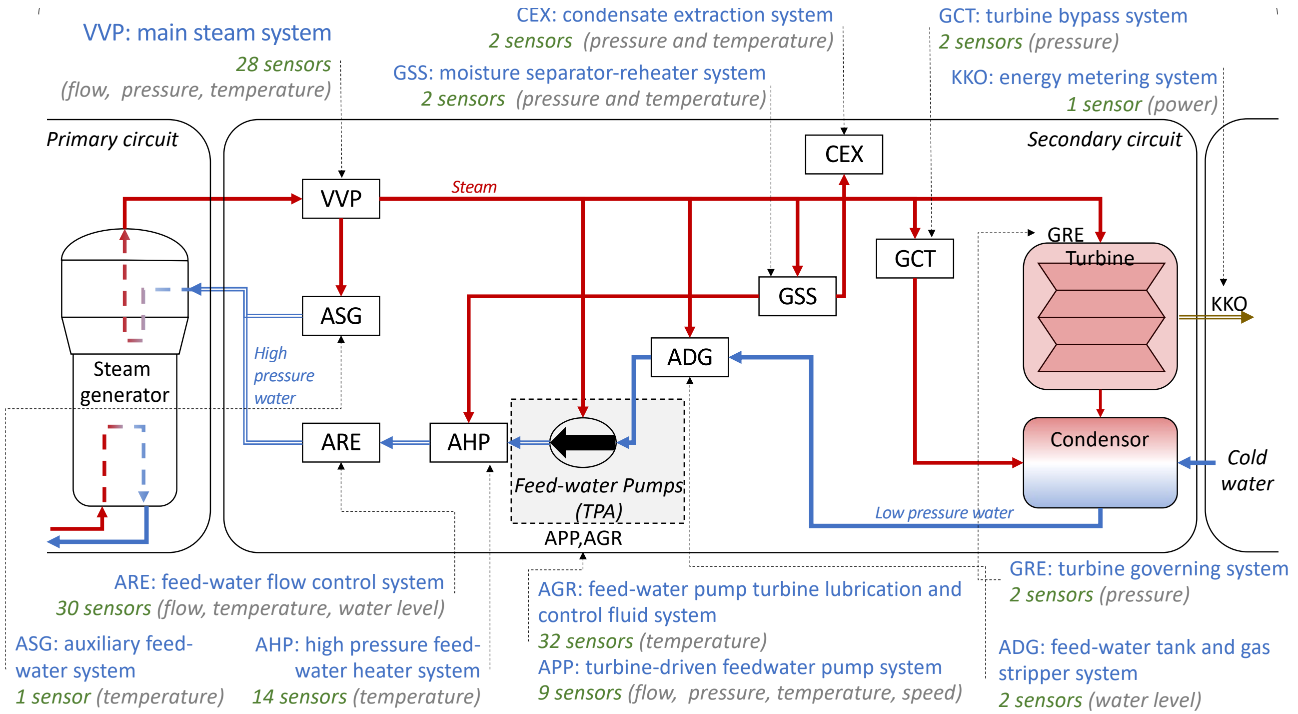

Class 1: Time series without any vibrations

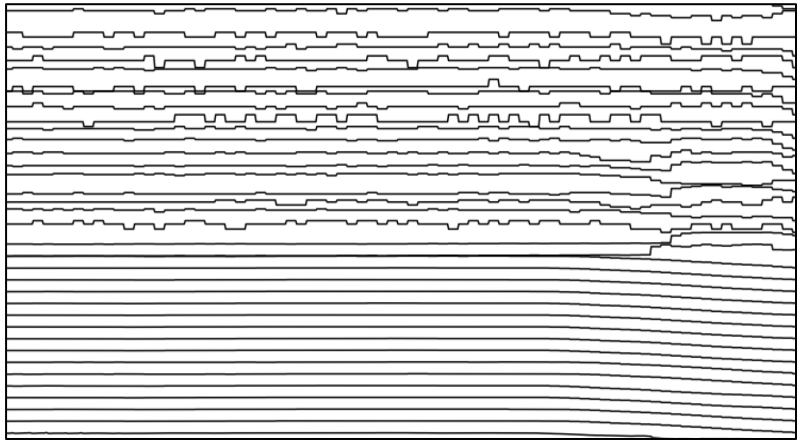

Class 2: Time series with a vibrations

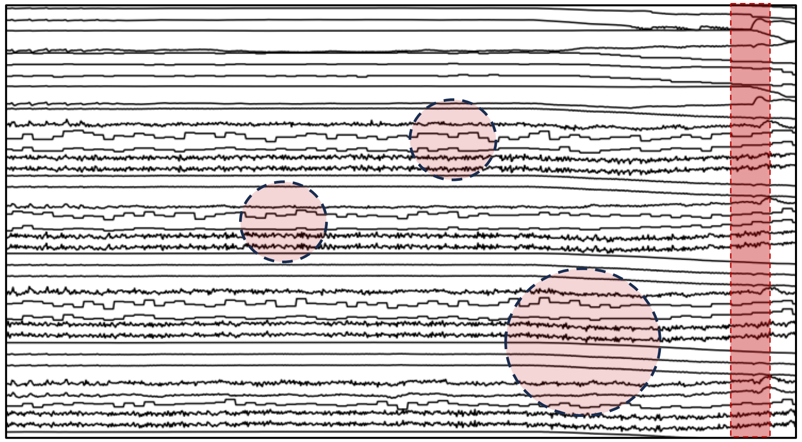

Vibration

Supervised

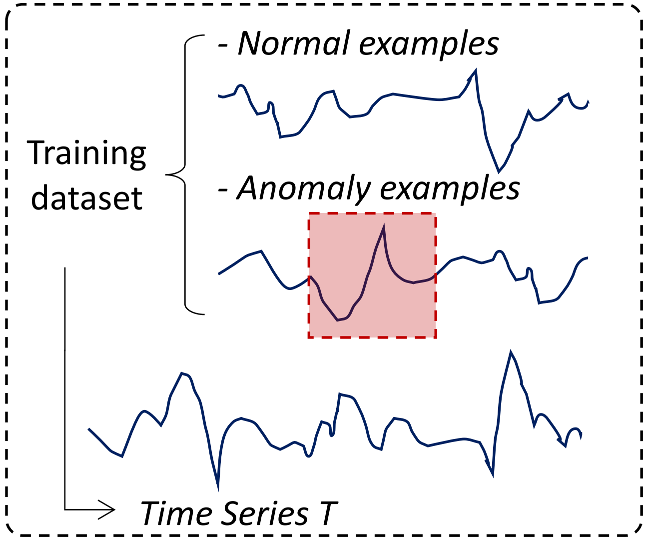

Anomaly Dete

By inputs…

Time

Supervised

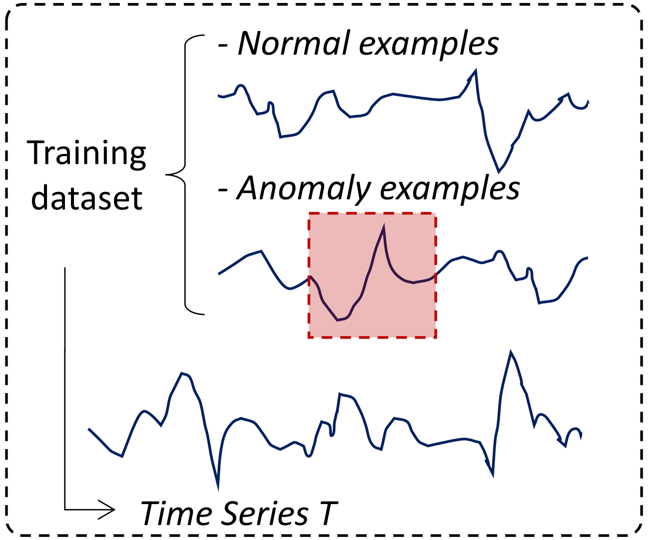

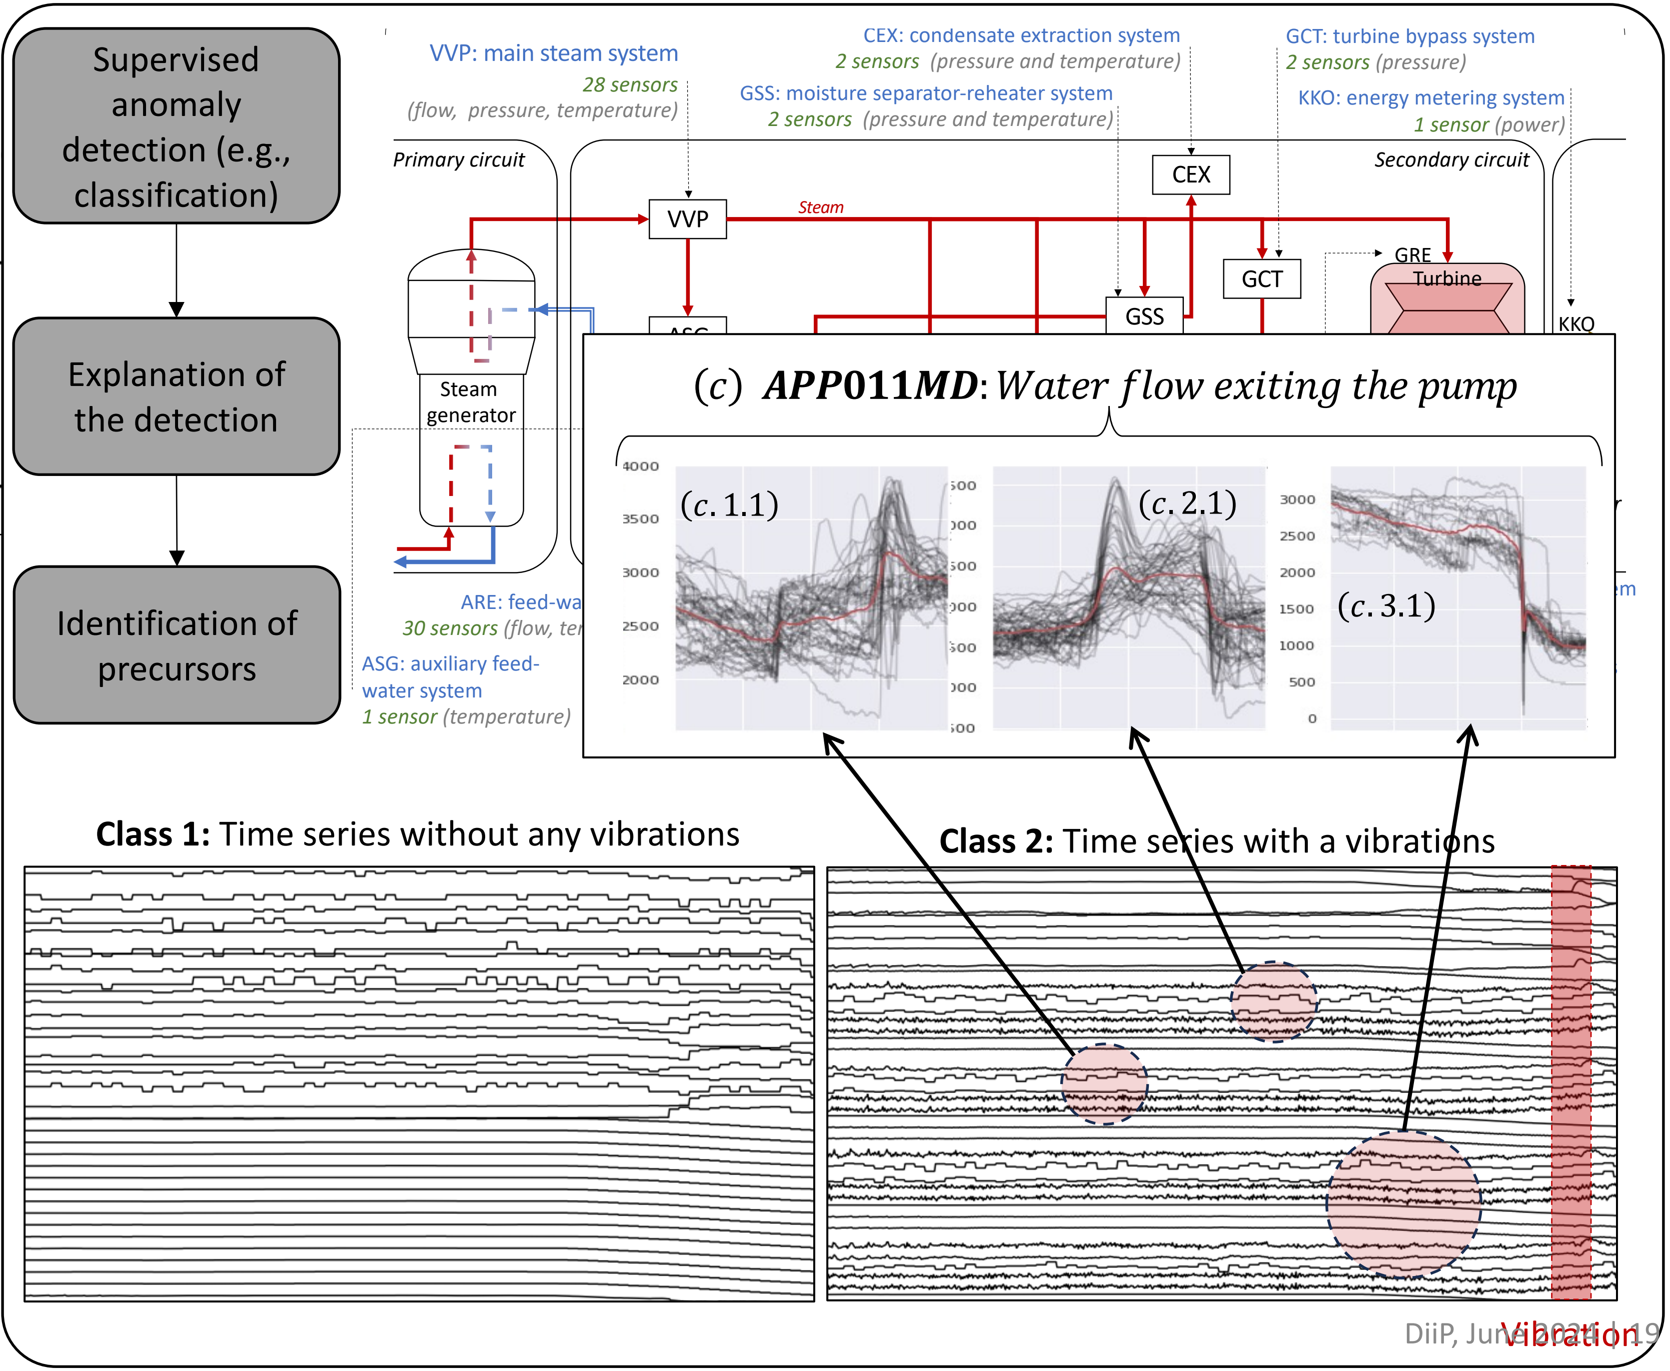

Anomaly Dete

By inputs…

Time

Supervised anomaly detection (e.g., classification)

Explanation the detectio

Identification precursors

Supervised

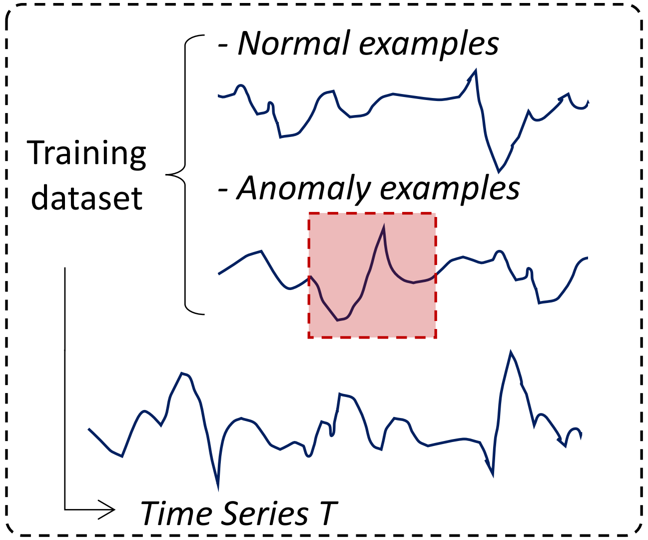

Class 1: Ti

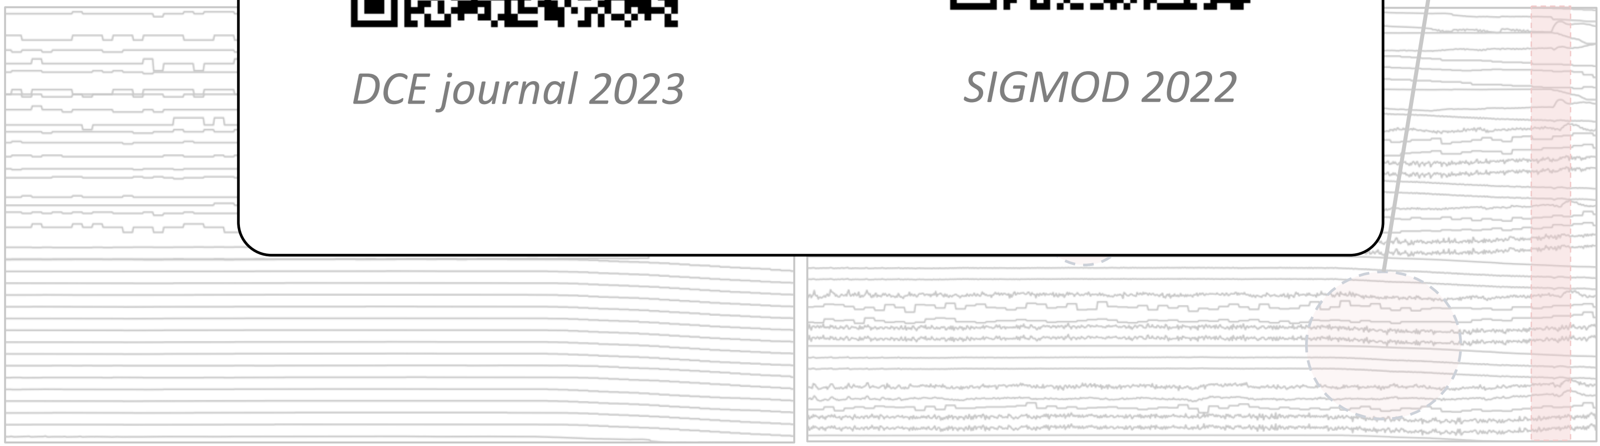

More info :

On the use case

DCE journal 2023

On the method

SIGMOD 2022

Anomaly Detection methods: *A taxonomy*

By methods…

Time series anomaly detection methods

- Distance-based
  - Proximity-based
  - Clustering-based
  - Discord-based

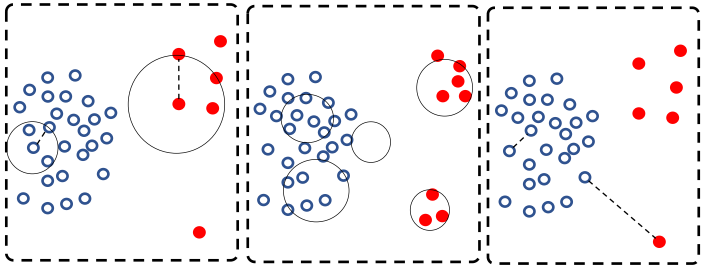

E.g., LOF

E.g., NormA, SAND

E.g., MP, DAMP

Anomaly Detection methods: *A taxonomy*

By methods…

- Time series anomaly detection methods
  - Distance-based
    - Proximity-based
    - Clustering-based
    - Discord-based
  - Density-based
    - Distribution-based
    - Graph-based
    - Tree-based
    - Encoding-based

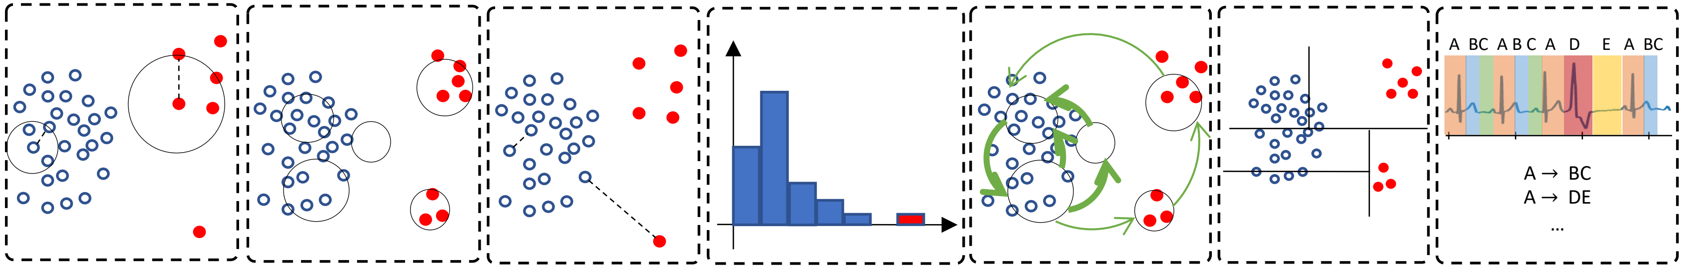

Proximity-based	Clustering-based	Discord-based	Distribution-based	Graph-based	Tree-based	Encoding-based
E.g., LOF	*E.g.,* NormA SAND	*E.g.,* MP DAMP	*E.g.,* HOBS OCSVM	*E.g.,* Series2Graph	*E.g.,* Isolation-Forest	*E.g.,* GrammarViz POLY, PCA

Anomaly Detection methods: *A taxonomy*

By methods…

- **Time series anomaly detection methods**
  - **Distance-based**
    - Proximity-based — *E.g.,* LOF
    - Clustering-based — *E.g.,* NormA, SAND
    - Discord-based — *E.g.,* MP, DAMP
  - **Density-based**
    - Distribution-based — *E.g.,* HOBS, OCSVM
    - Graph-based — *E.g.,* Series2Graph
    - Tree-based — *E.g.,* Isolation-Forest
    - Encoding-based — *E.g.,* GrammarViz, POLY, PCA
  - **Prediction-based**
    - Forecasting-based — *E.g.,* LSTM,CNN
    - Reconstruction-based — *E.g.,* AutoEncoder

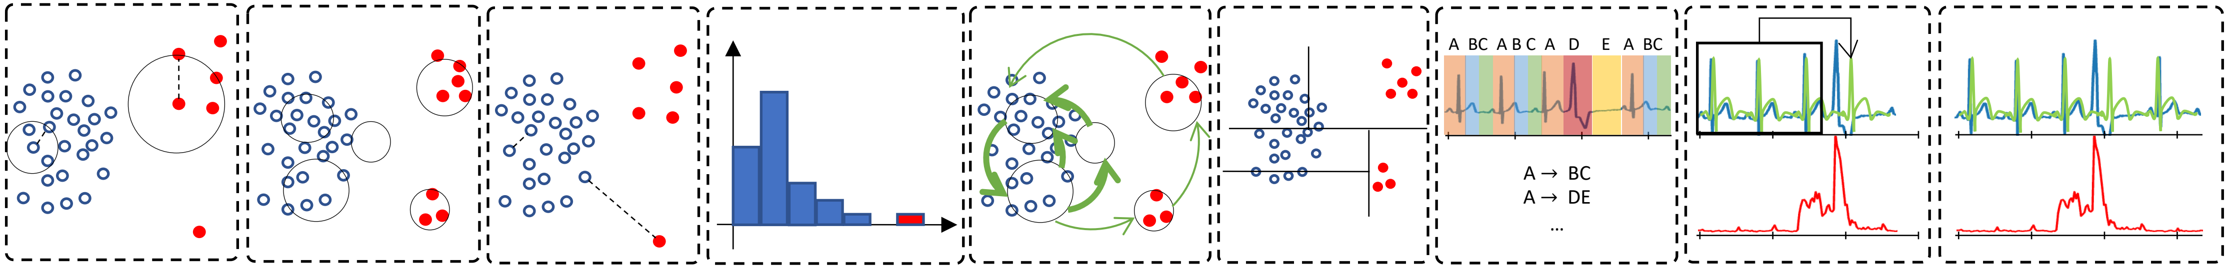

Anomaly Detection methods: *A taxonomy*

By time…

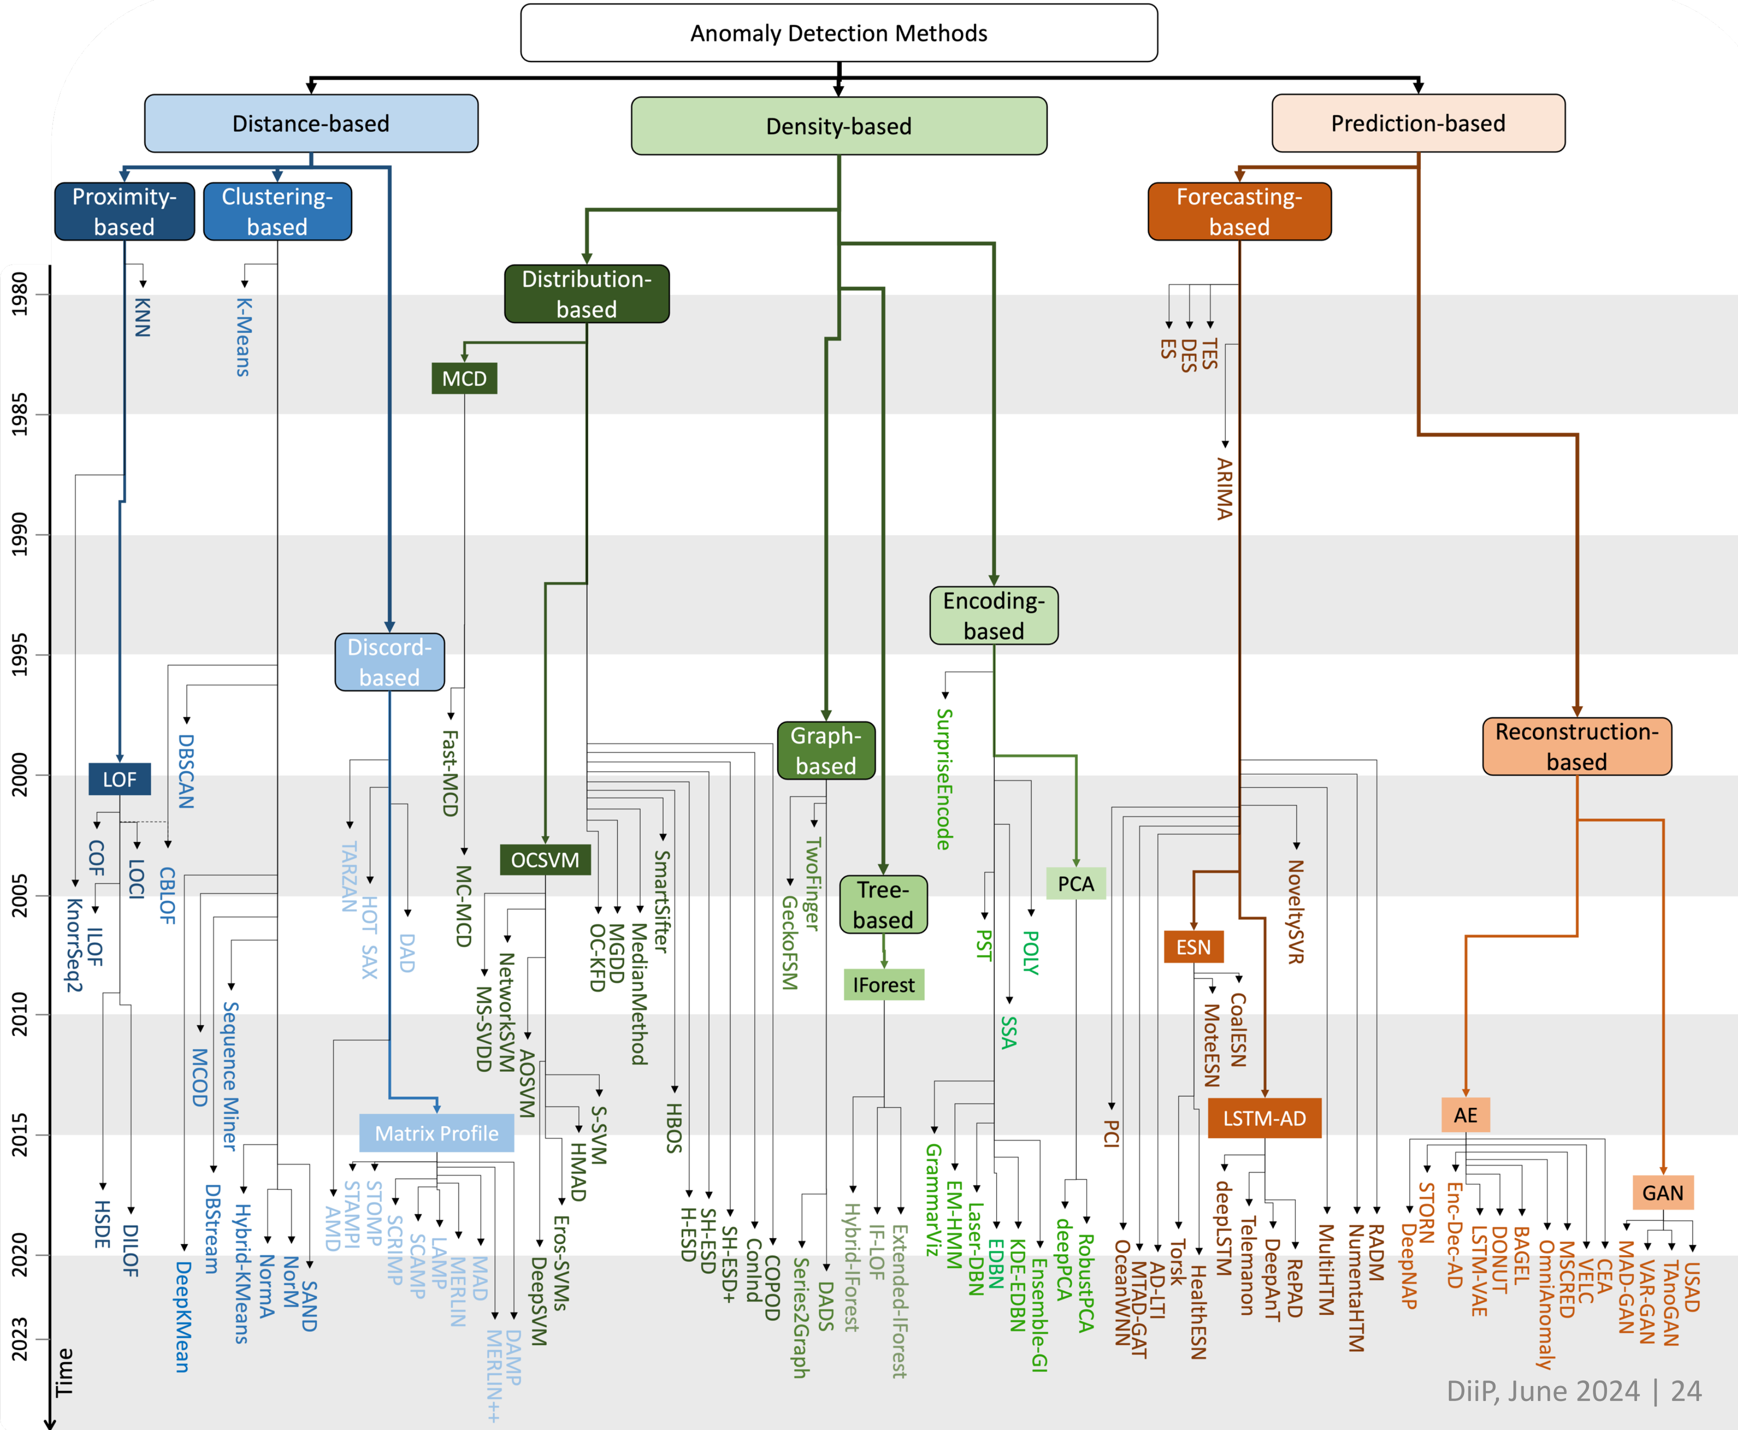

Anomaly Detection methods: *A taxonomy*

By time…

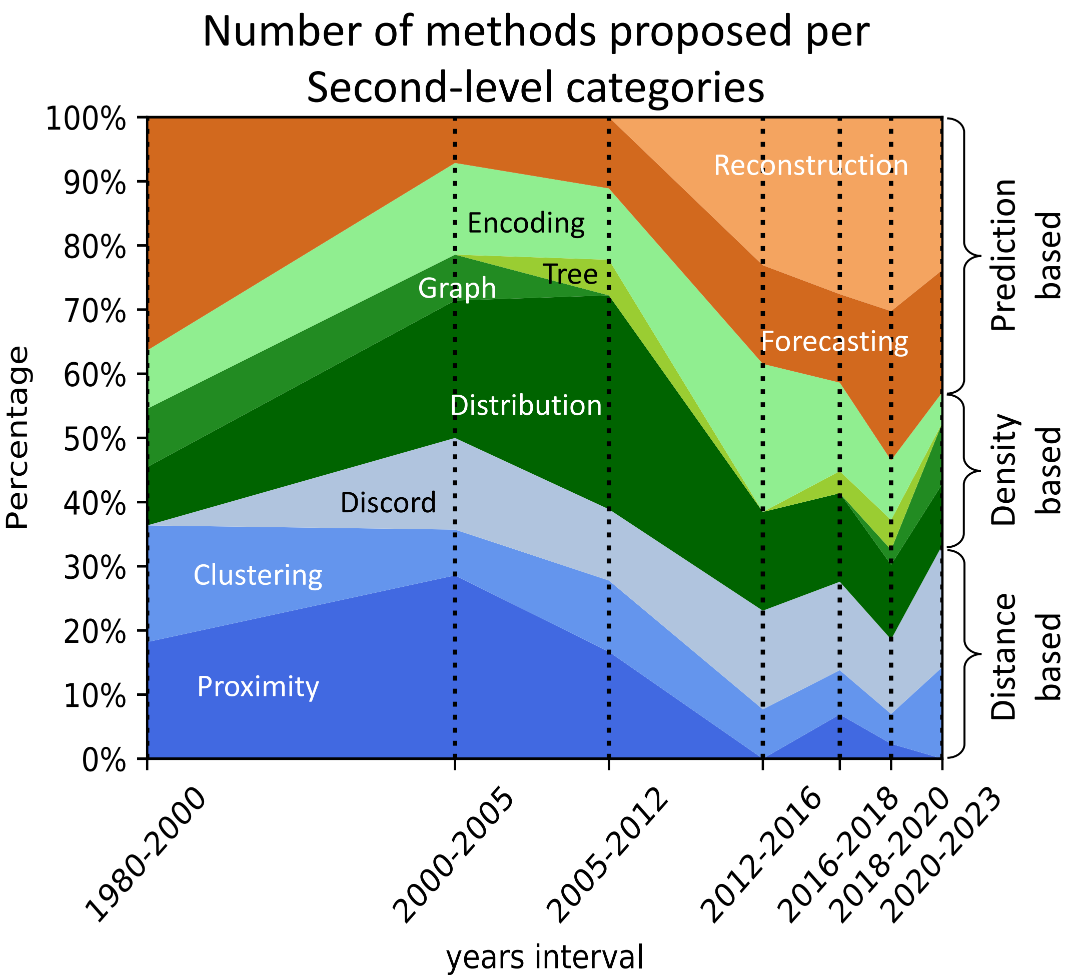

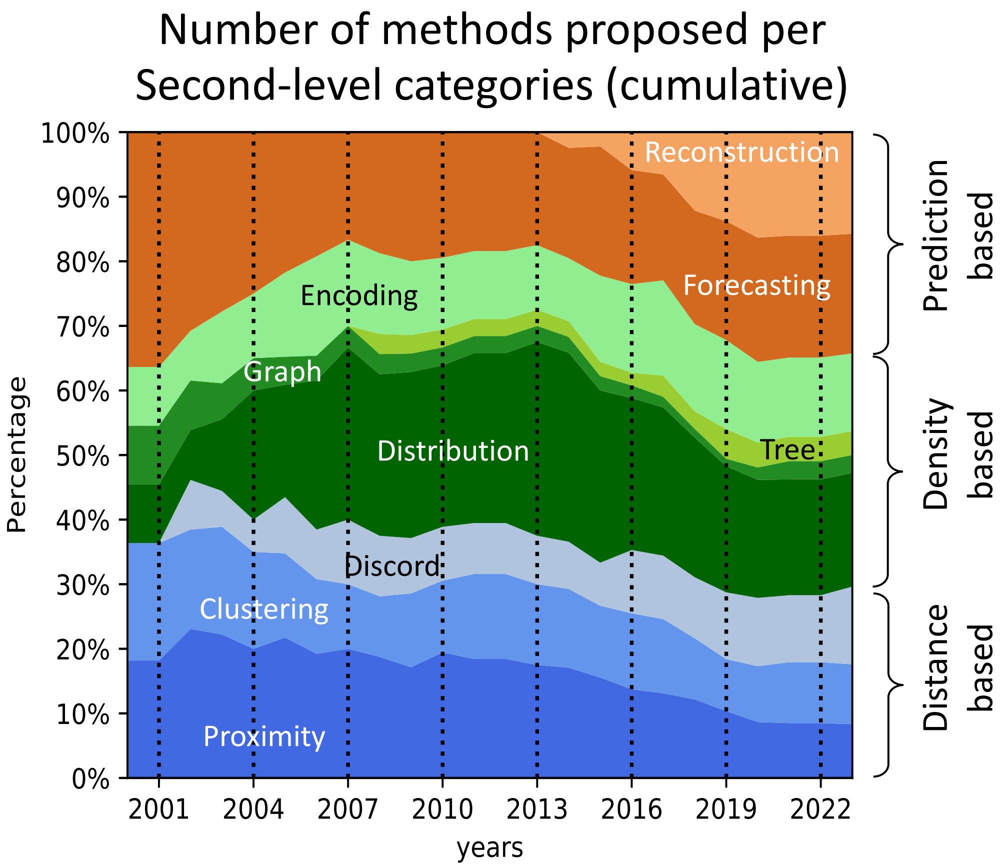

Anomaly Detection methods: *A taxonomy*

By time…

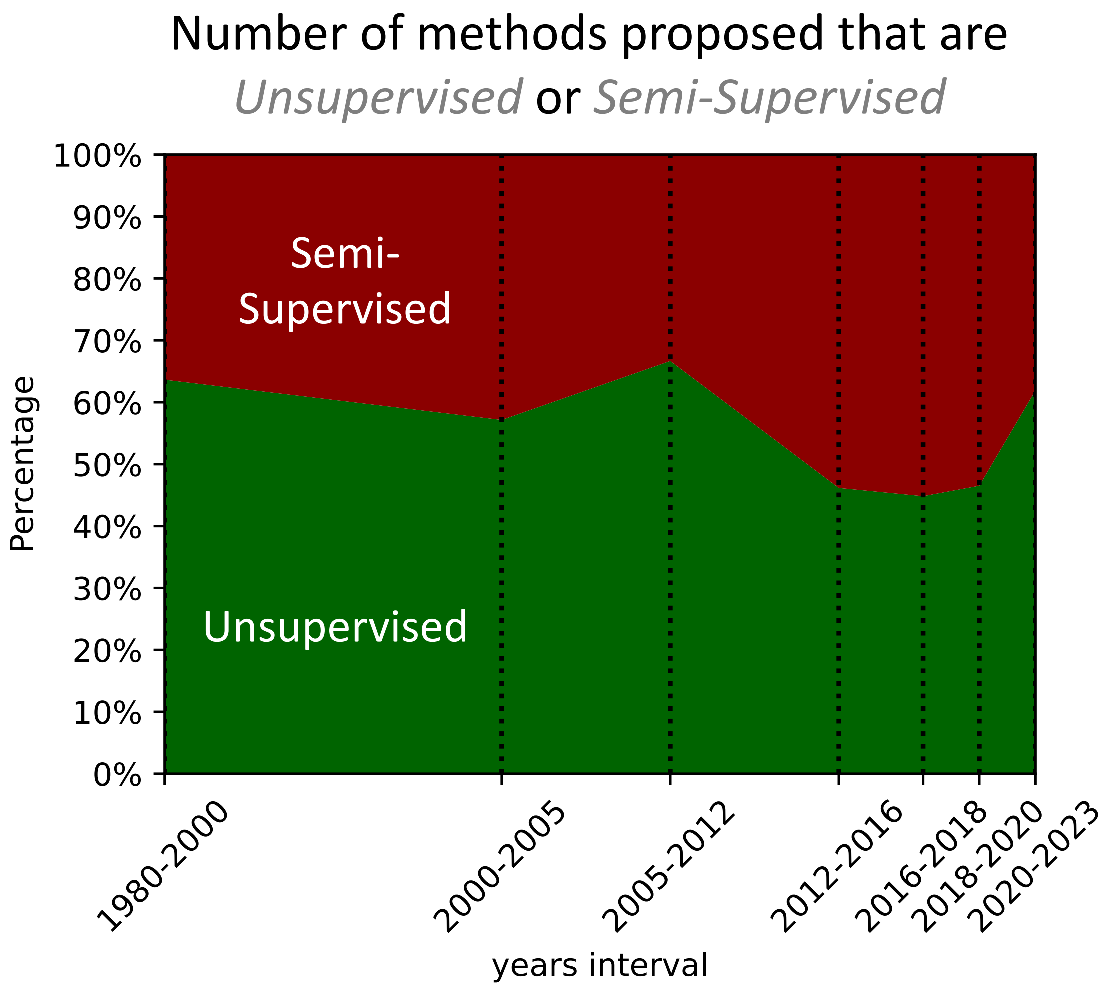

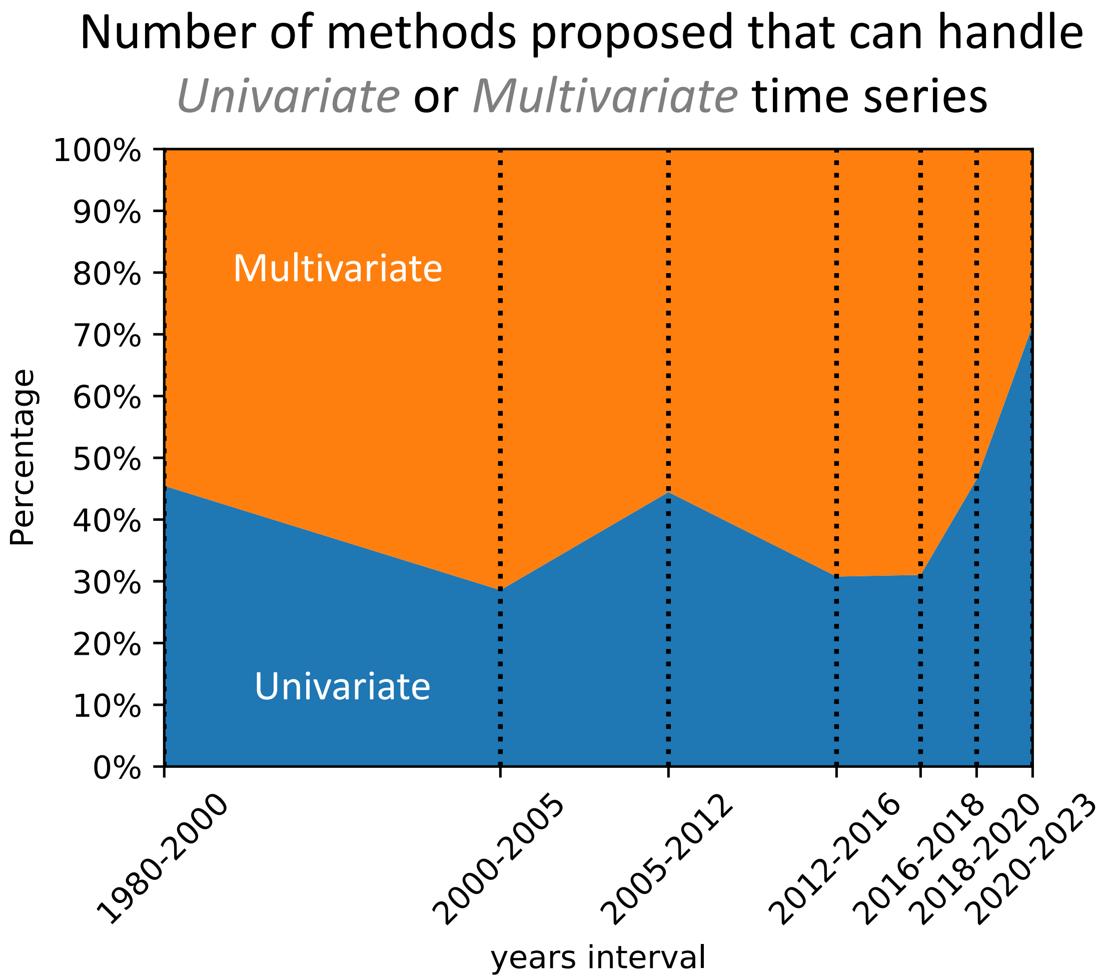

Anomaly Detection methods: *Distance-based*

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

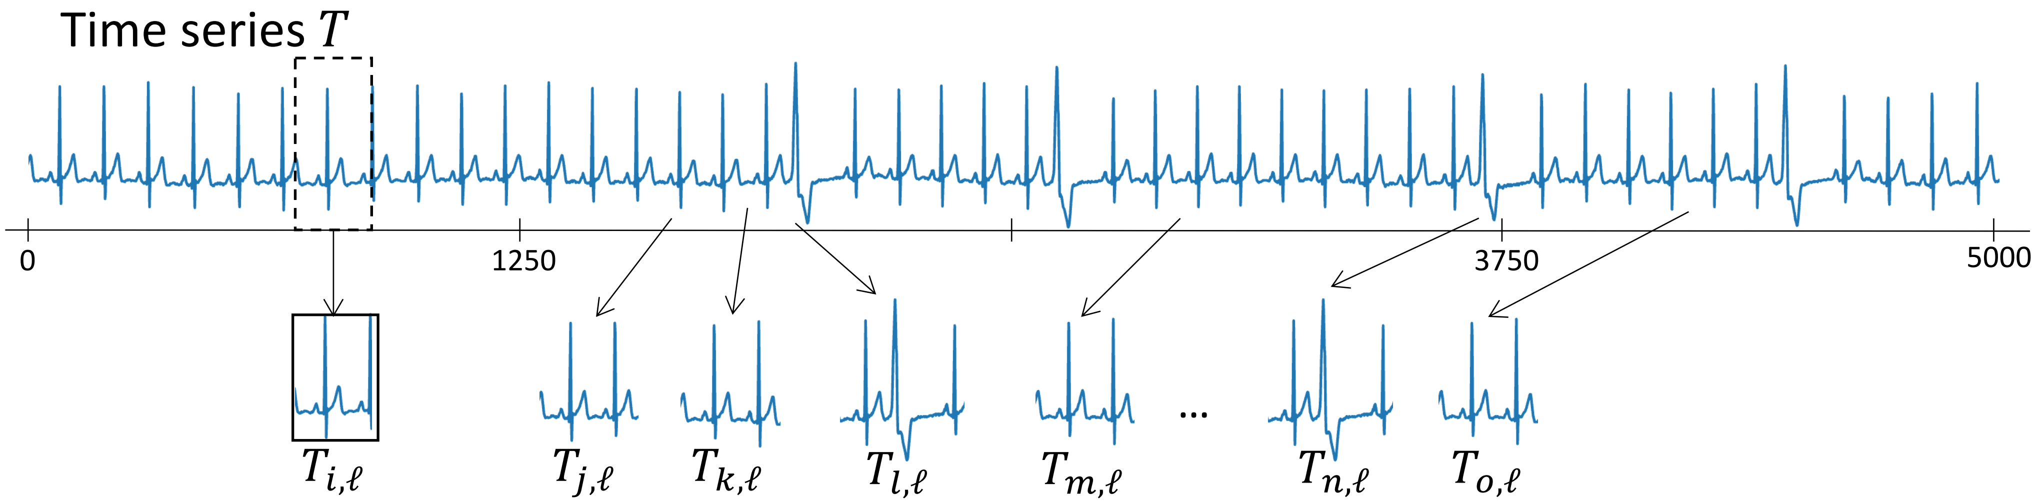

Anomaly Detection methods: *Distance-based*

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

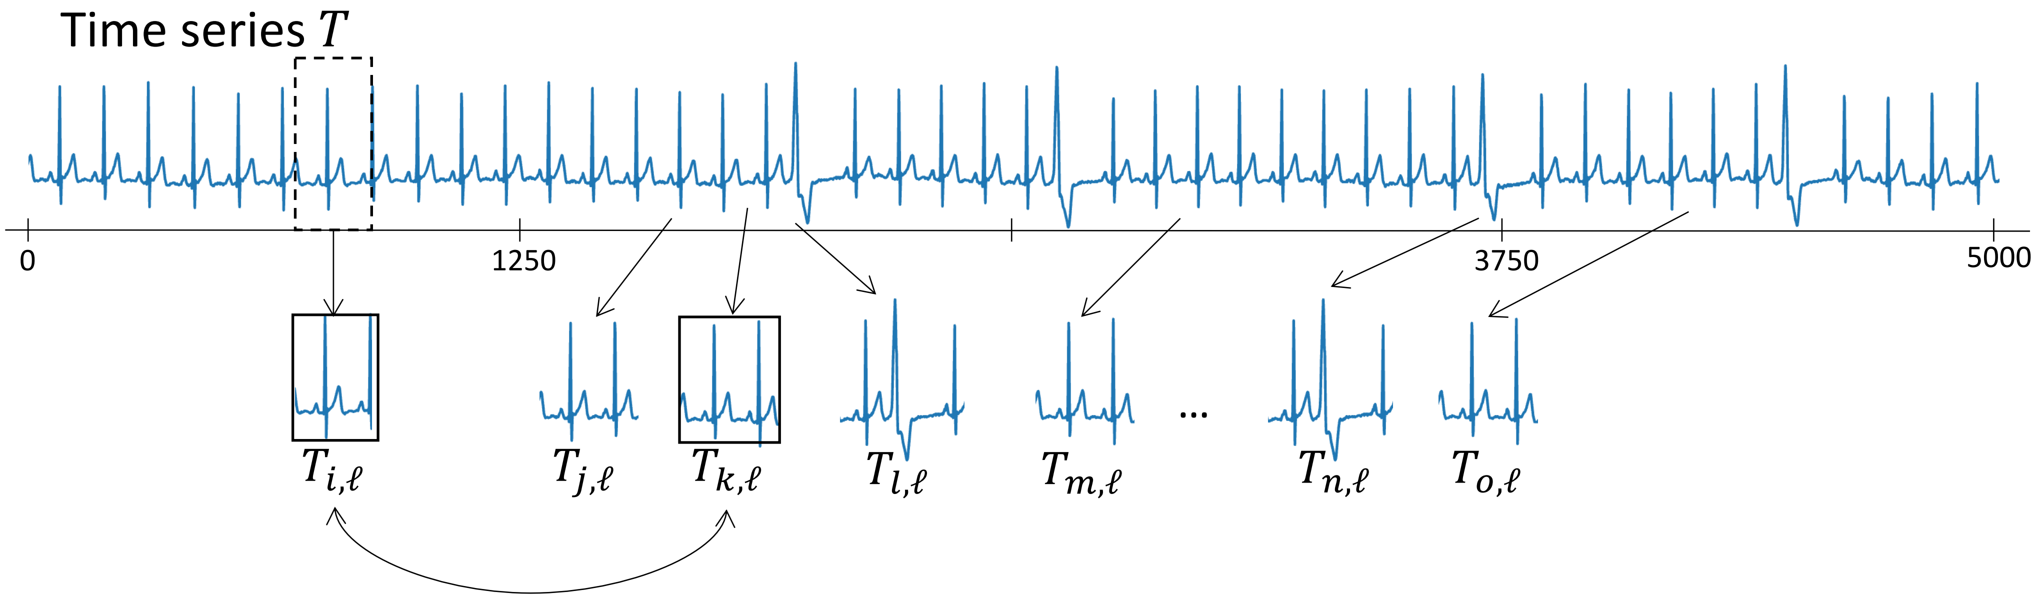

$S_{T,i} = d(T_{i,\ell}, T_{j,\ell})$

Anomaly Detection methods: *Distance-based*

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

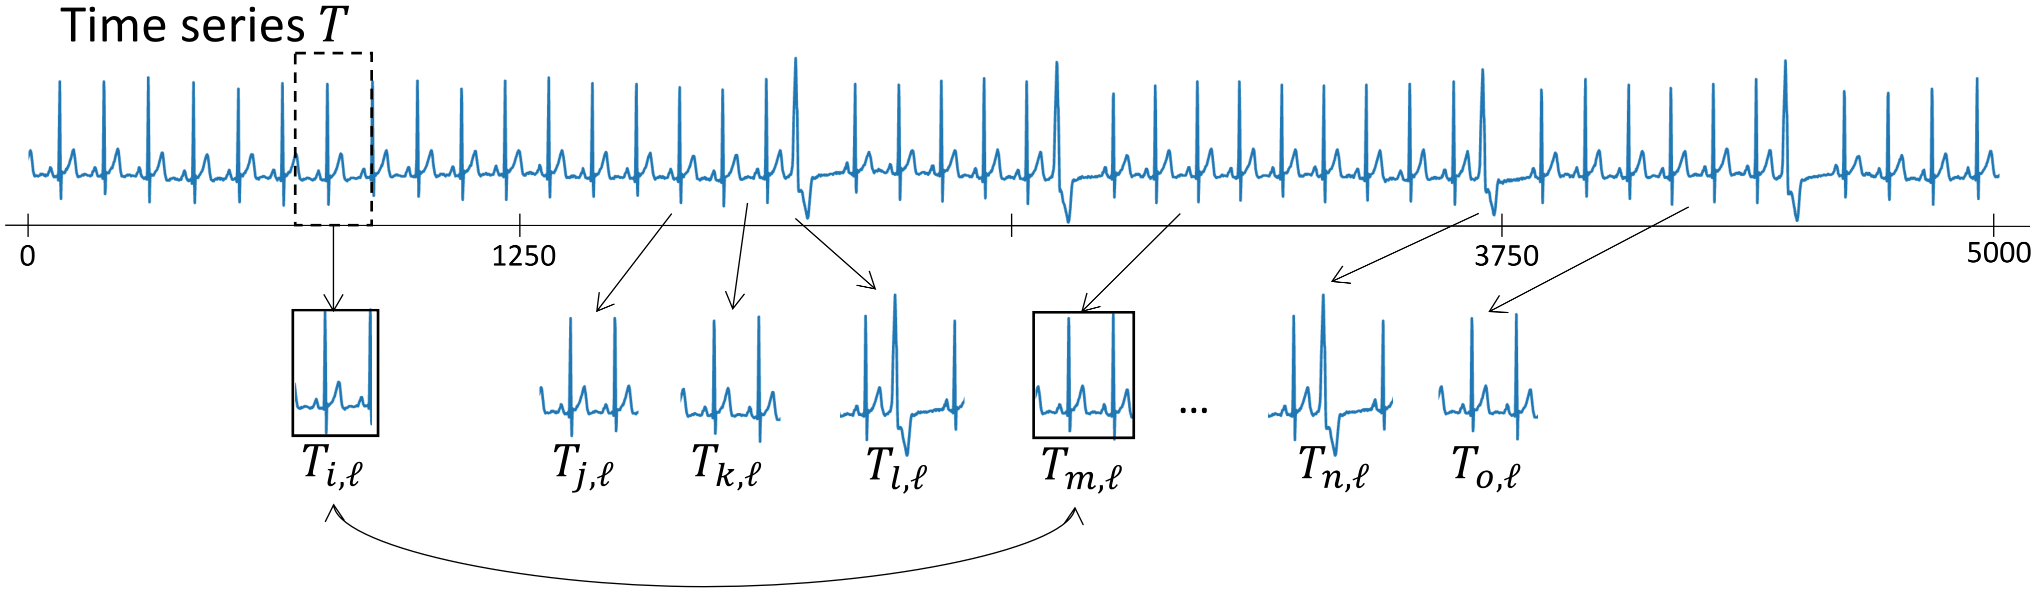

$$S_{T,i} = d(T_{i,\ell}, T_{m,\ell})$$

Anomaly Detection methods: *Distance-based*

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

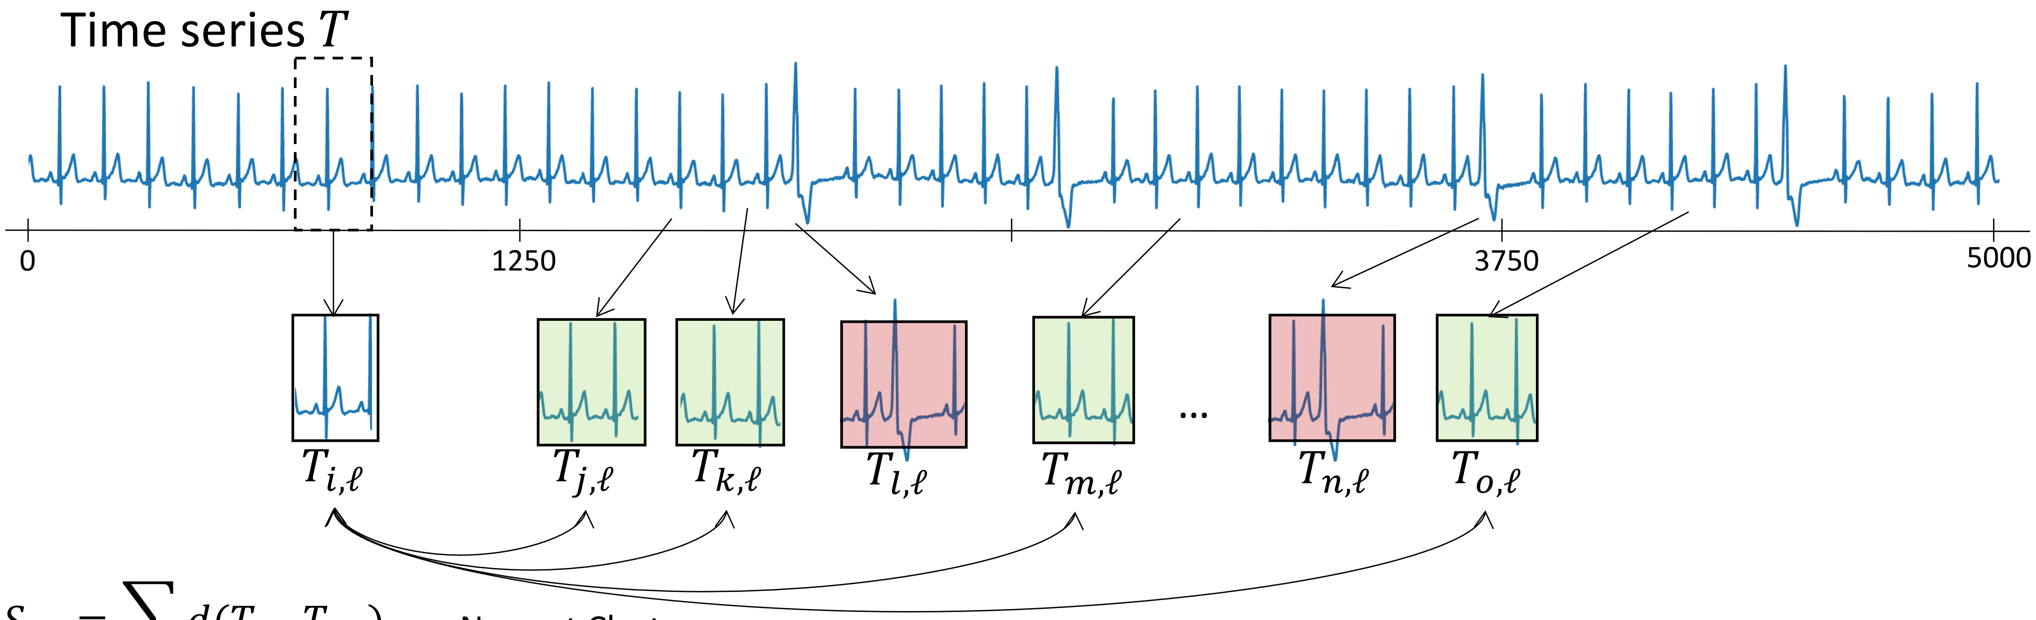

$$S_{T,i} = \sum_{k \in \mathcal{C}} d(T_{i,\ell}, T_{k,\ell})$$

Nearest Cluster

Anomaly Detection methods: *Distance-based*

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

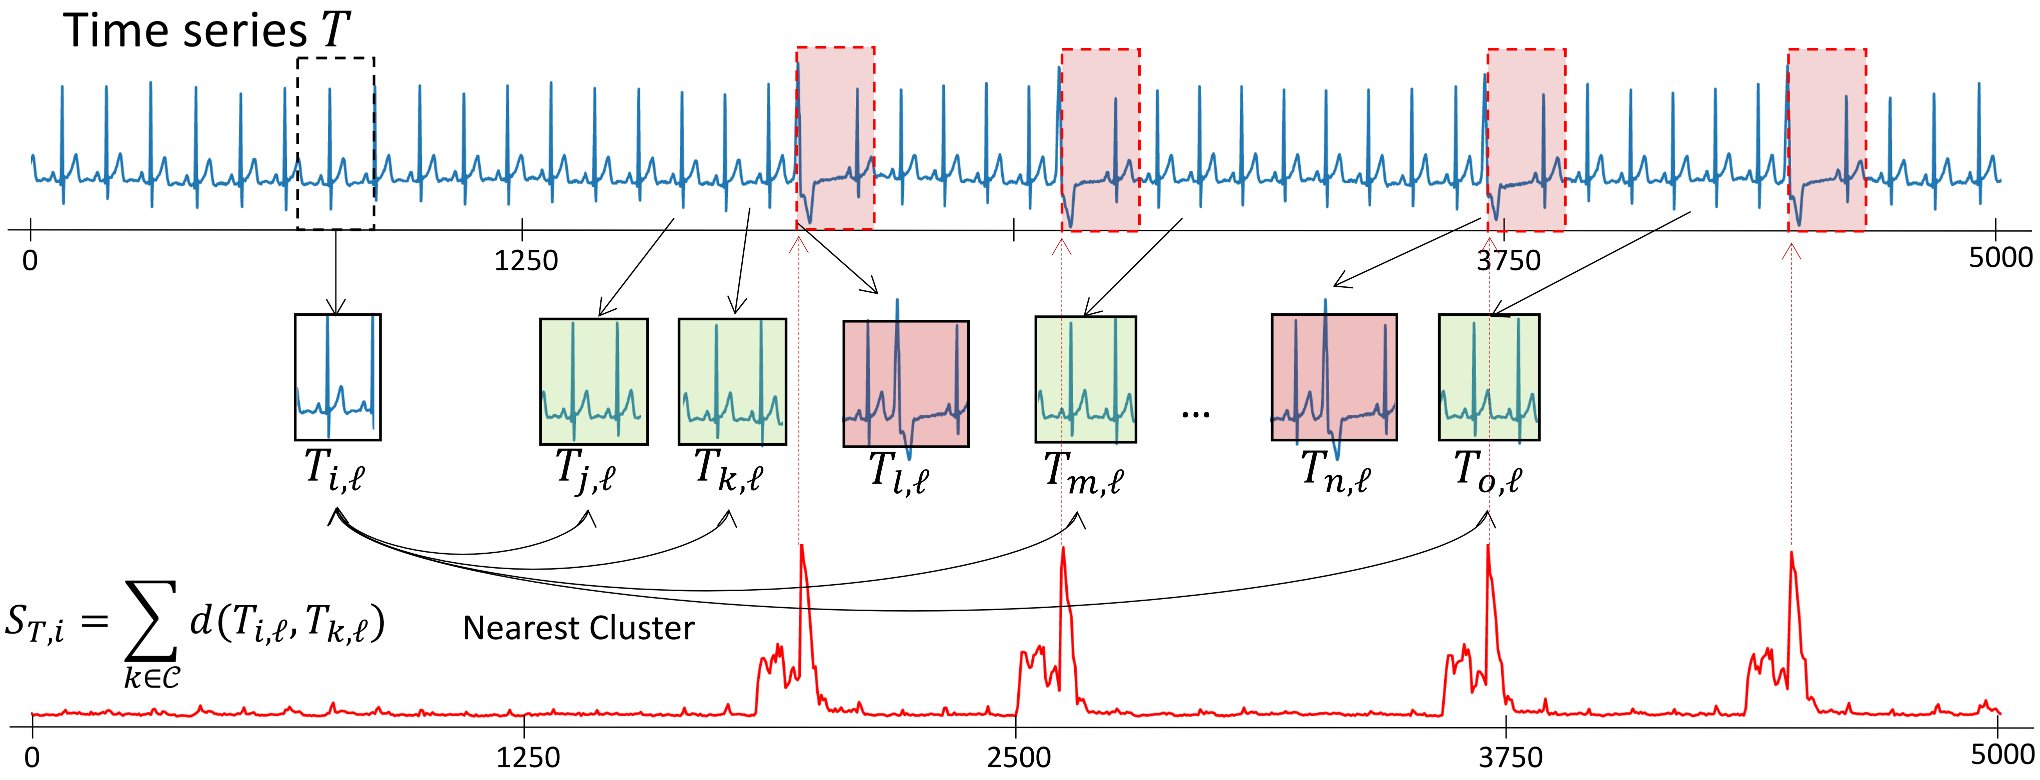

Anomaly Detection methods: *Distance-based*

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

Example of distance computation

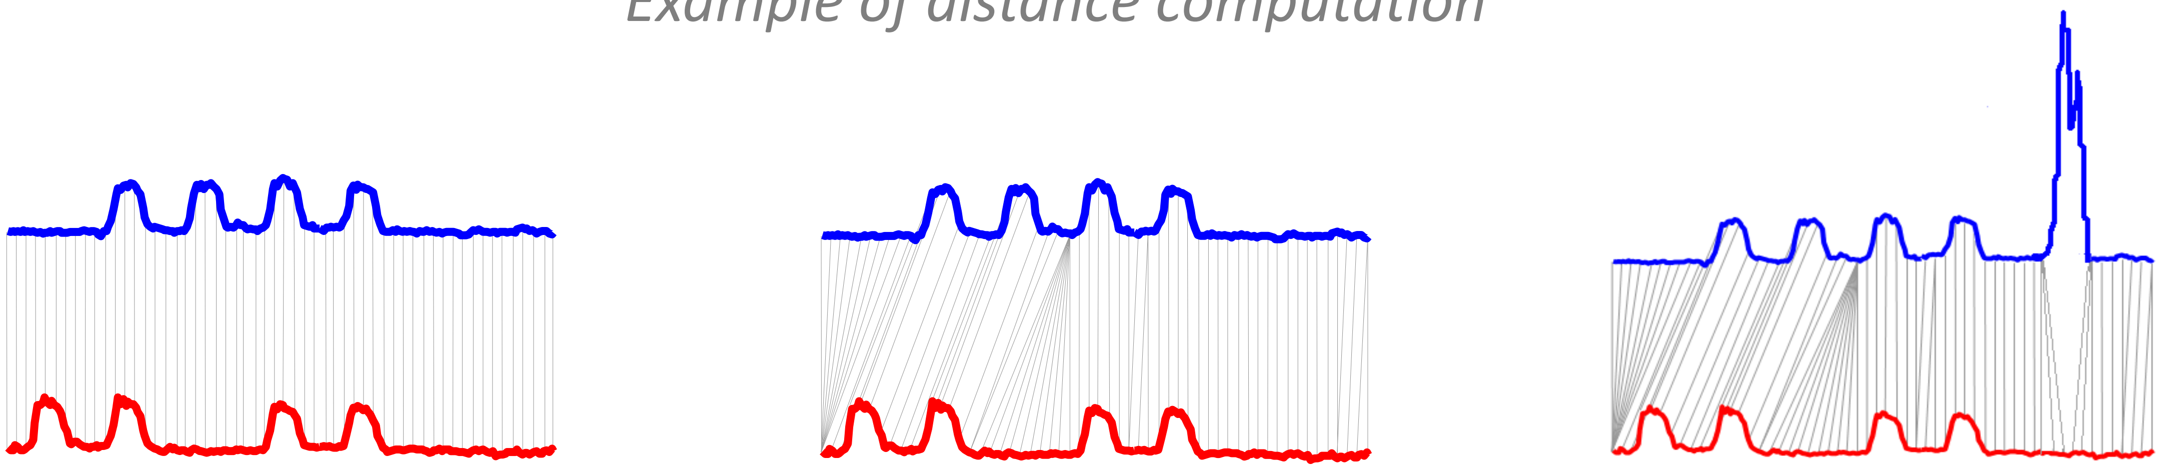

(a) Euclidean Distance (b) DTW Distance (c) LCSS Distance

$S_{T,i} = \sum_{k \in \mathcal{C}} d(T_{i,\ell}, T_{k,\ell})$ Nearest Cluster

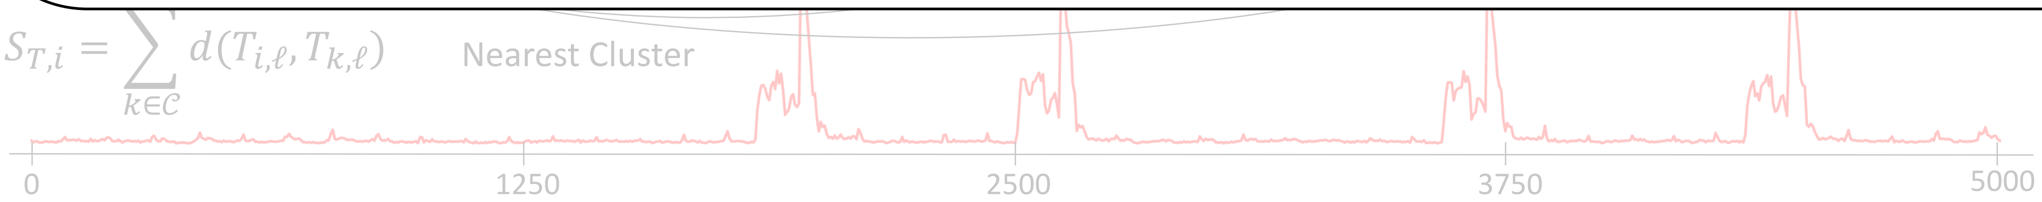

Anomaly Detection methods: *an Example*

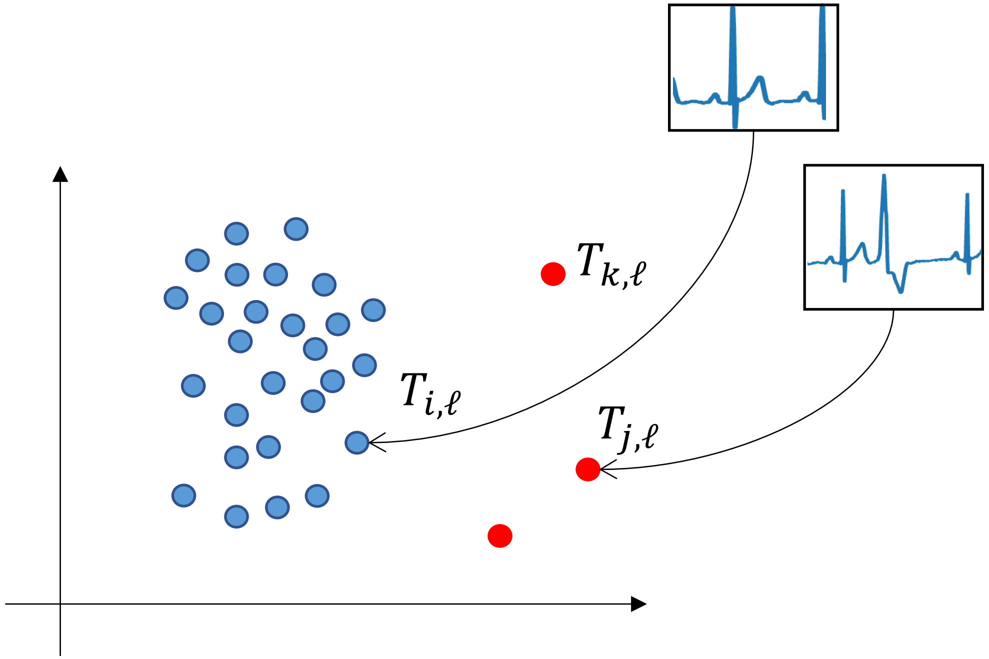

Matrix Profile [6] (MP)

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

[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 Prole I: All Pairs Similarity Joins for Time Series. In ICDM.

Anomaly Detection methods: *an Example*

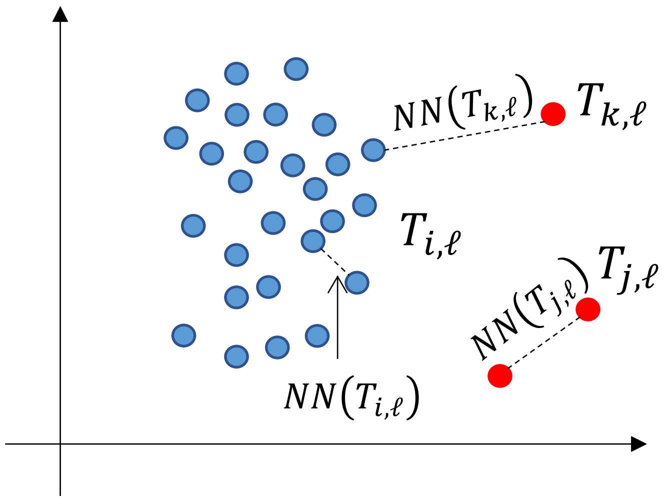

The matrix Profile is computed as follows:

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

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

[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 Prole I: All Pairs Similarity Joins for Time Series. In ICDM.

Anomaly Detection methods: *an Example*

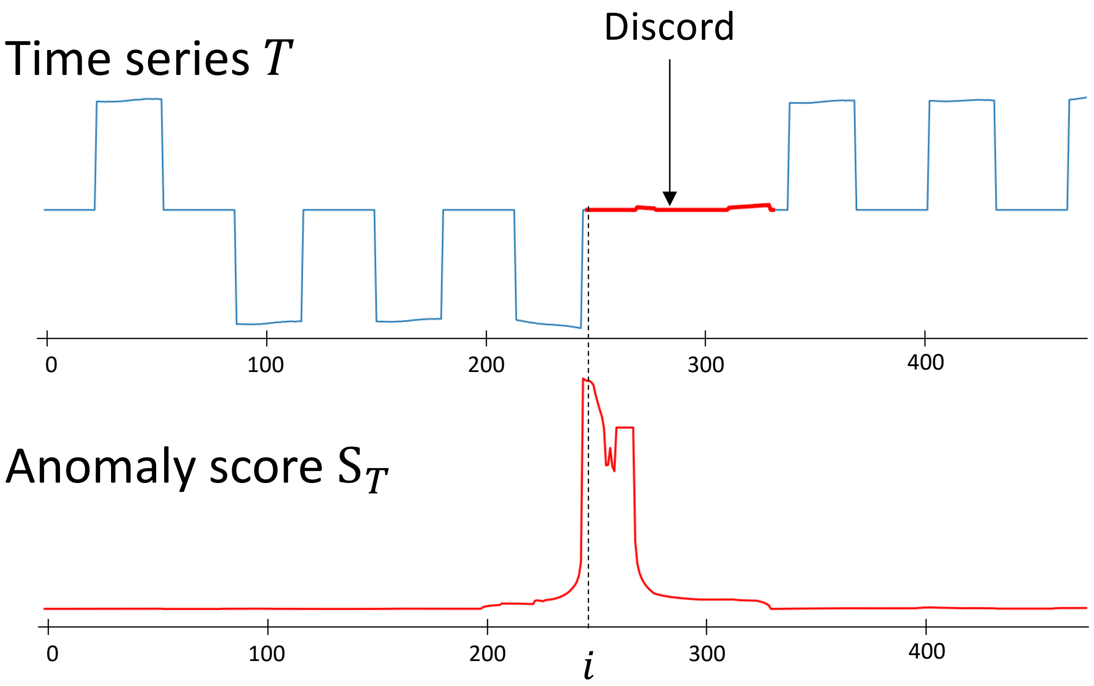

Matrix Profile [6] (MP)

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

[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 Prole I: All Pairs Similarity Joins for Time Series. In ICDM.

Anomaly Detection methods: *an Example*

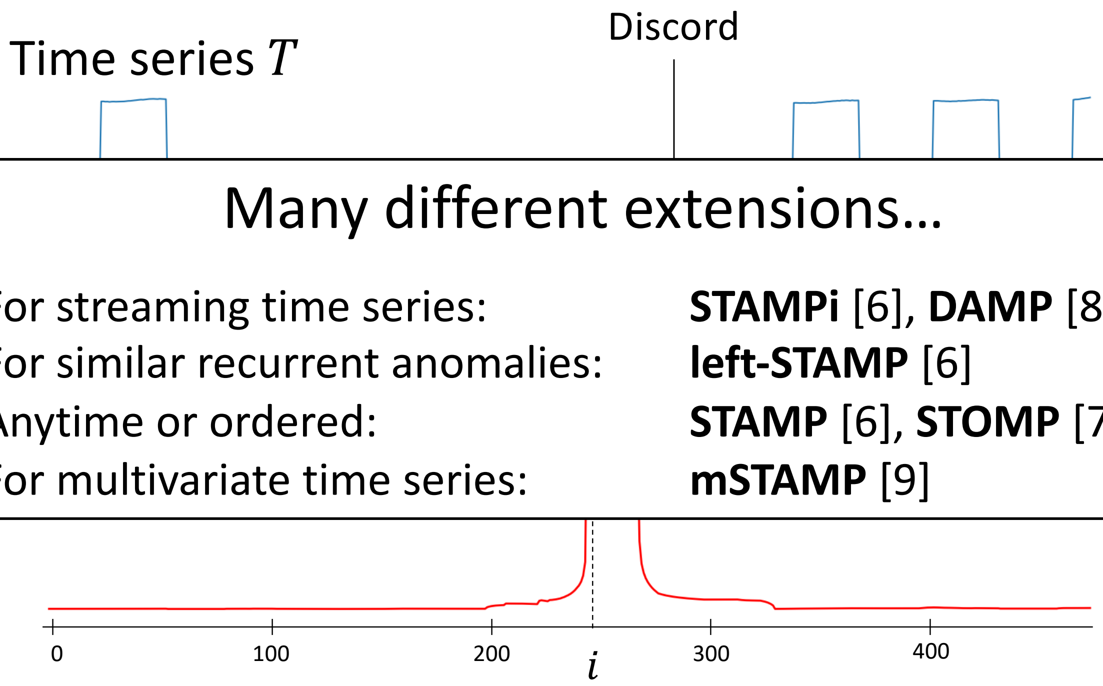

Many different extensions…

- For streaming time series: **STAMPi** [6], **DAMP** [8]
- For similar recurrent anomalies: **left-STAMP** [6]
- Anytime or ordered: **STAMP** [6], **STOMP** [7]
- For multivariate time series: **mSTAMP** [9]

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

[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 Prole I: All Pairs Similarity Joins for Time Series. In ICDM.

Anomaly Detection methods: *an Example*

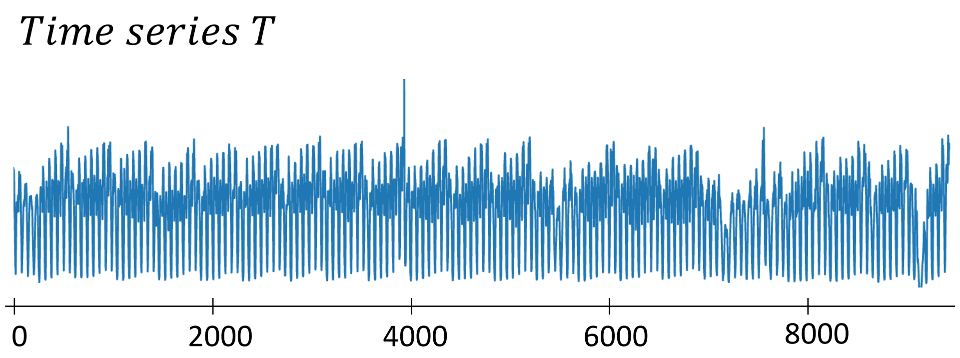

NormA [10]

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

[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.

Anomaly Detection methods: *an Example*

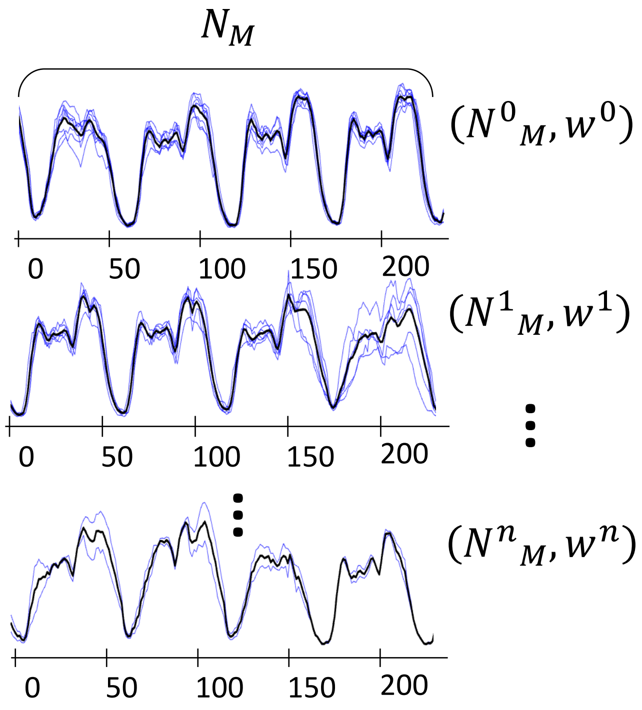

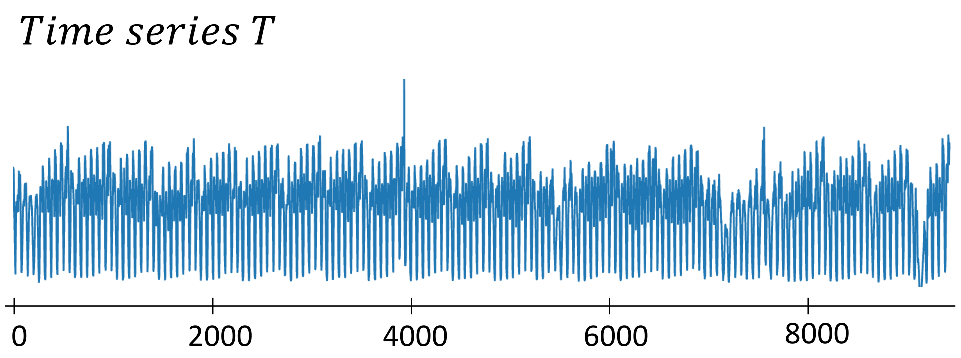

NormA [10]

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

[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.

Anomaly Detection methods: *an Example*

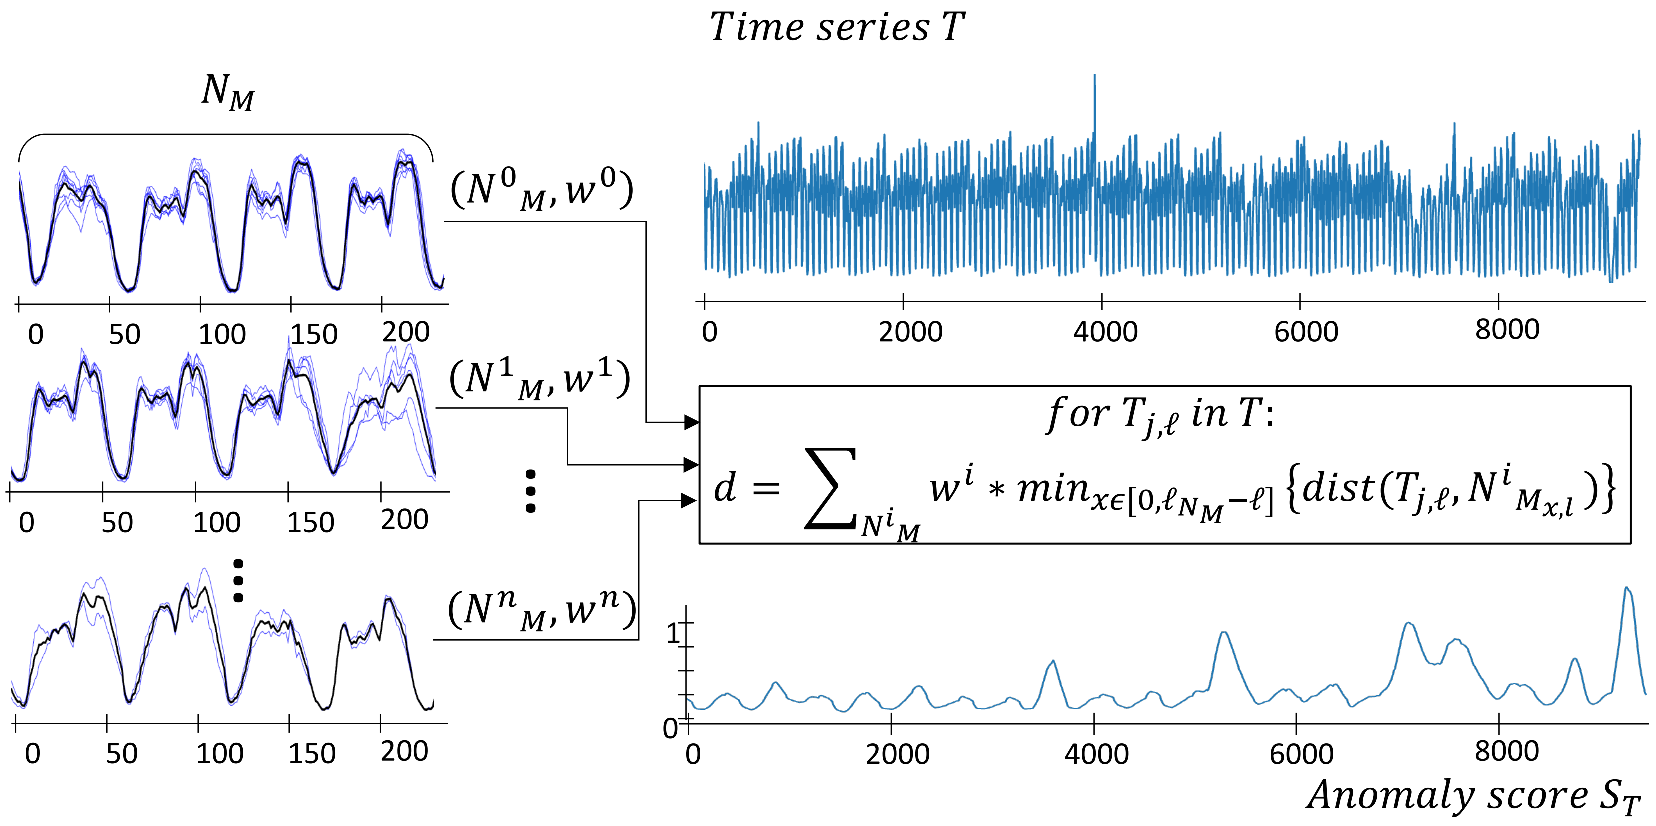

NormA [10]

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

[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.

Anomaly Detection methods: *an Example*

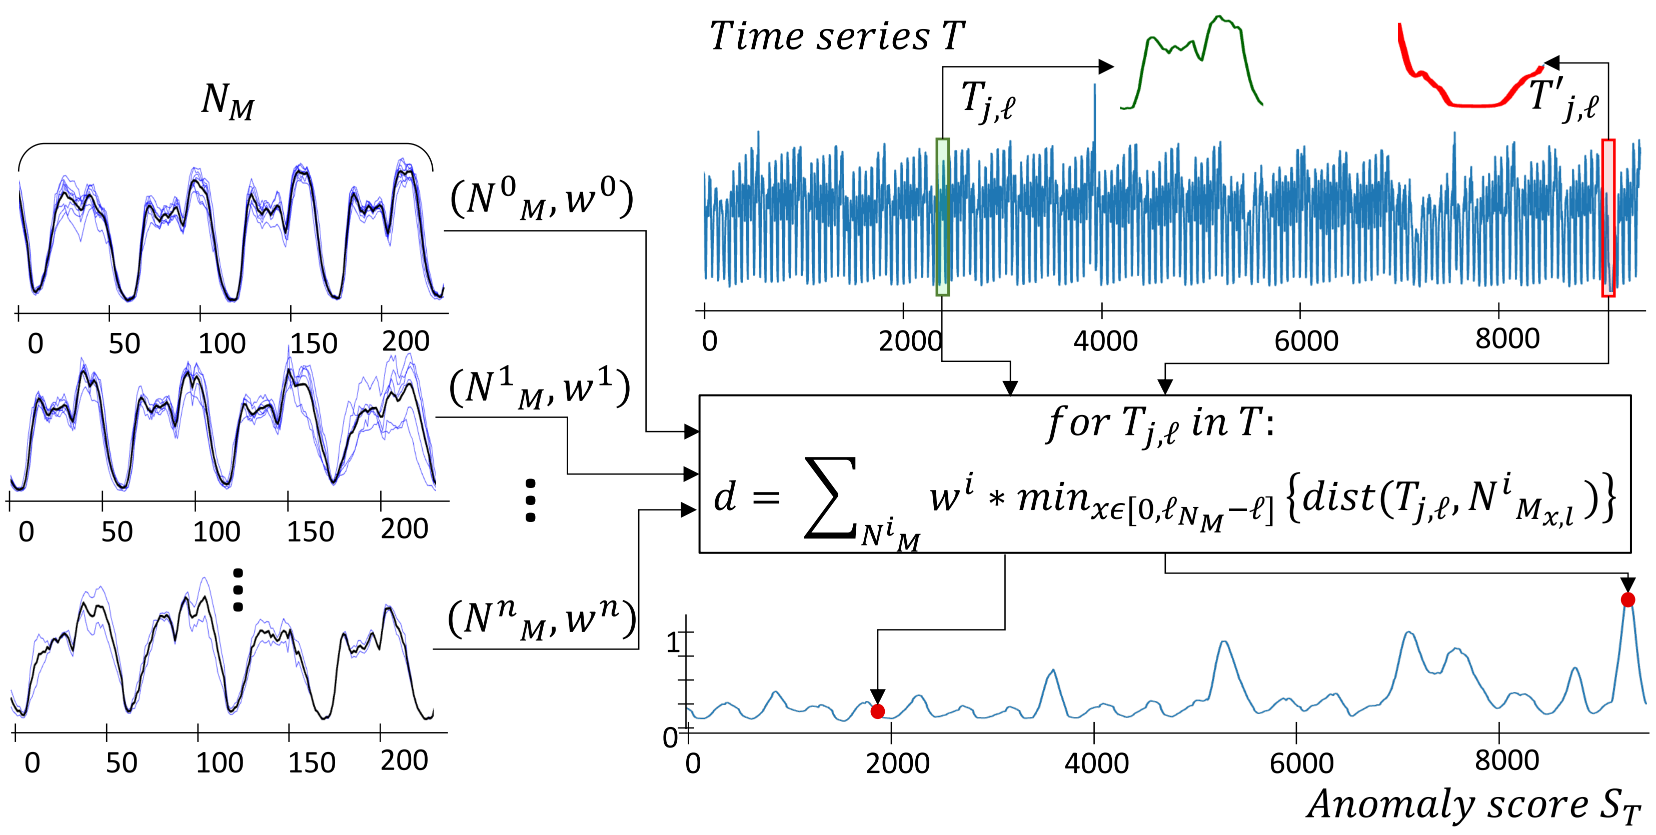

NormA [10]

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

[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.

Anomaly Detection methods: *an Example*

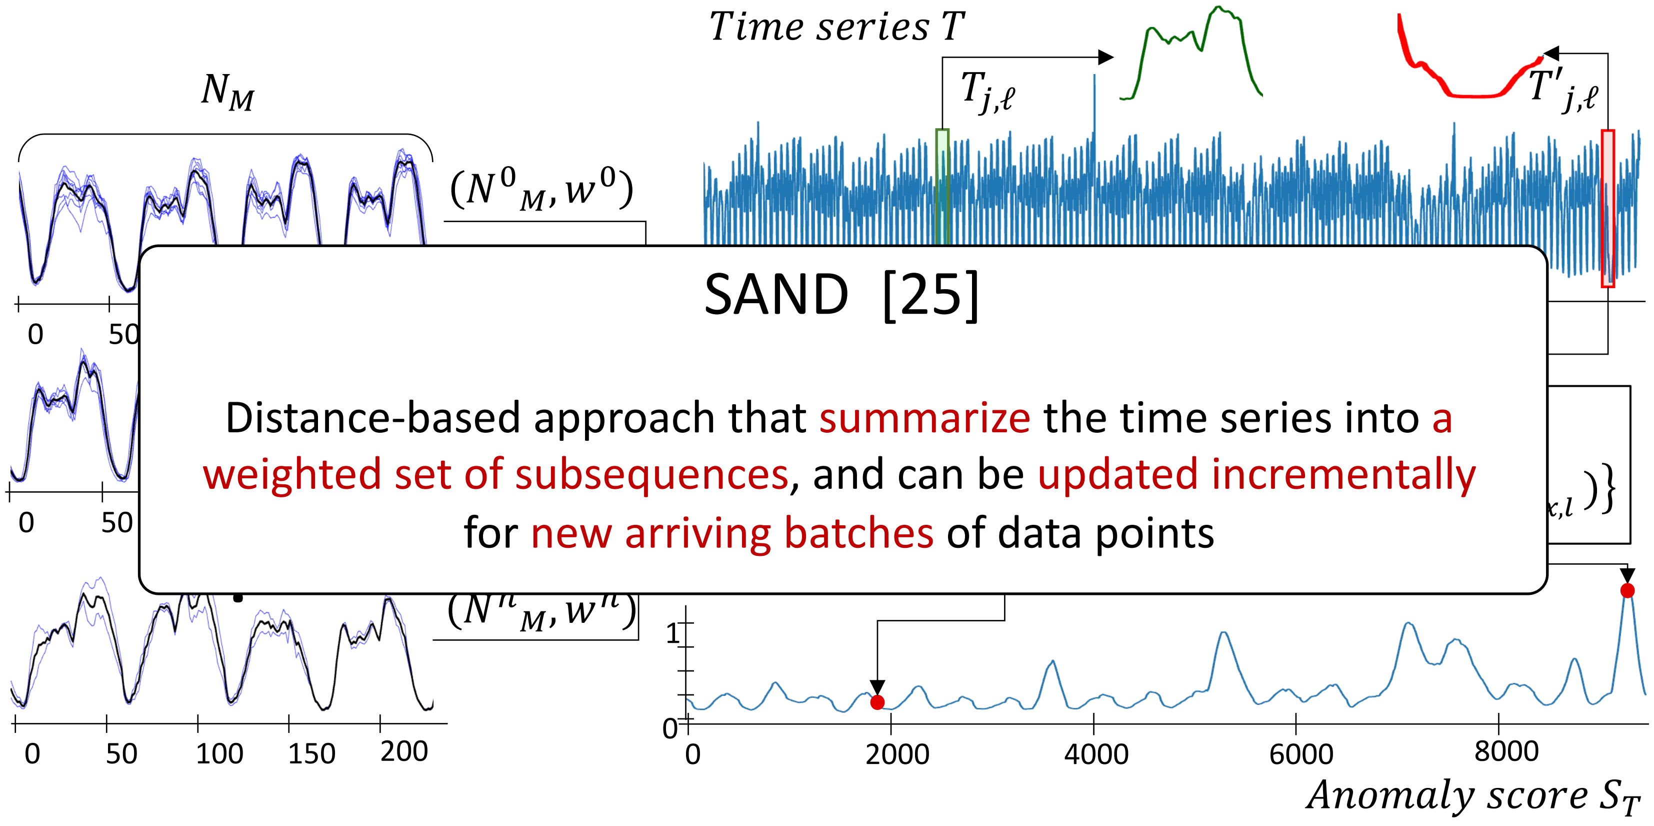

SAND [25]

Distance-based approach that summarize the time series into a weighted set of subsequences, and can be updated incrementally for new arriving batches of data points

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

[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.

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

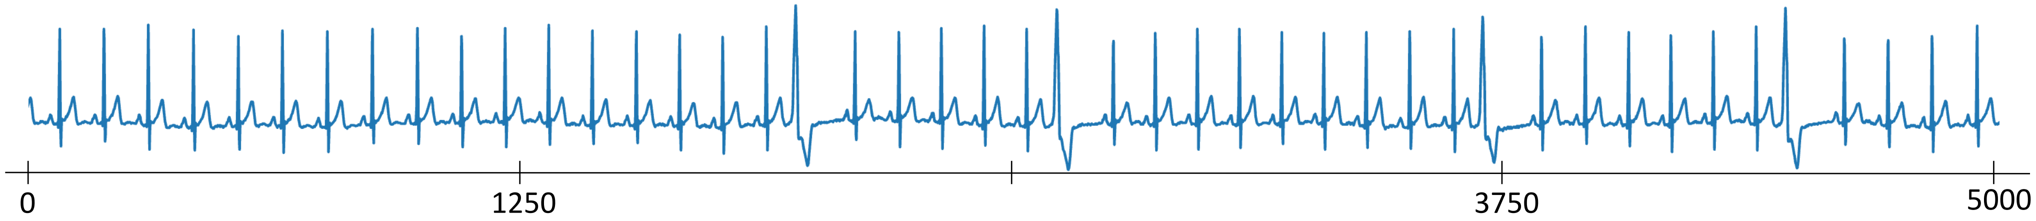

Anomaly Detection methods: *Density-based*

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

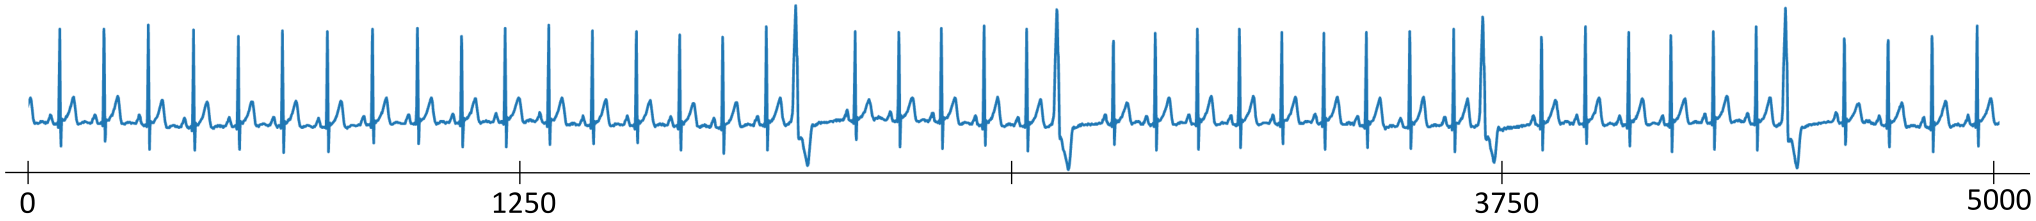

Tree-based approaches [11]

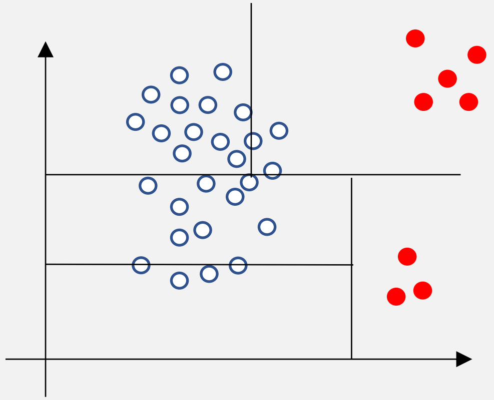

Distribution-based Approaches [12]

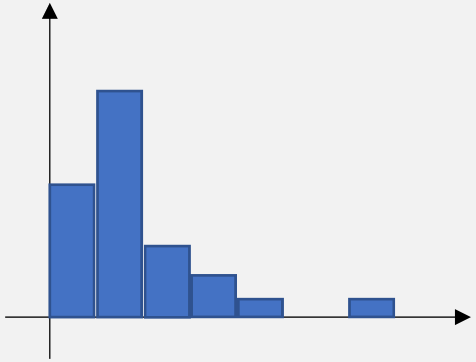

Graph-based approaches [13]

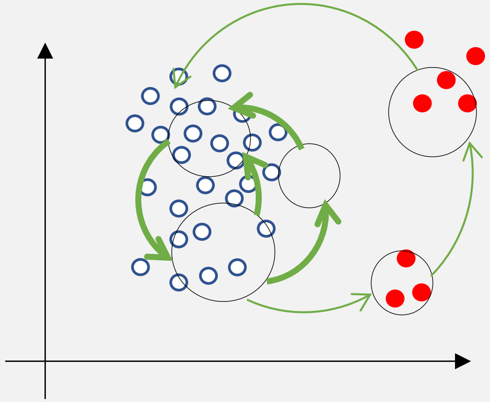

…

Anomaly Detection methods: *an Example*

0 splits

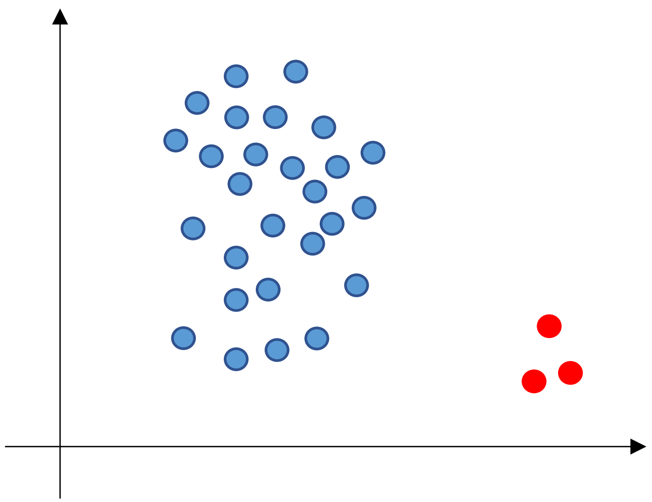

0 splits

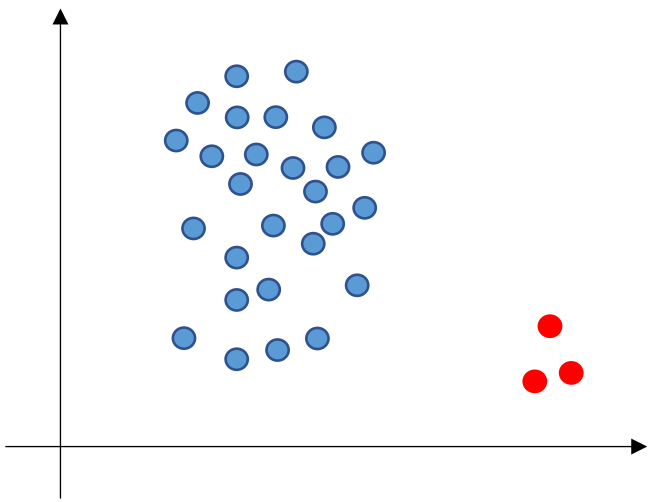

Isolation Forest [11]

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

Unsupervised

Univariate/Multivariate

Point/sequence

[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

Anomaly Detection methods: *an Example*

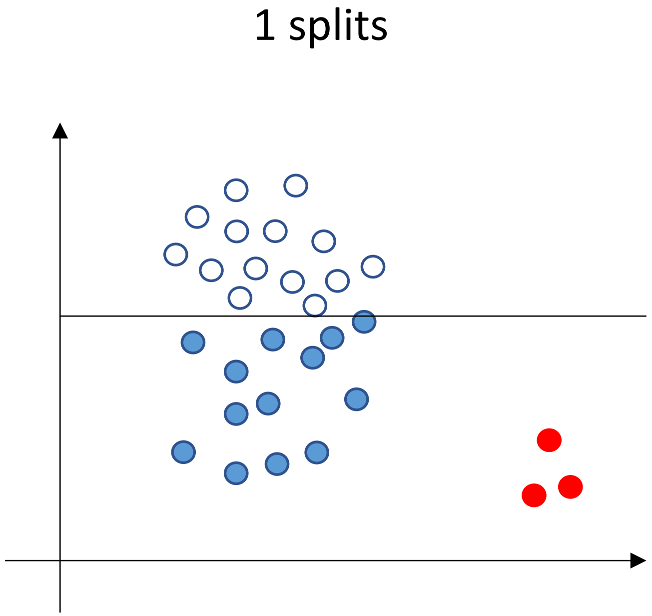

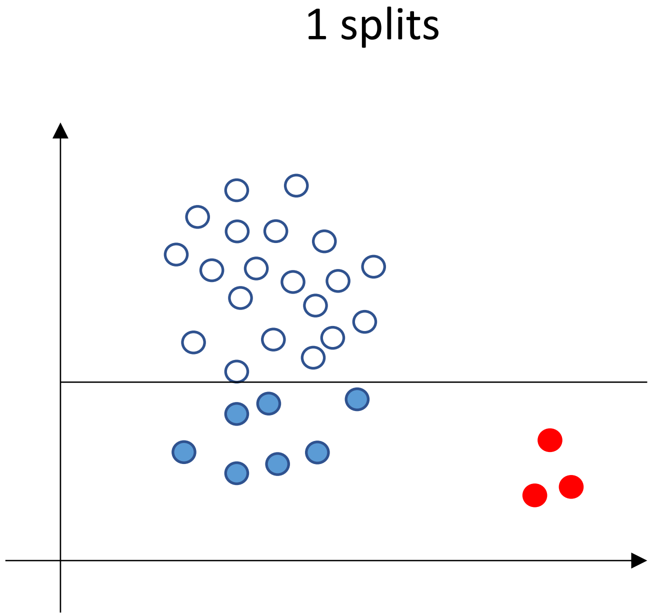

Isolation Forest [11]

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

Unsupervised

Univariate/Multivariate

Point/sequence

[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

Anomaly Detection methods: *an Example*

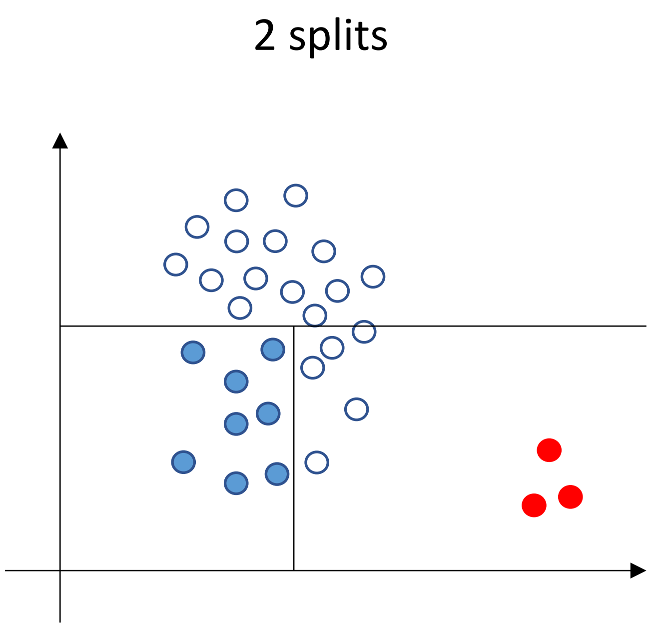

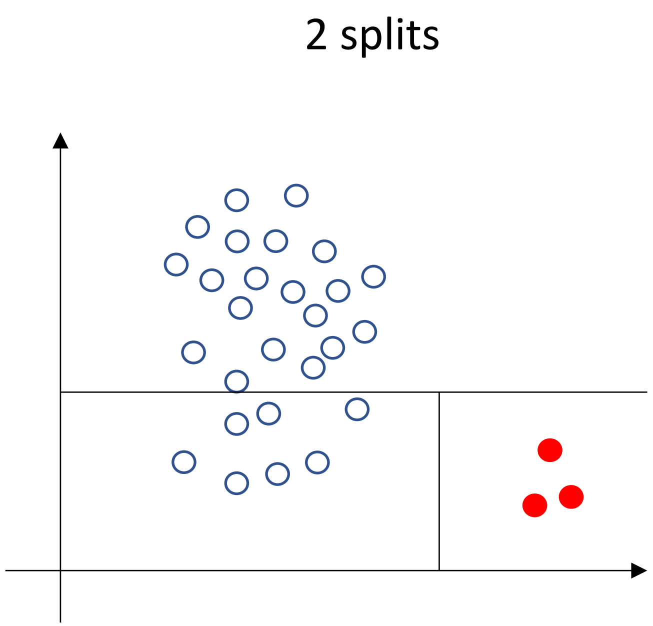

Isolation Forest [11]

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

Unsupervised

Univariate/Multivariate

Point/sequence

[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

Anomaly Detection methods: *an Example*

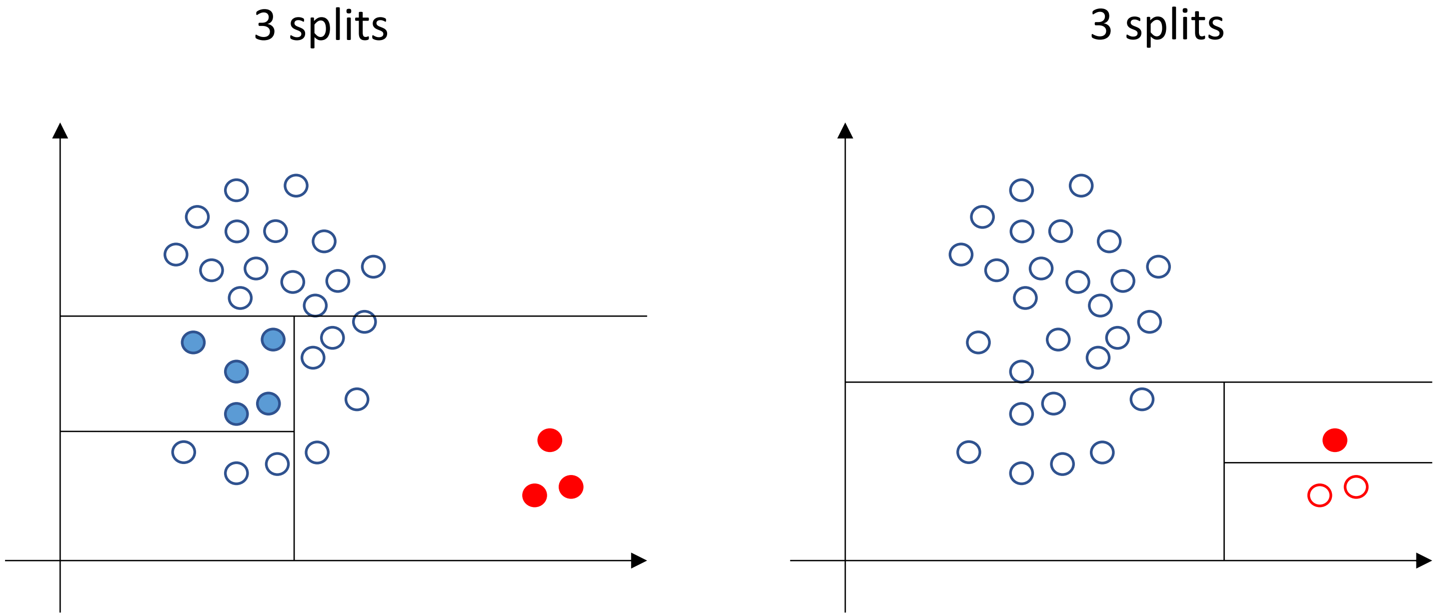

Isolation Forest [11]

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

Unsupervised

Univariate/Multivariate

Point/sequence

[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

Anomaly Detection methods: *an Example*

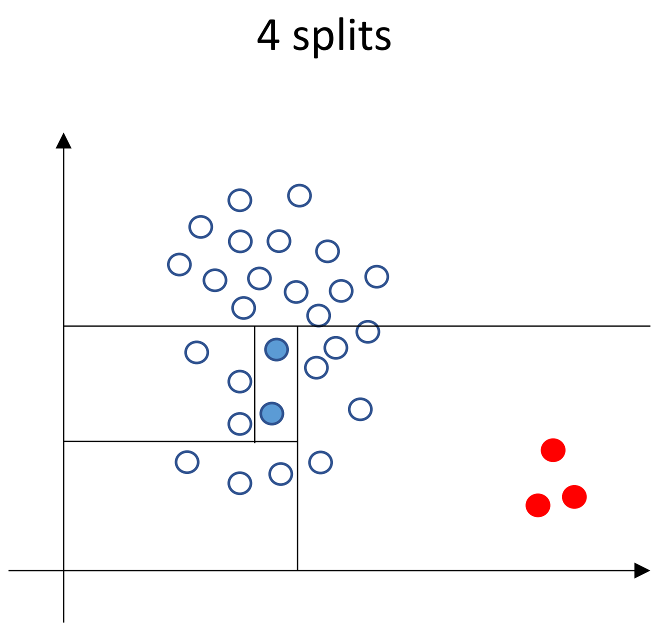

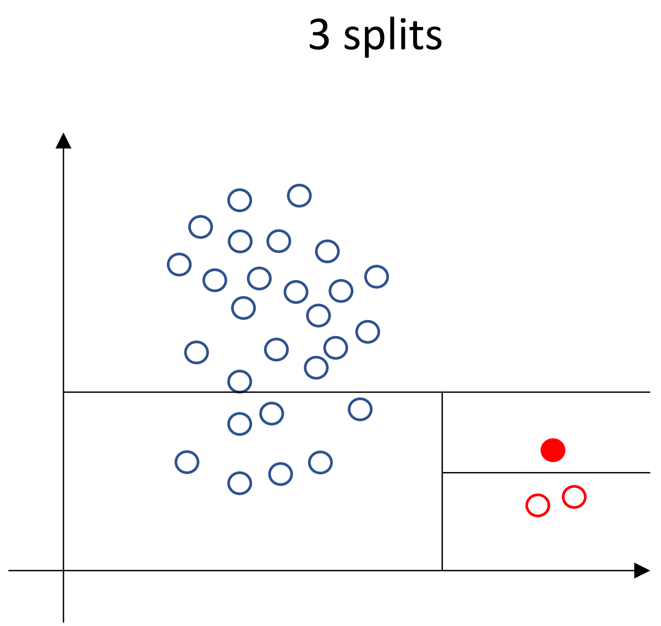

Isolation Forest [11]

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

Unsupervised

Univariate/Multivariate

Point/sequence

[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

Anomaly Detection methods: *an Example*

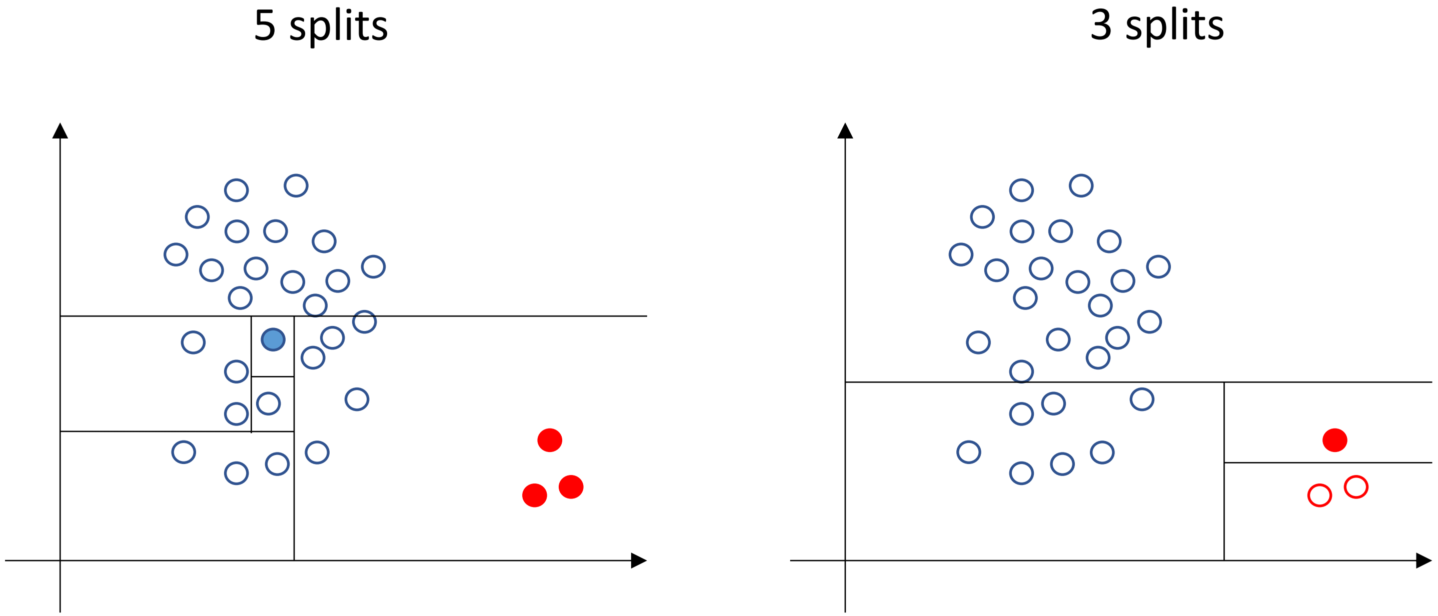

Isolation Forest [11]

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

Unsupervised

Univariate/Multivariate

Point/sequence

[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

Anomaly Detection methods: *an Example*

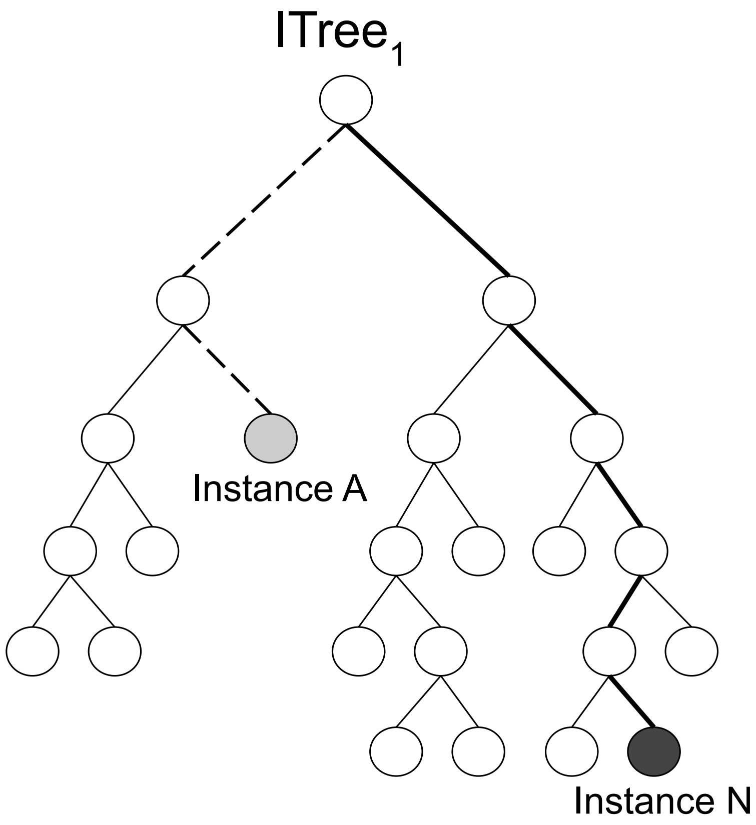

Isolation Forest [11]

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

Unsupervised

Univariate/Multivariate

Point/sequence

[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

Anomaly Detection methods: *an Example*

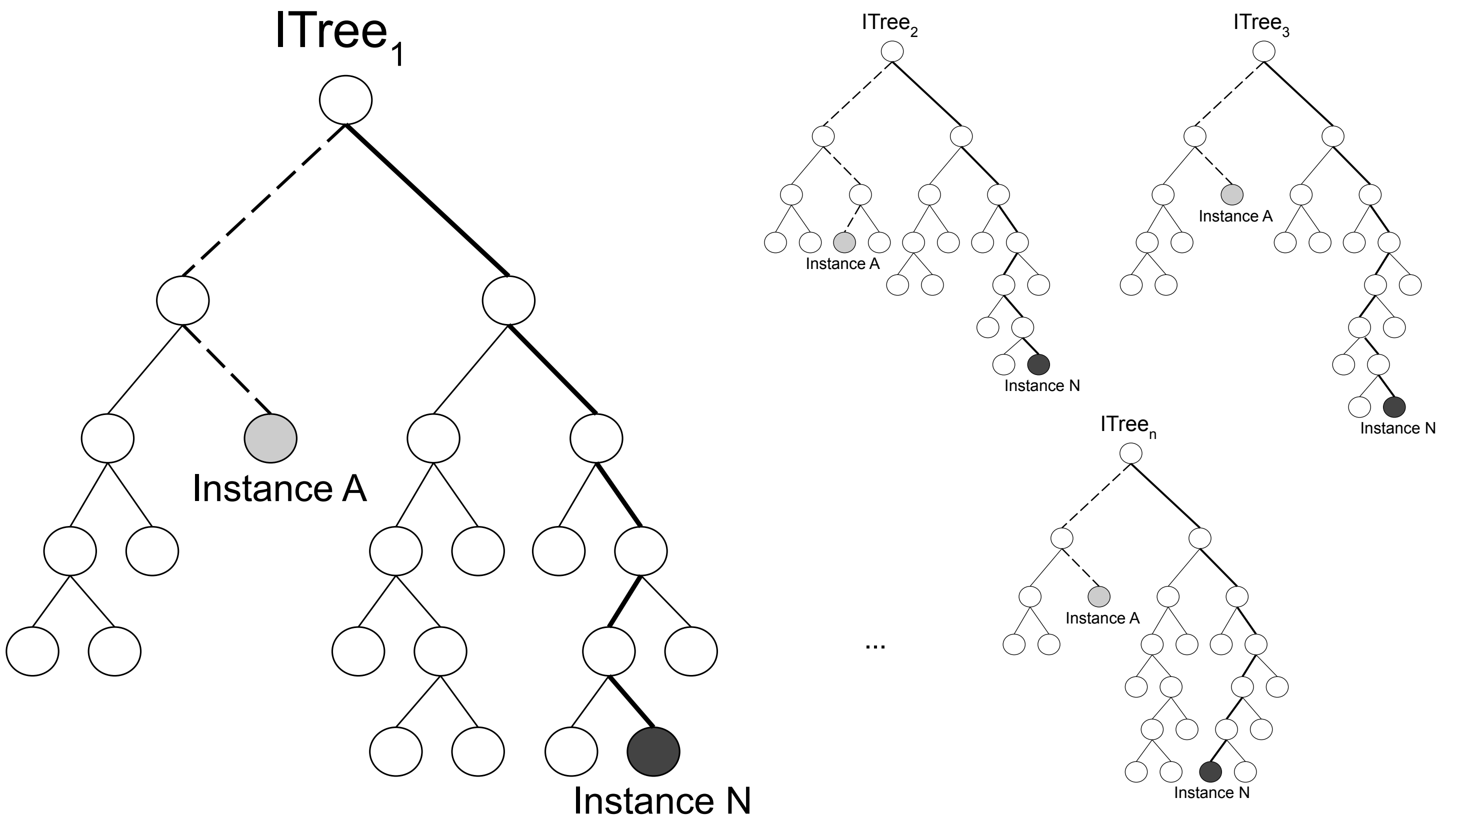

Isolation Forest [11]

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

Unsupervised

Univariate/Multivariate

Point/sequence

[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

Anomaly Detection methods: *an Example*

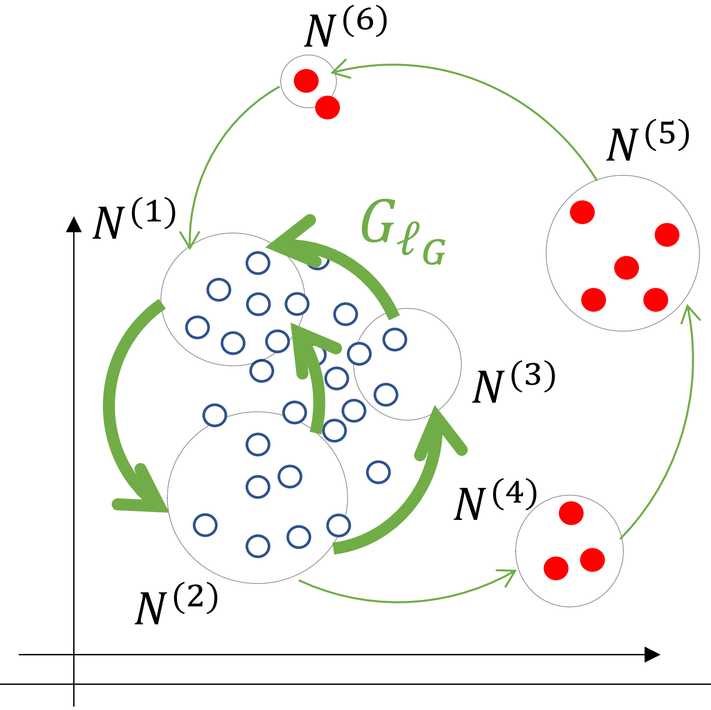

Each **node** is an ensemble of similar subsequences.

Each **edge** is associated to a weight w that corresponds to the number of times a subsequence move from one node to another.

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

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

Series2Graph [13]

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

Unsupervised

Univariate

subsequence

[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

Anomaly Detection methods: *an Example*

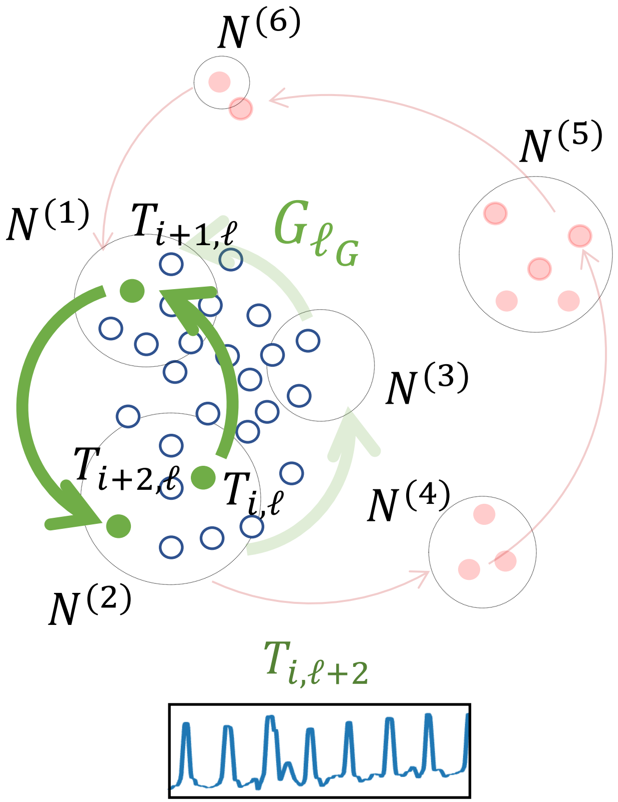

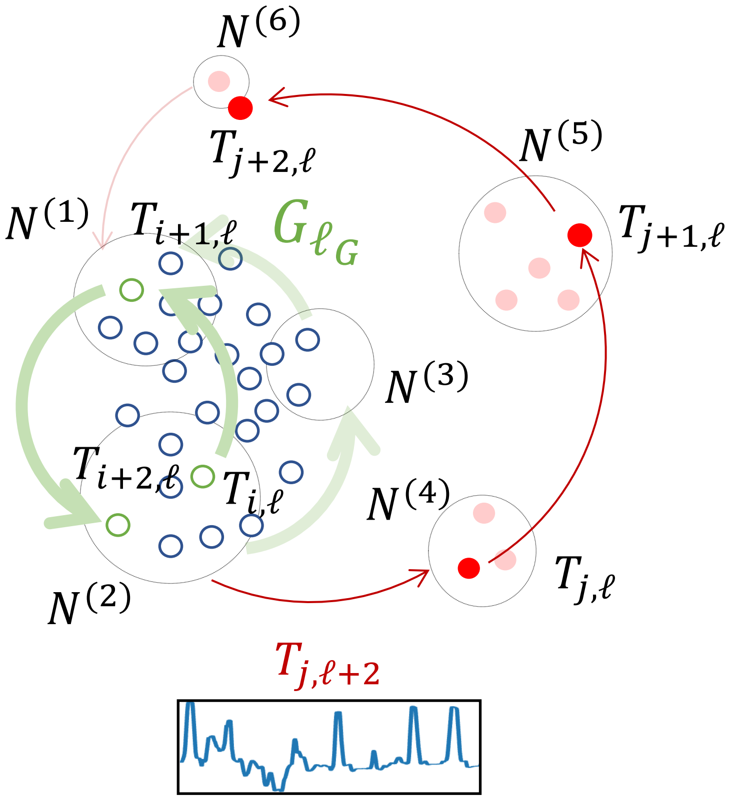

Series2Graph [13]

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

Unsupervised

Univariate

subsequence

[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

Anomaly Detection methods: *an Example*

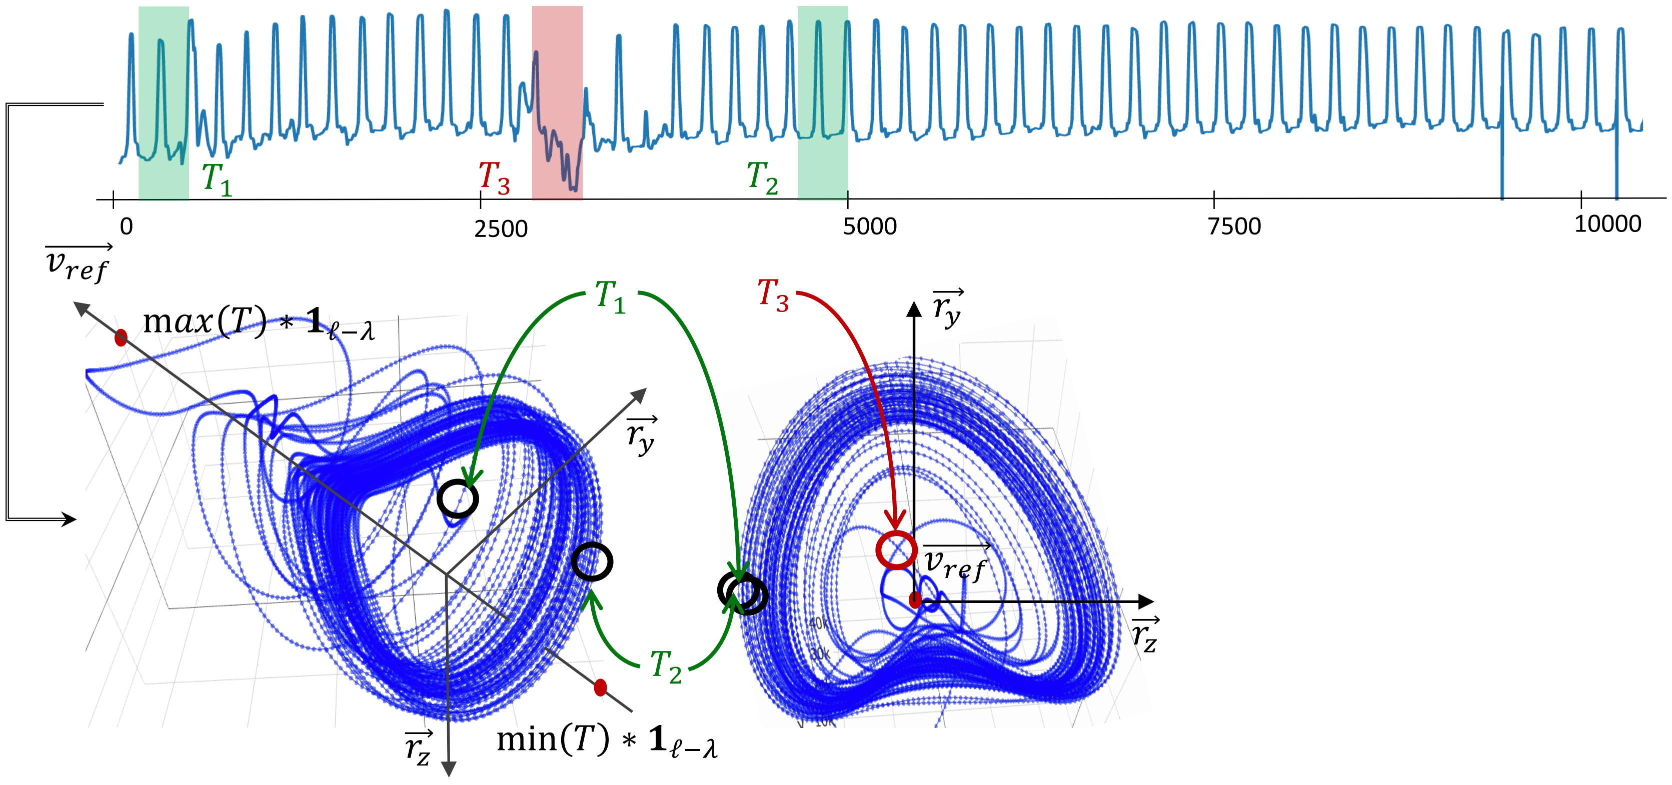

(1) subsequence embedding

Series2Graph [13]

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

Unsupervised

Univariate

subsequence

[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

Anomaly Detection methods: *an Example*

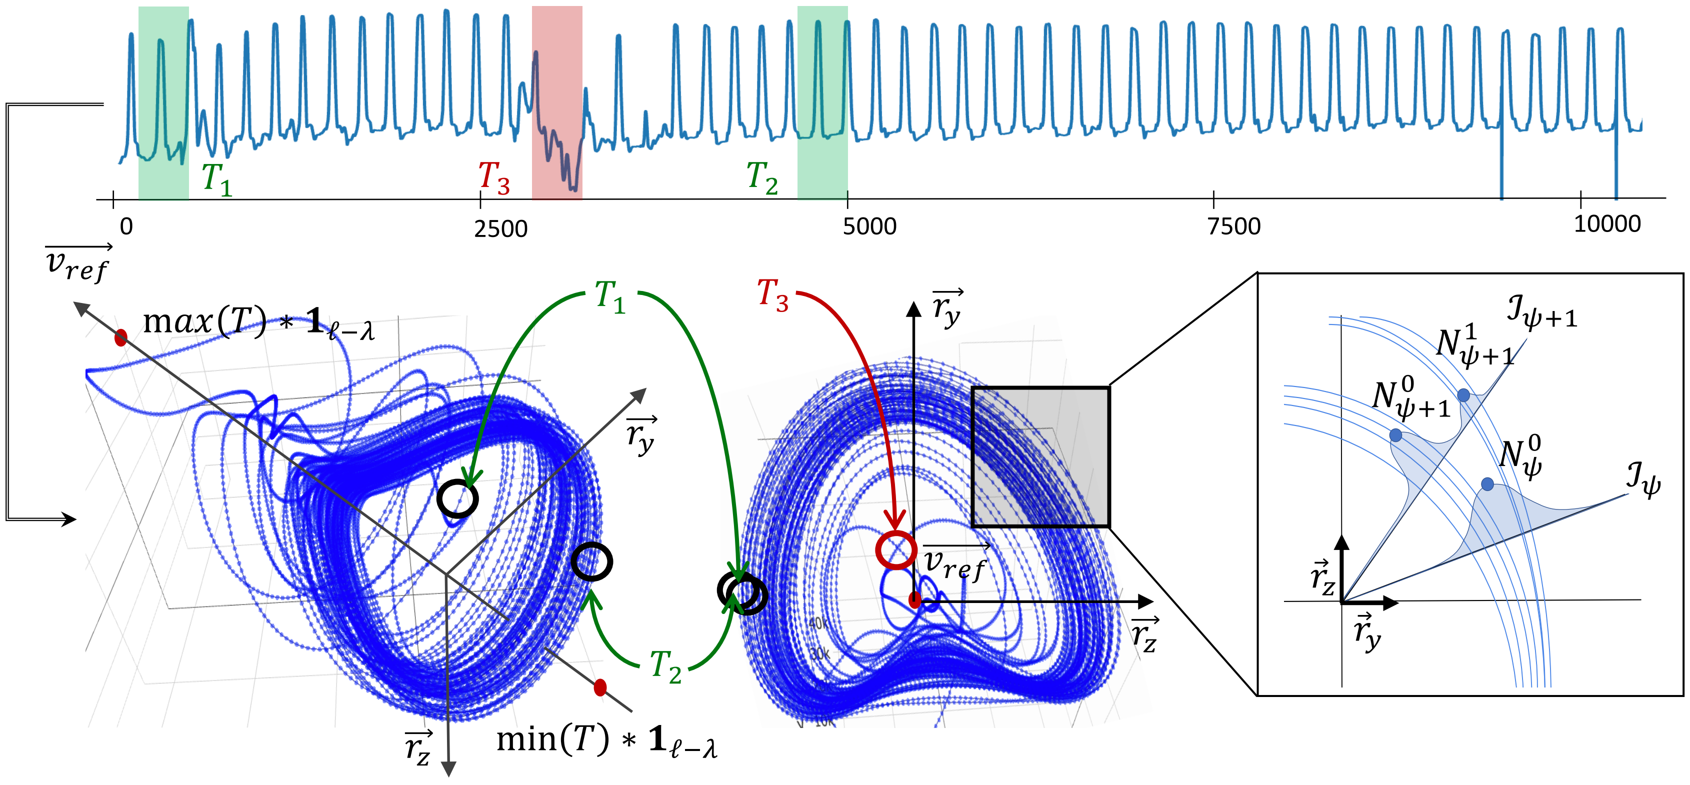

(1) subsequence embedding

(2) node creation

Series2Graph [13]

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

Unsupervised

Univariate

subsequence

[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

Anomaly Detection methods: *an Example*

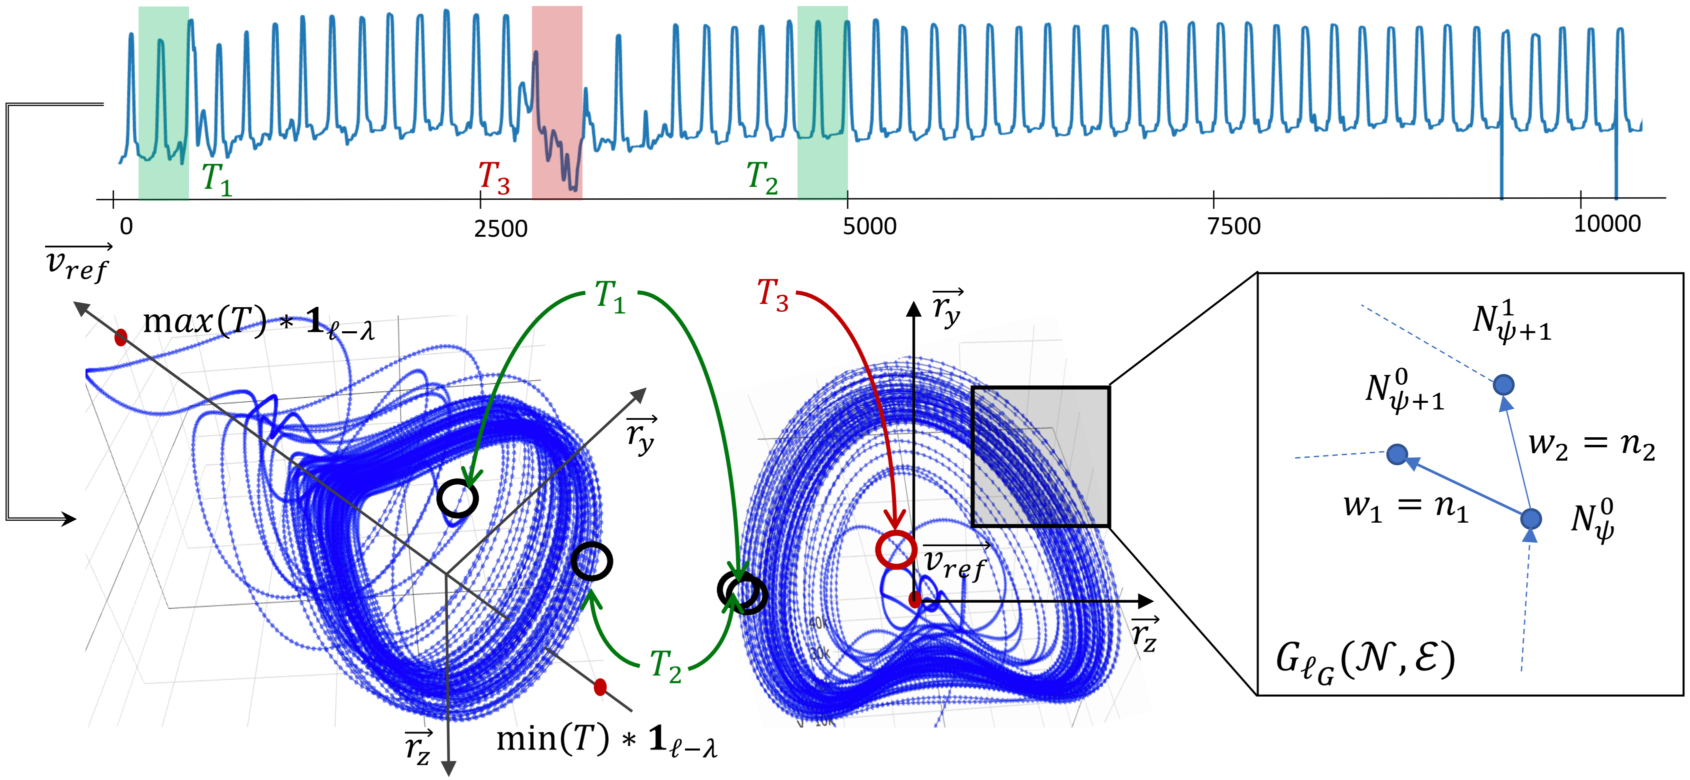

(1) subsequence embedding

(3) edge creation

Series2Graph [13]

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

Unsupervised

Univariate

subsequence

[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

Anomaly Detection methods: *an Example*

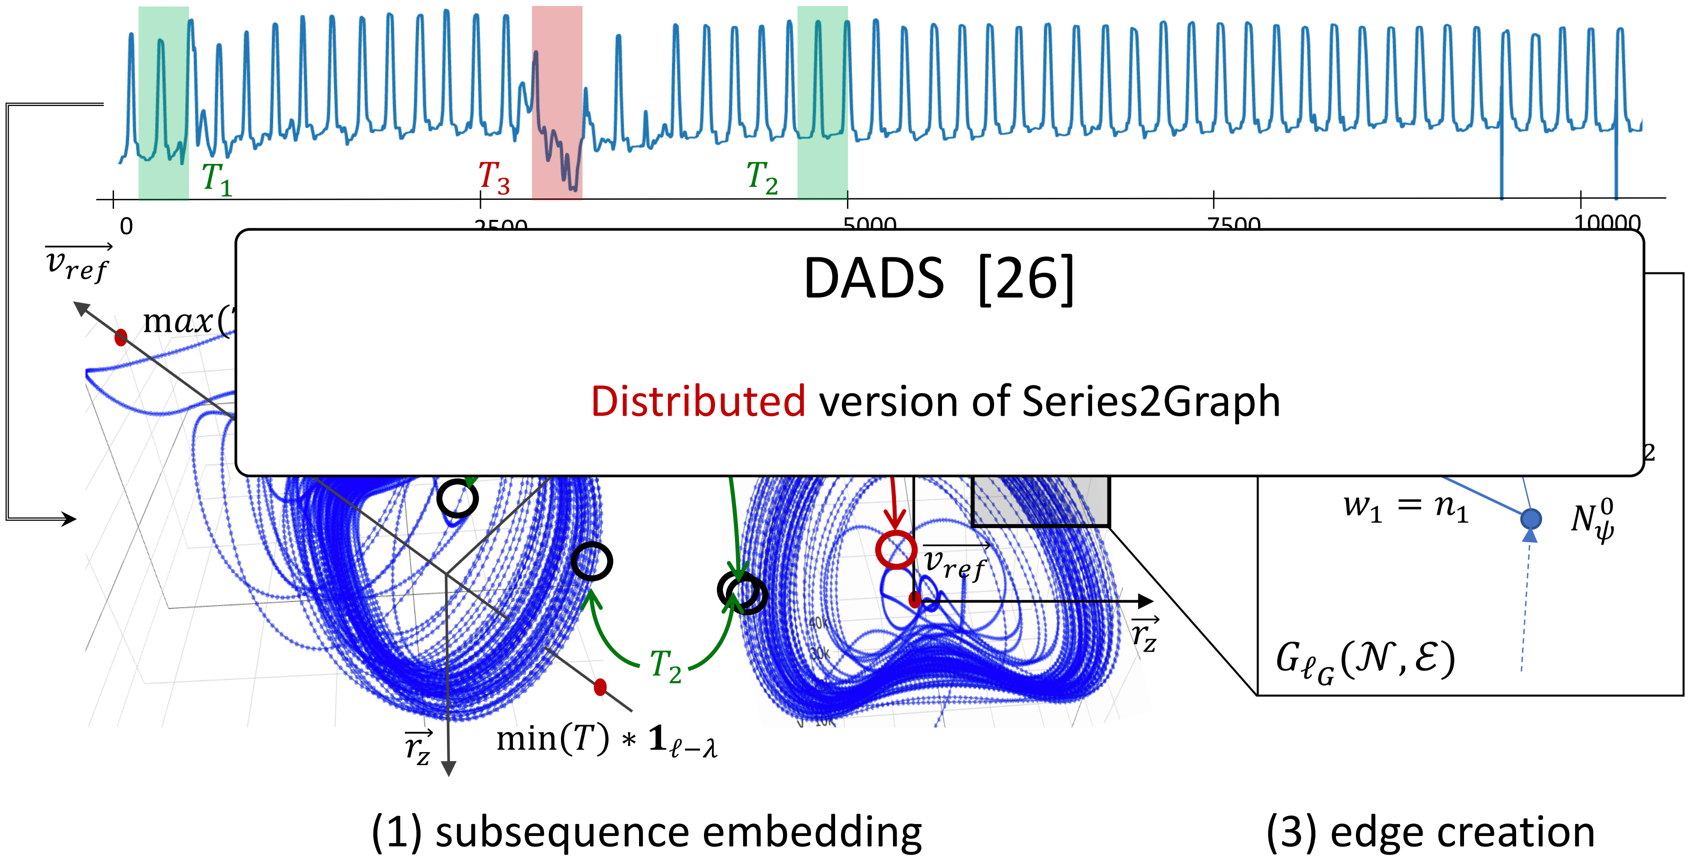

DADS [26]

Distributed version of Series2Graph

Series2Graph [13]

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

- Unsupervised
- Univariate
- subsequence

[26] Schneider, J., Wenig, P. & Papenbrock, T. Distributed detection of sequential anomalies in univariate time series. The VLDB Journal **30**, 579–602 (2021).

Anomaly Detection methods: *an Example*

Snippet of SED time series [14]

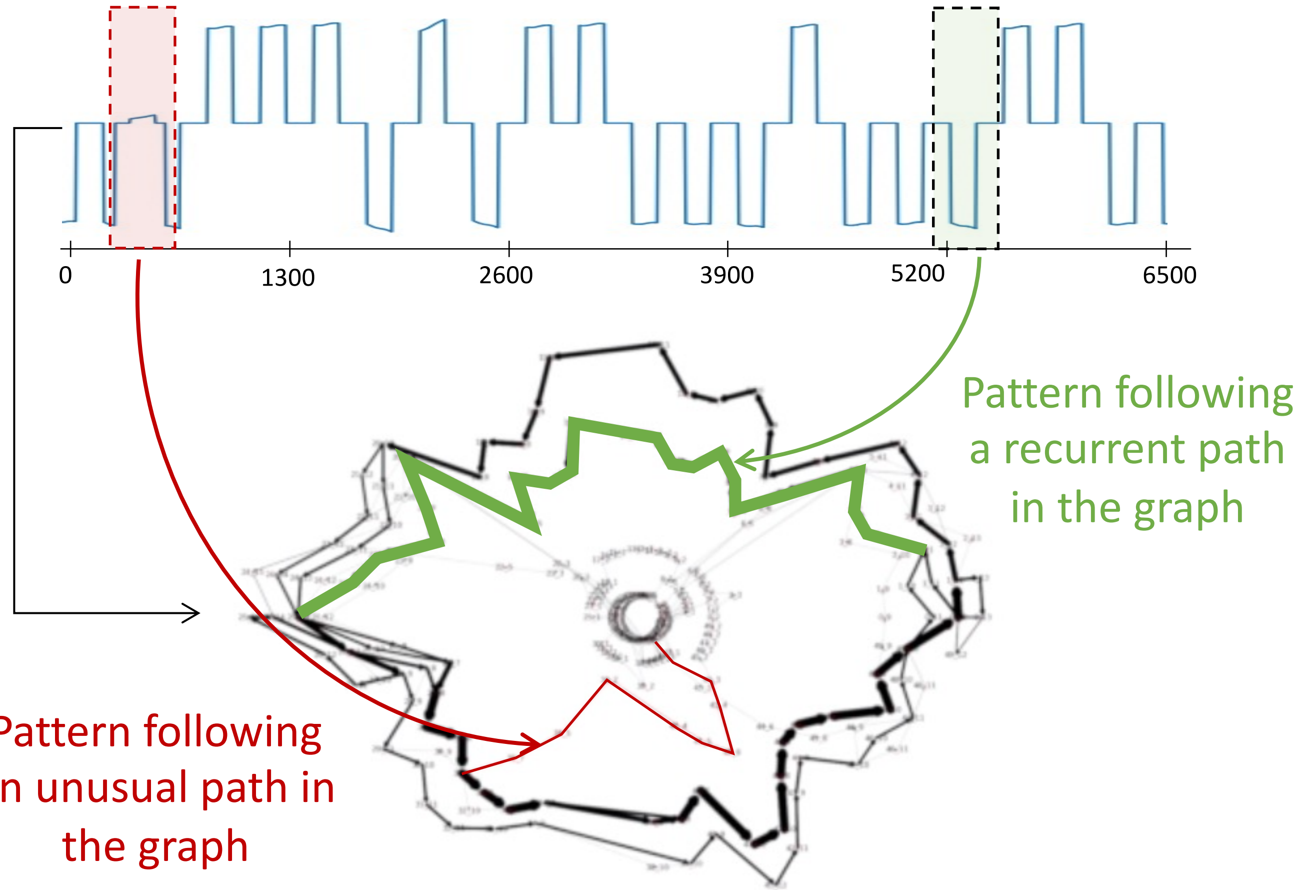

Series2Graph [13]

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

Unsupervised

Univariate

subsequence

[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

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.

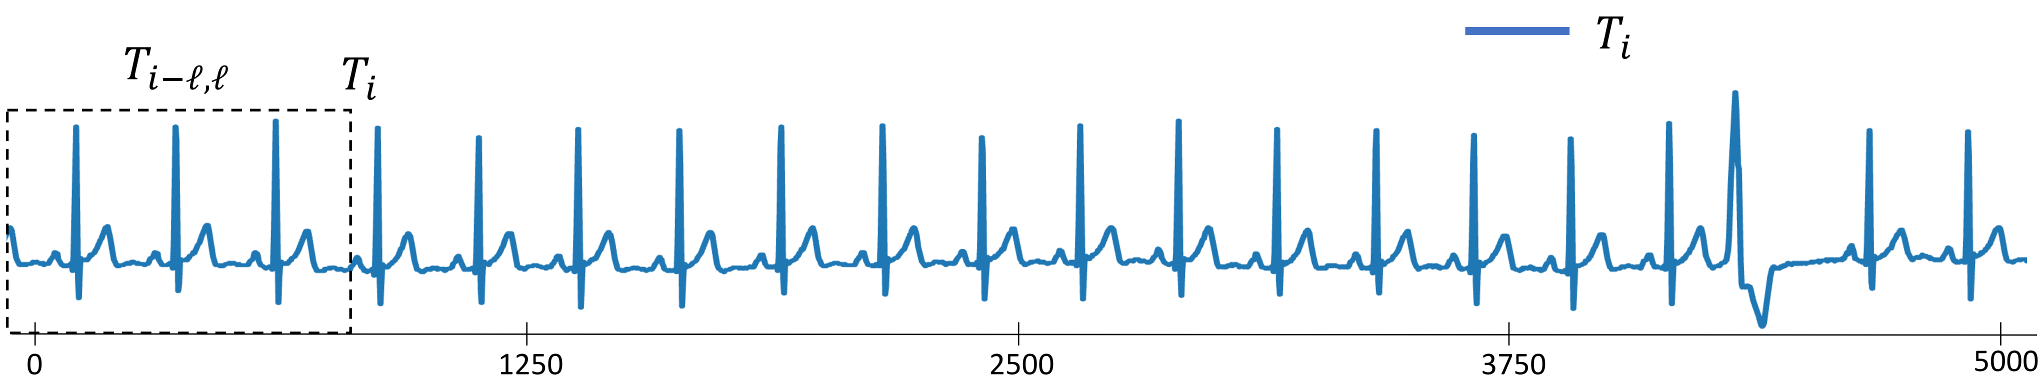

Anomaly Detection methods: *Forecasting-based*

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.

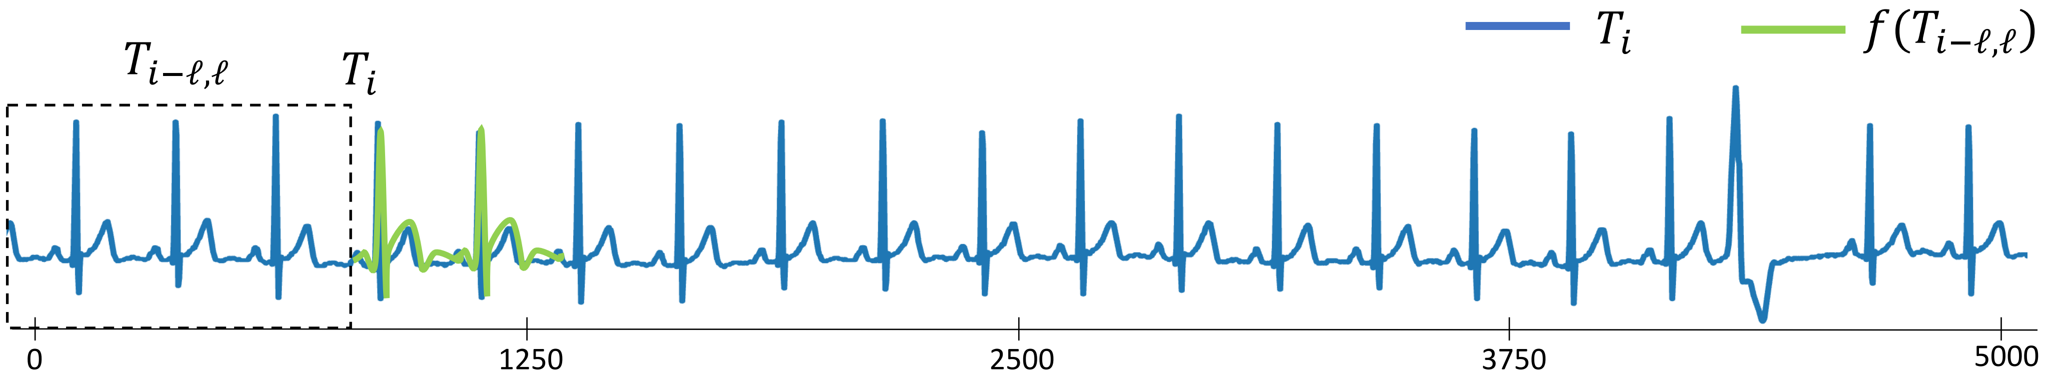

Anomaly Detection methods: *Forecasting-based*

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.

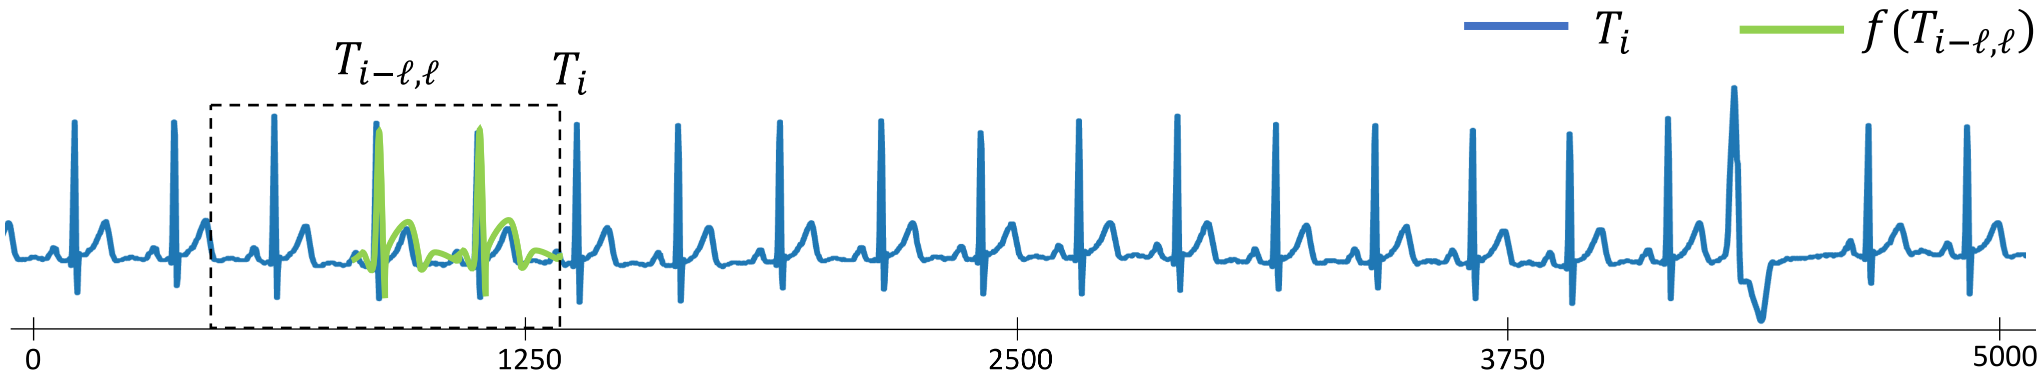

Anomaly Detection methods: *Forecasting-based*

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.

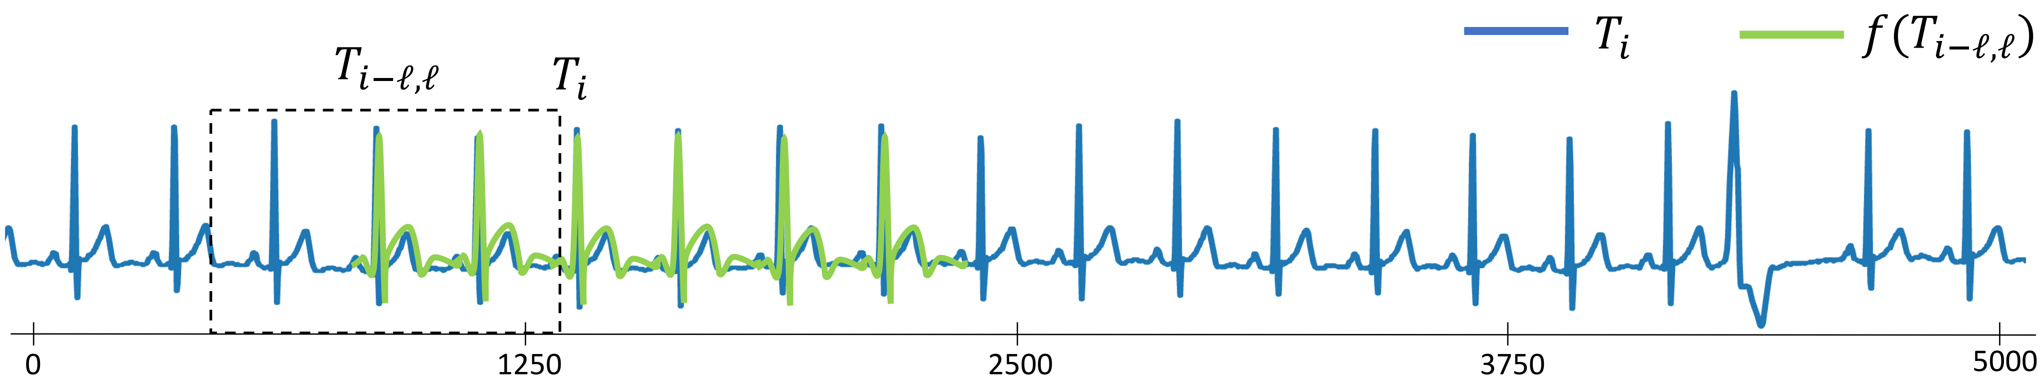

Anomaly Detection methods: *Forecasting-based*

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.

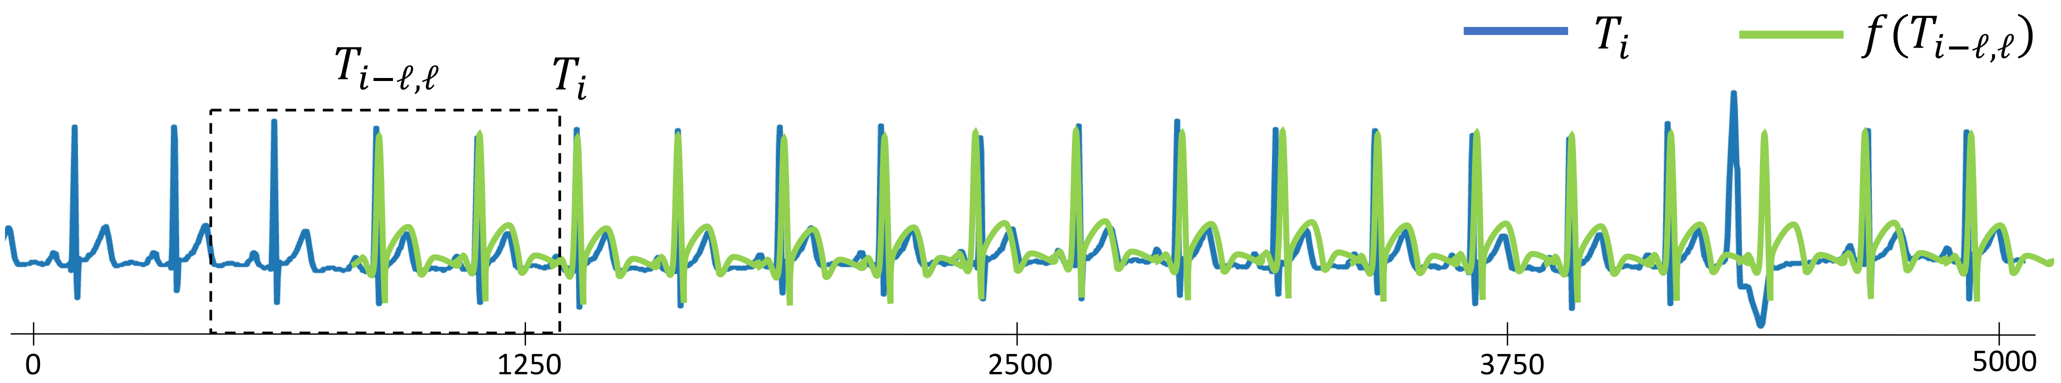

Anomaly Detection methods: *Forecasting-based*

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.

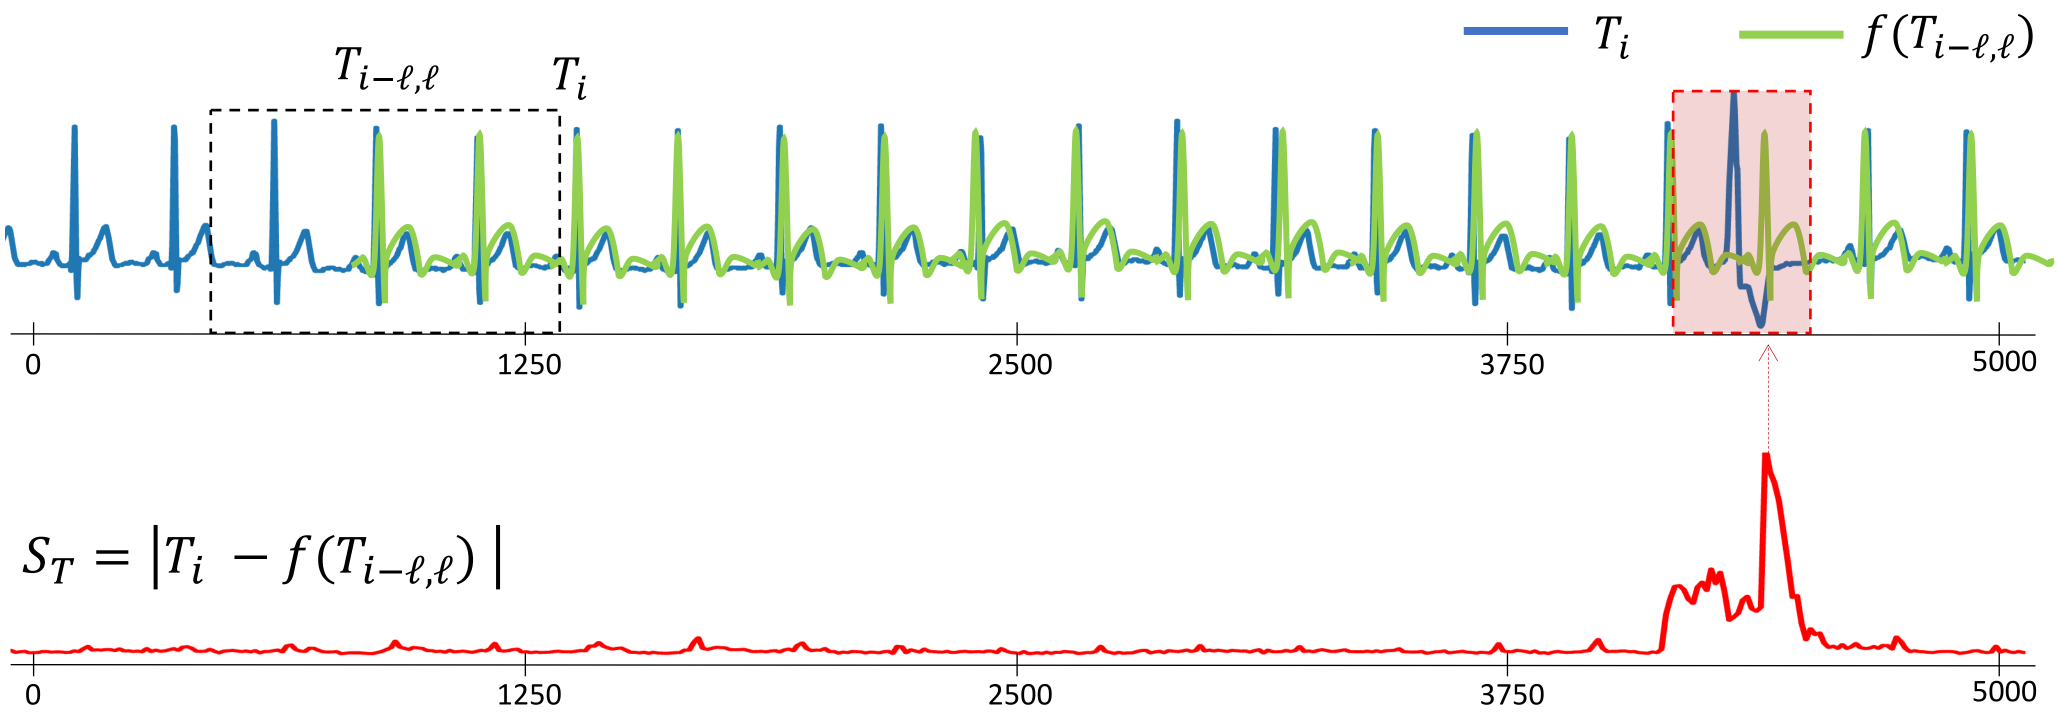

$$S_T = \left|T_i - f(T_{i-\ell,\ell})\right|$$

Anomaly Detection methods: *an Example*

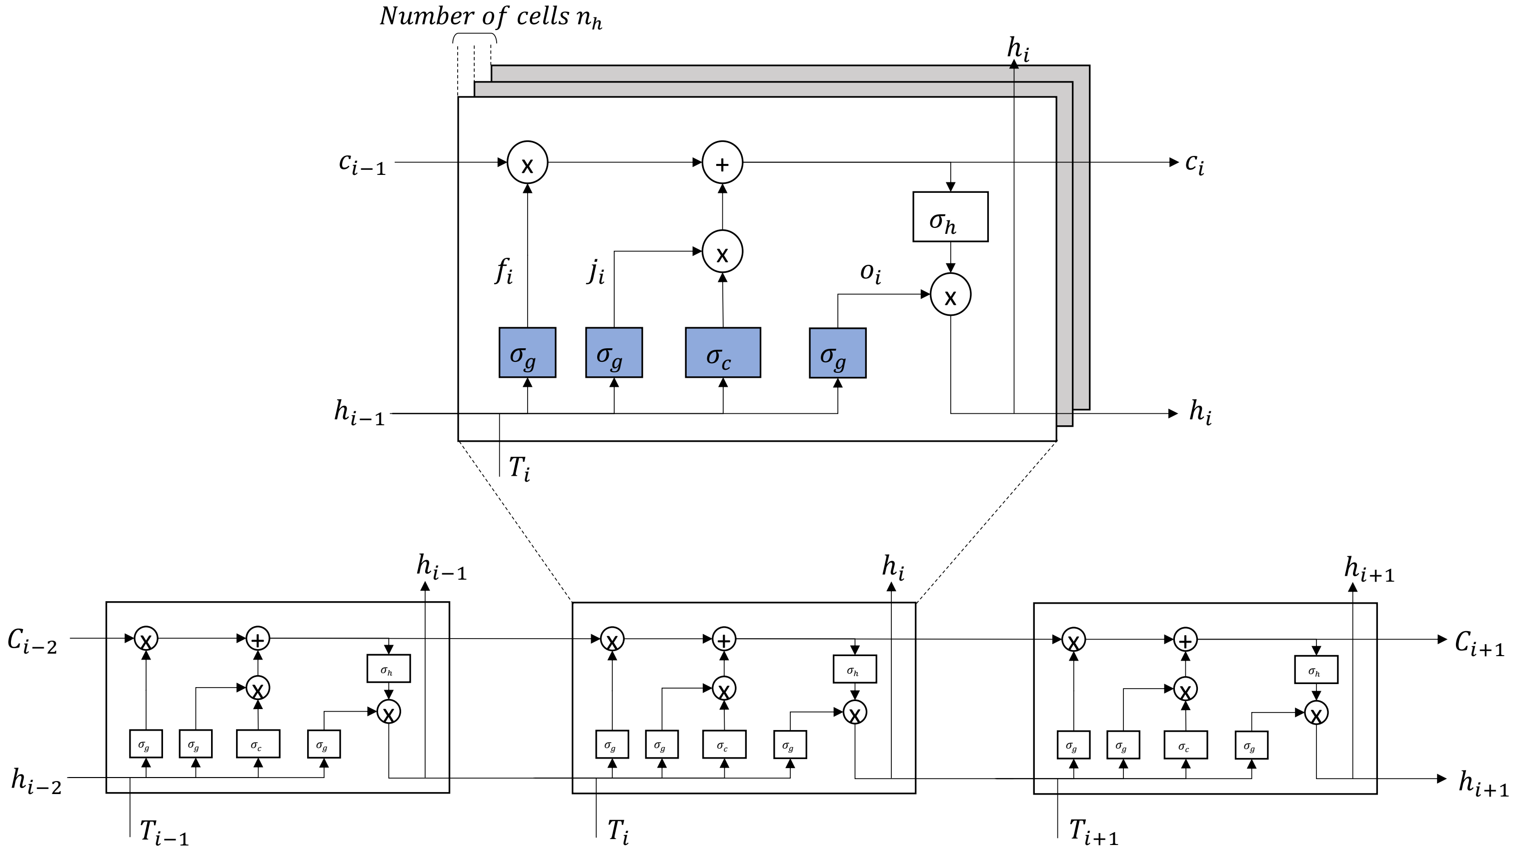

LSTM-AD [15]

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

Semi-supervised

Univariate/Multivariate

Point/sequence

[15] Pankaj Malhotra, Lovekesh Vig, Gautam Shro, and Puneet Agarwal. 2015. Long Short Term Memory Networks for Anomaly Detection in Time Series. (2015).

Anomaly Detection methods: *an Example*

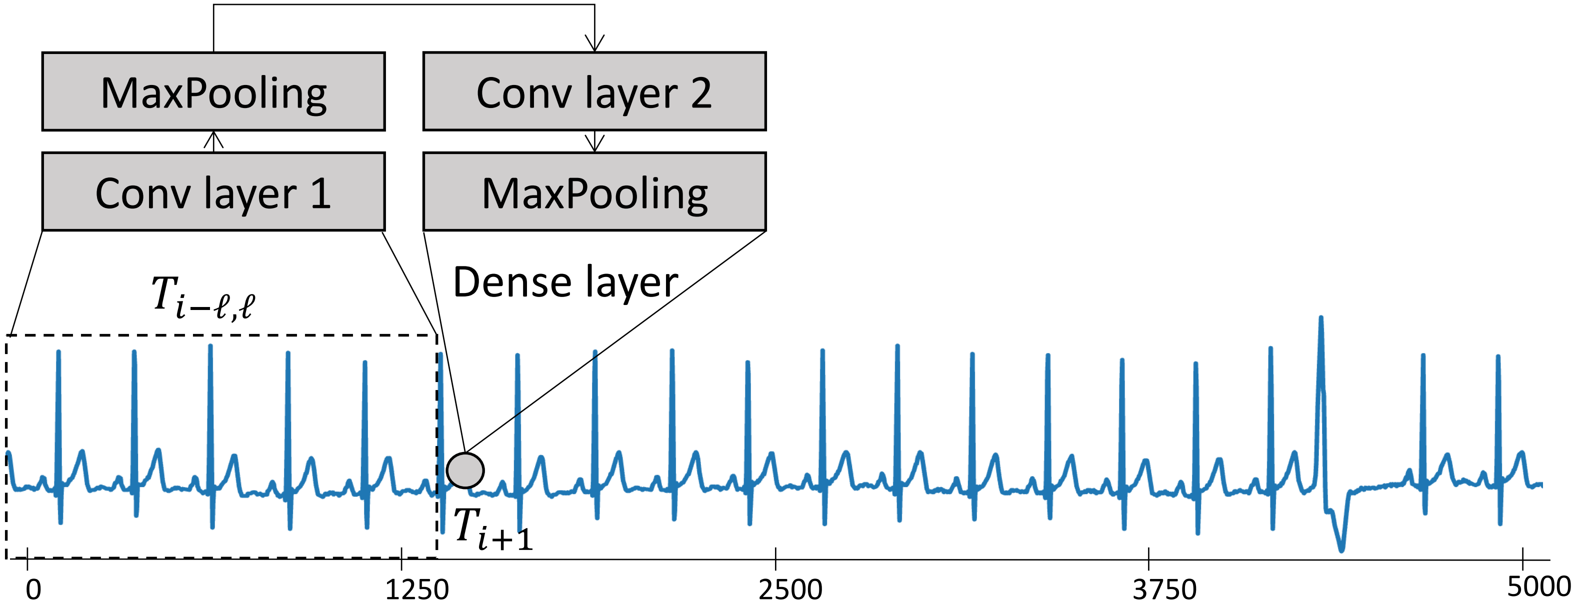

DeepAnT [16] (CNN)

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

Semi-supervised

Univariate/Multivariate

Point/sequence

[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.

Anomaly Detection methods: *an Example*

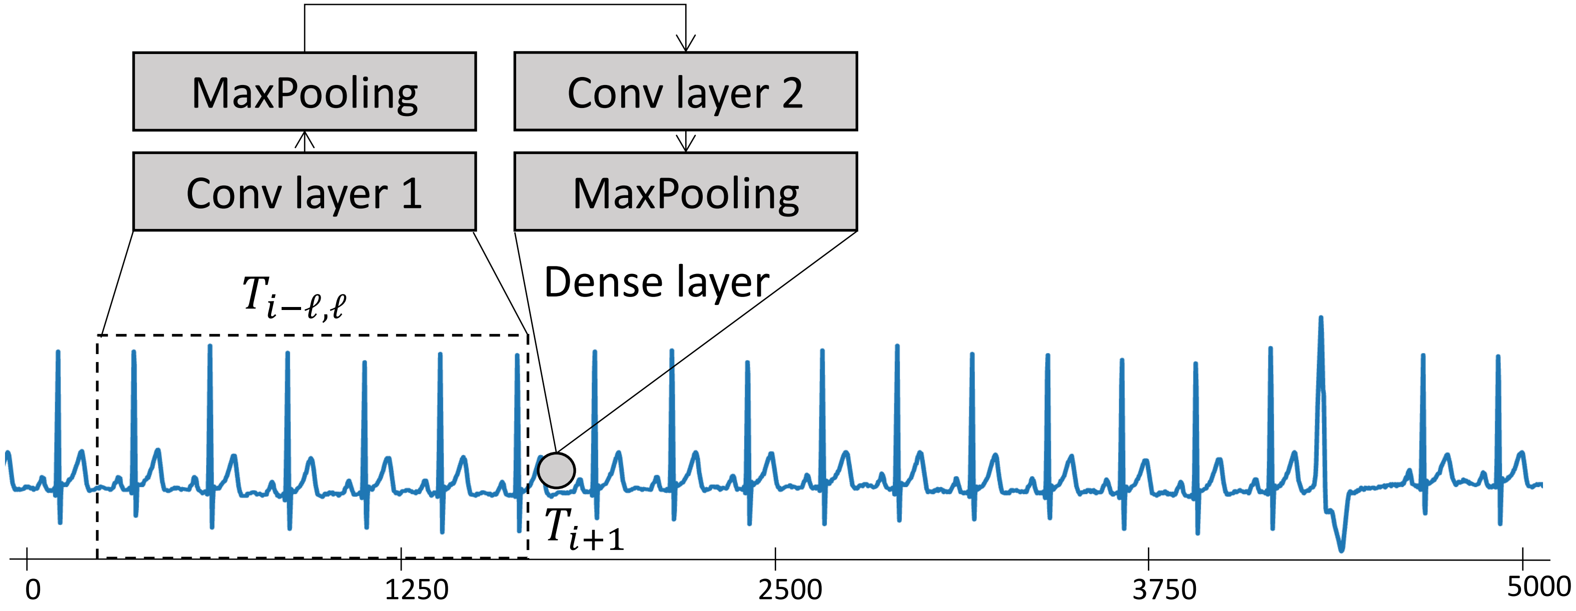

DeepAnT [16] (CNN)

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

- Semi-supervised
- Univariate/Multivariate
- Point/sequence

[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.

Anomaly Detection methods: *an Example*

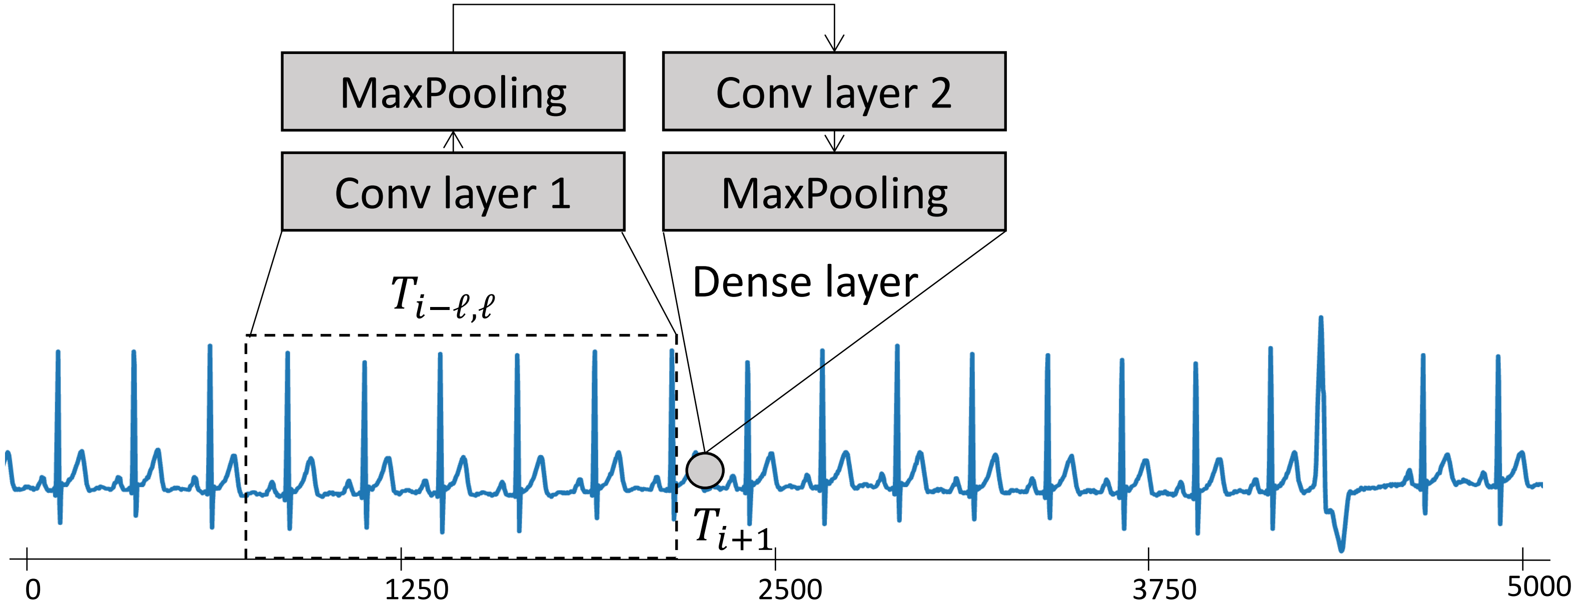

DeepAnT [16] (CNN)

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

Semi-supervised

Univariate/Multivariate

Point/sequence

[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.

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.

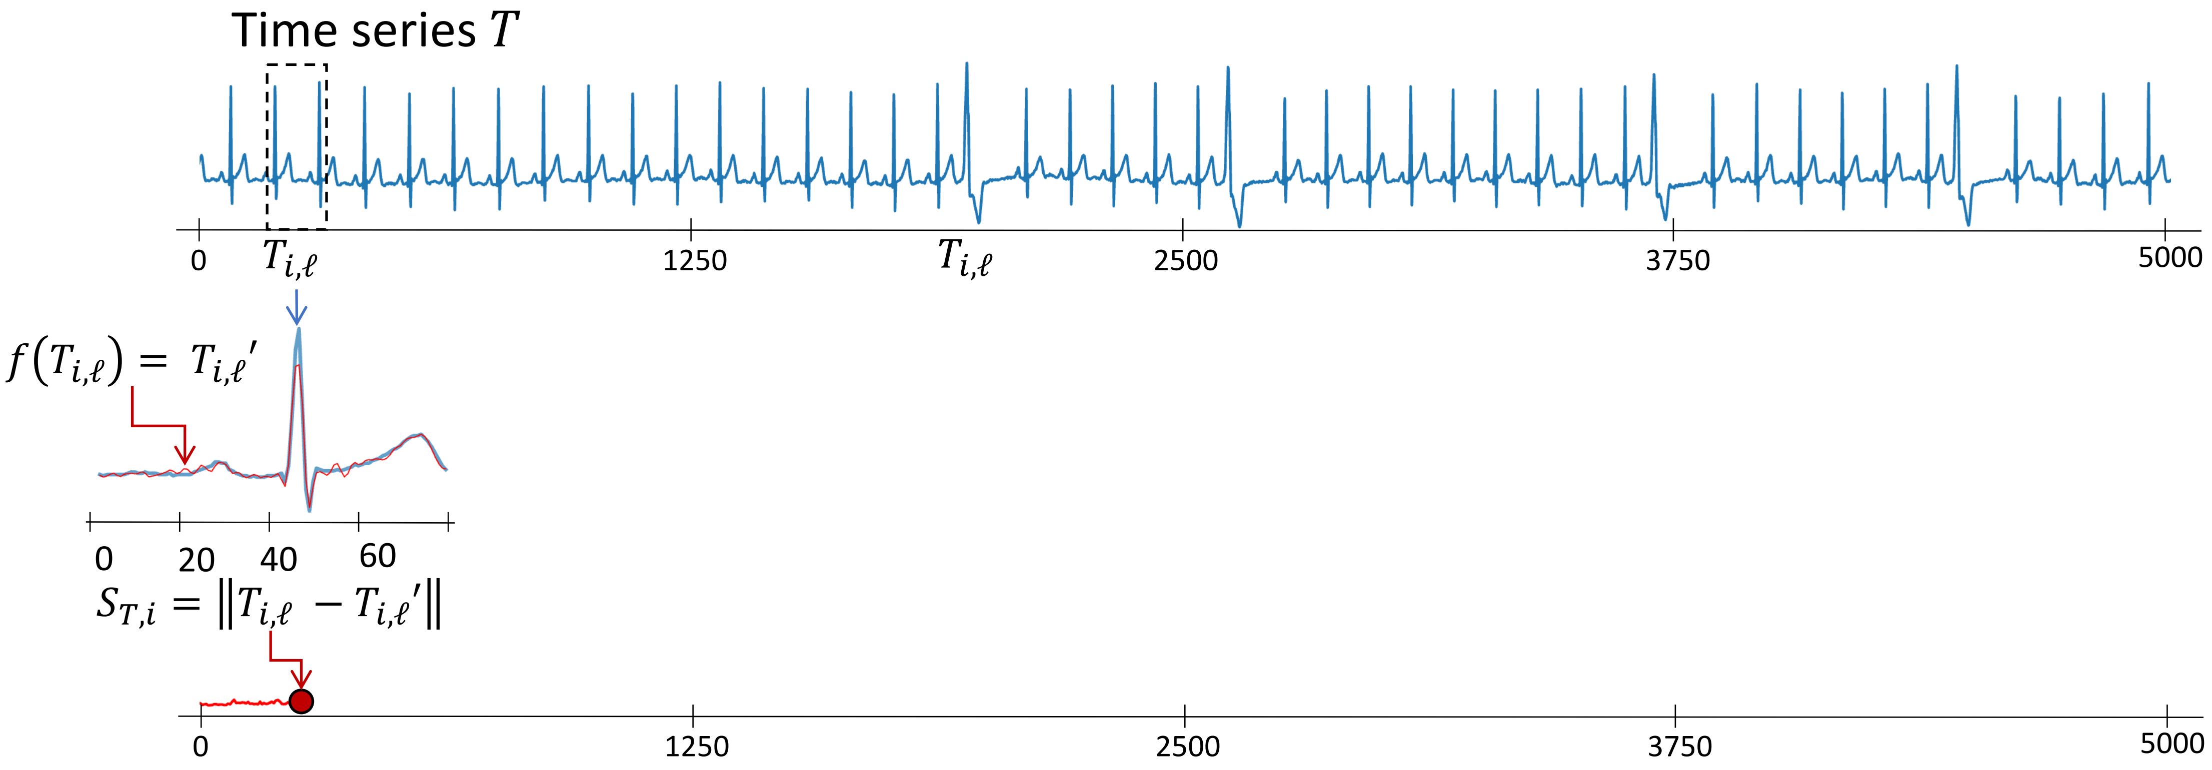

Anomaly Detection methods: *Reconstruction-based*

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.

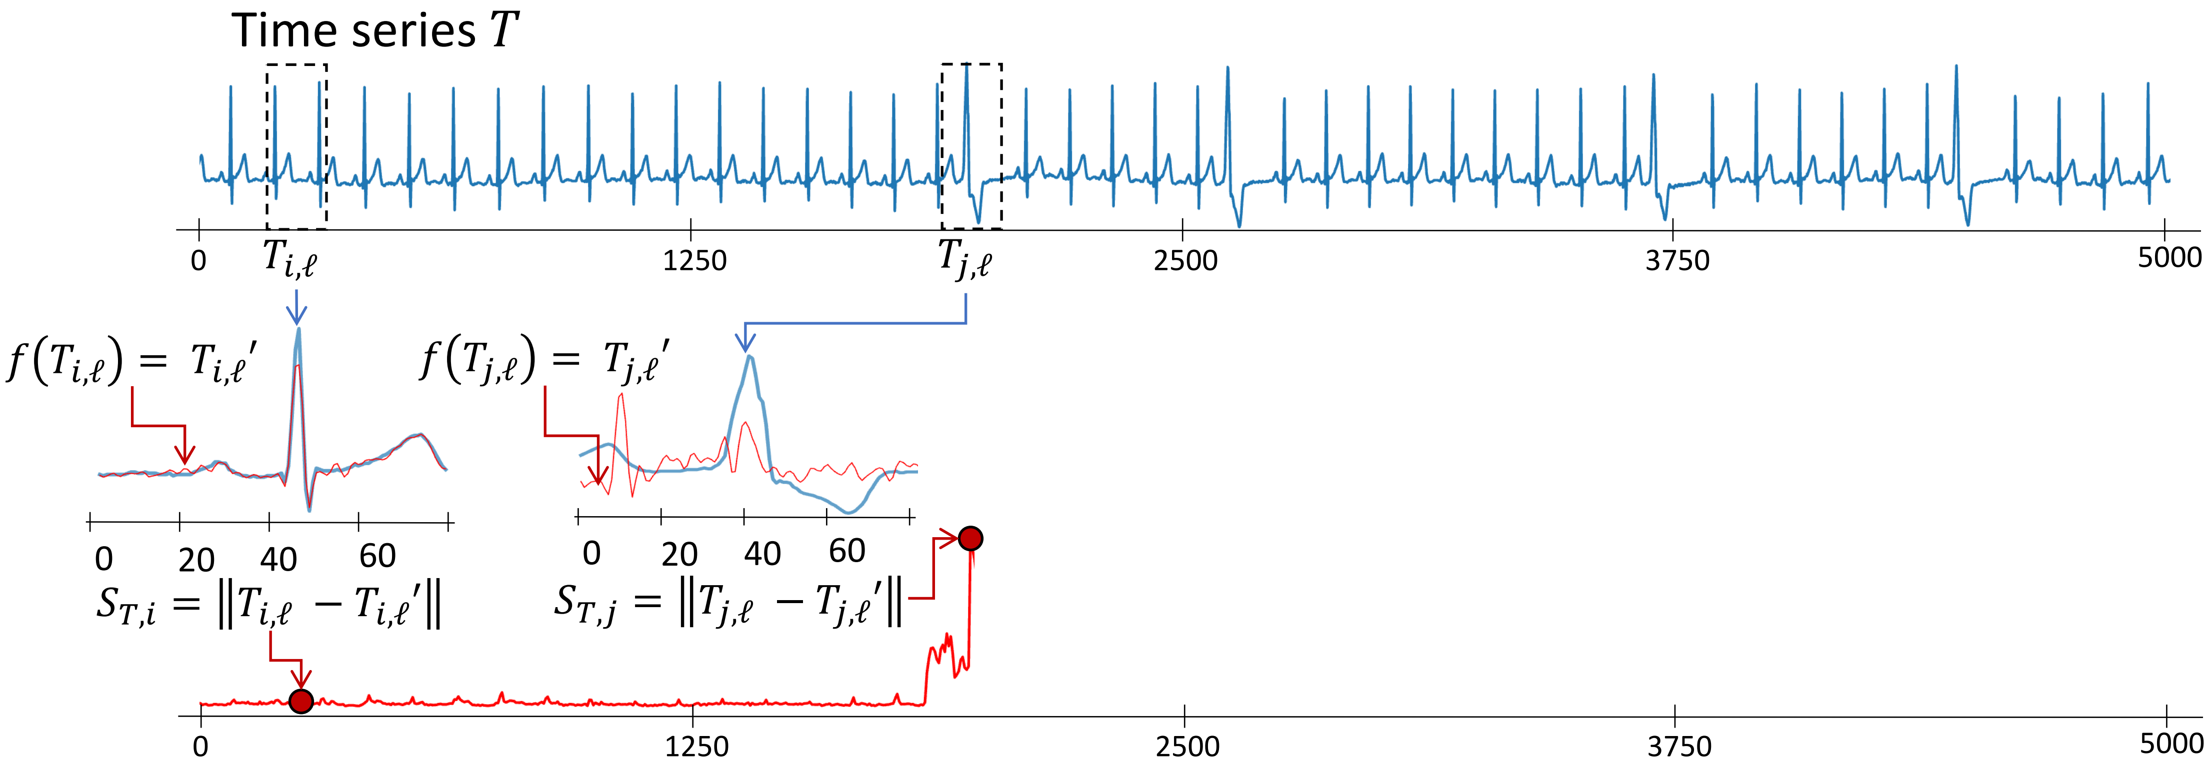

Anomaly Detection methods: *Reconstruction-based*

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.

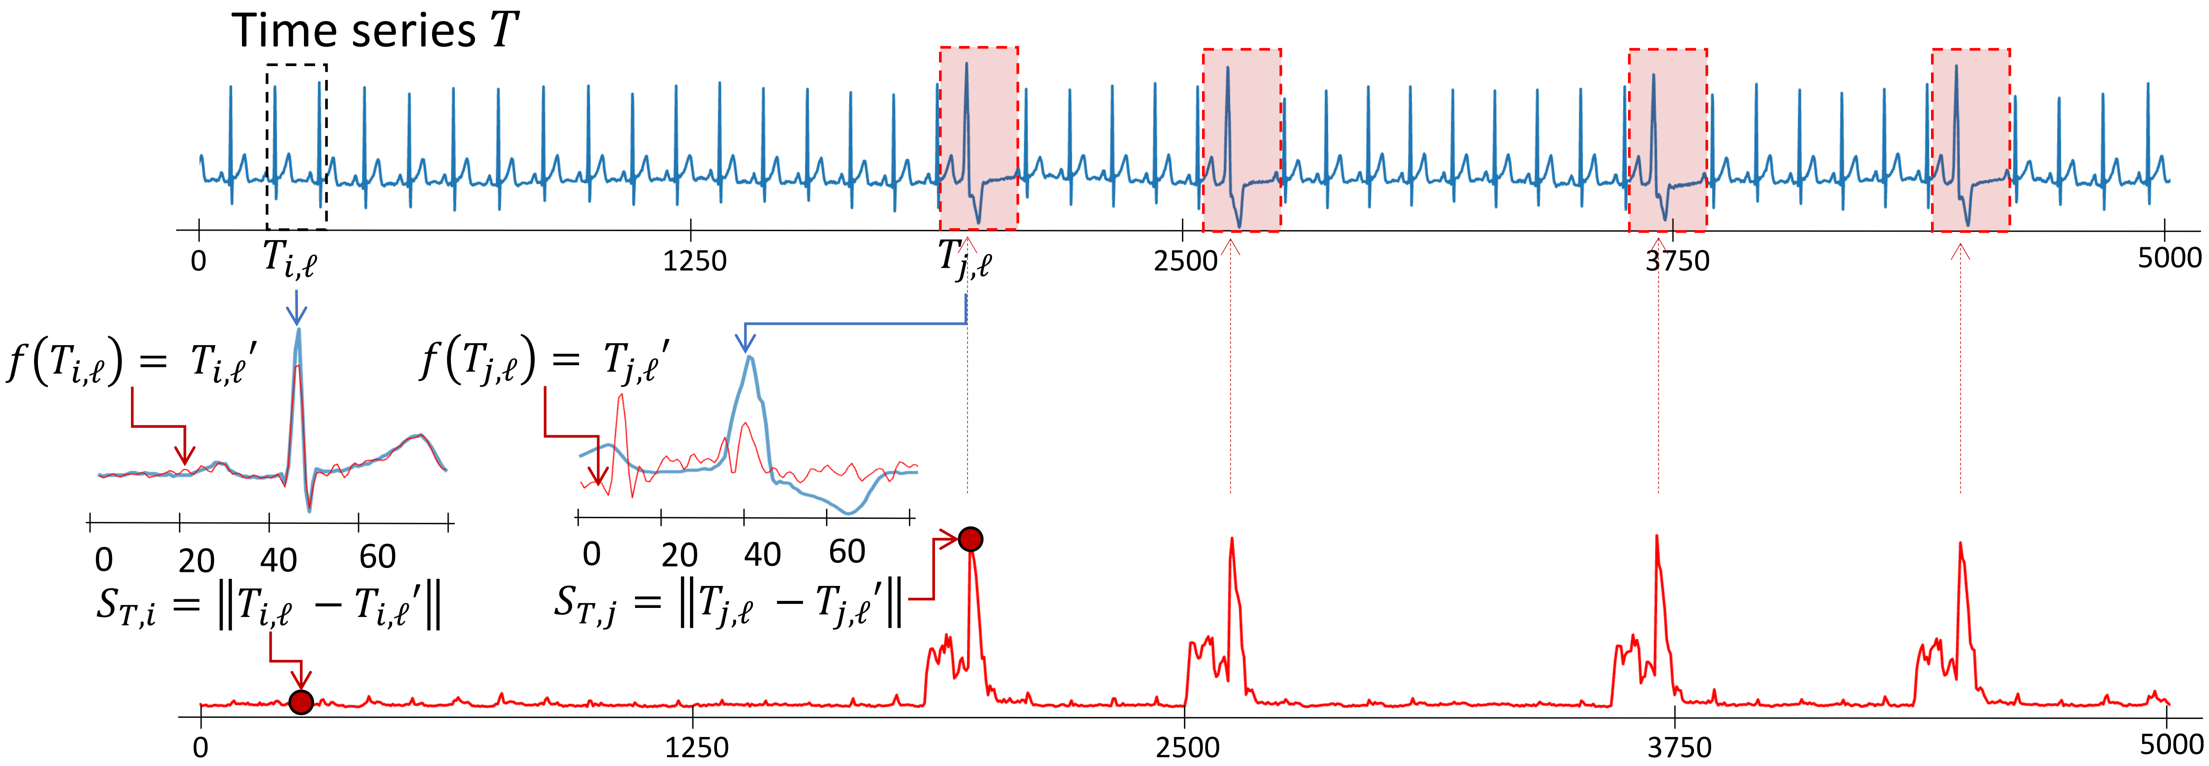

Anomaly Detection methods: *an Example*

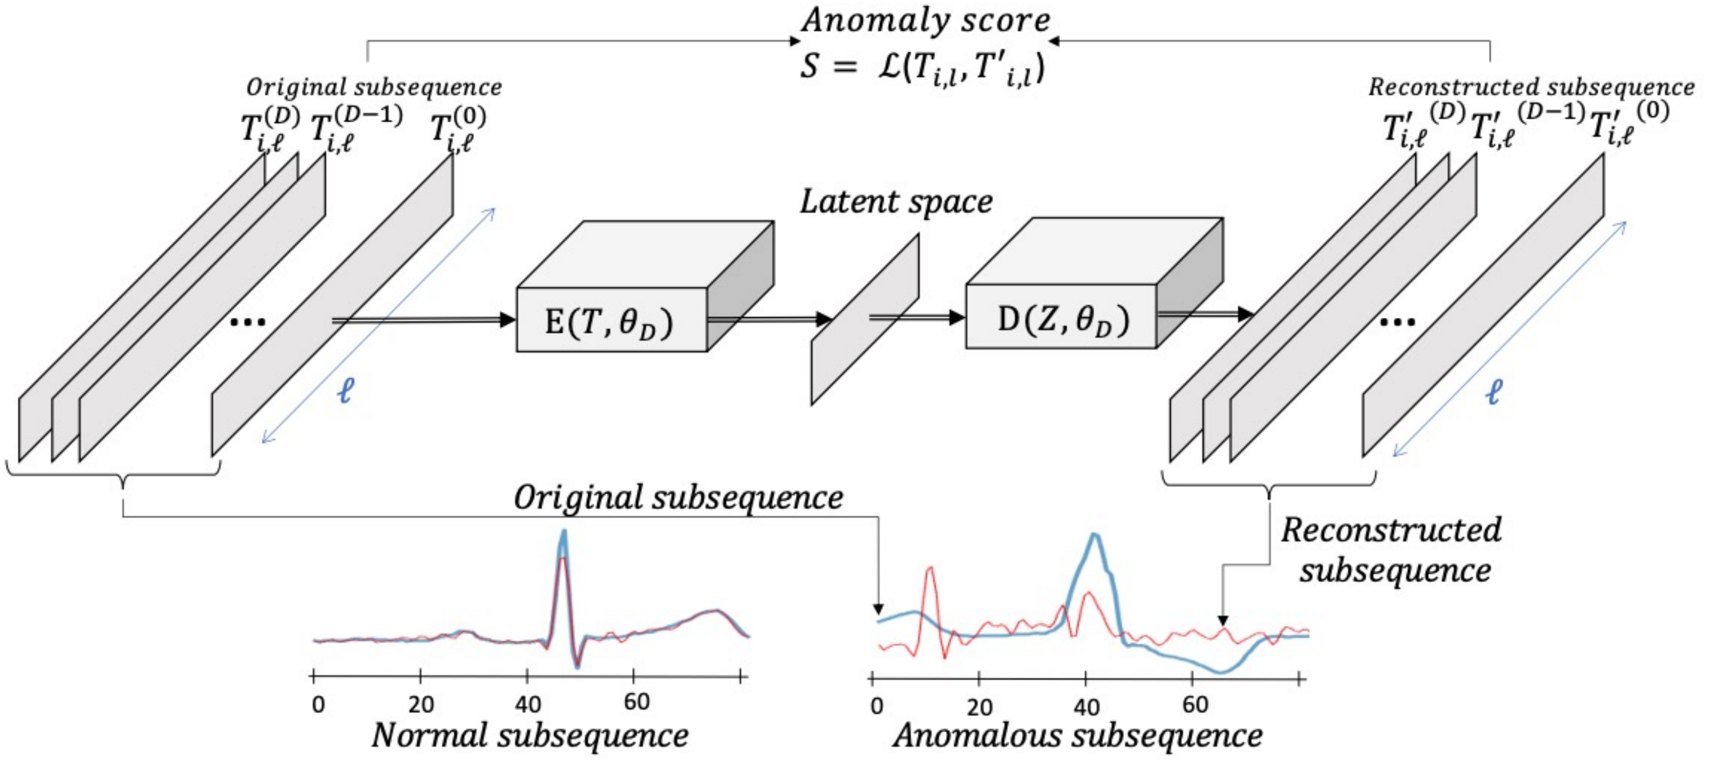

AutoEncoders [17] (AE)

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

[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).

Anomaly Detection methods: *Existing benchmark*

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

TimeEval [5]

Real datasets collection

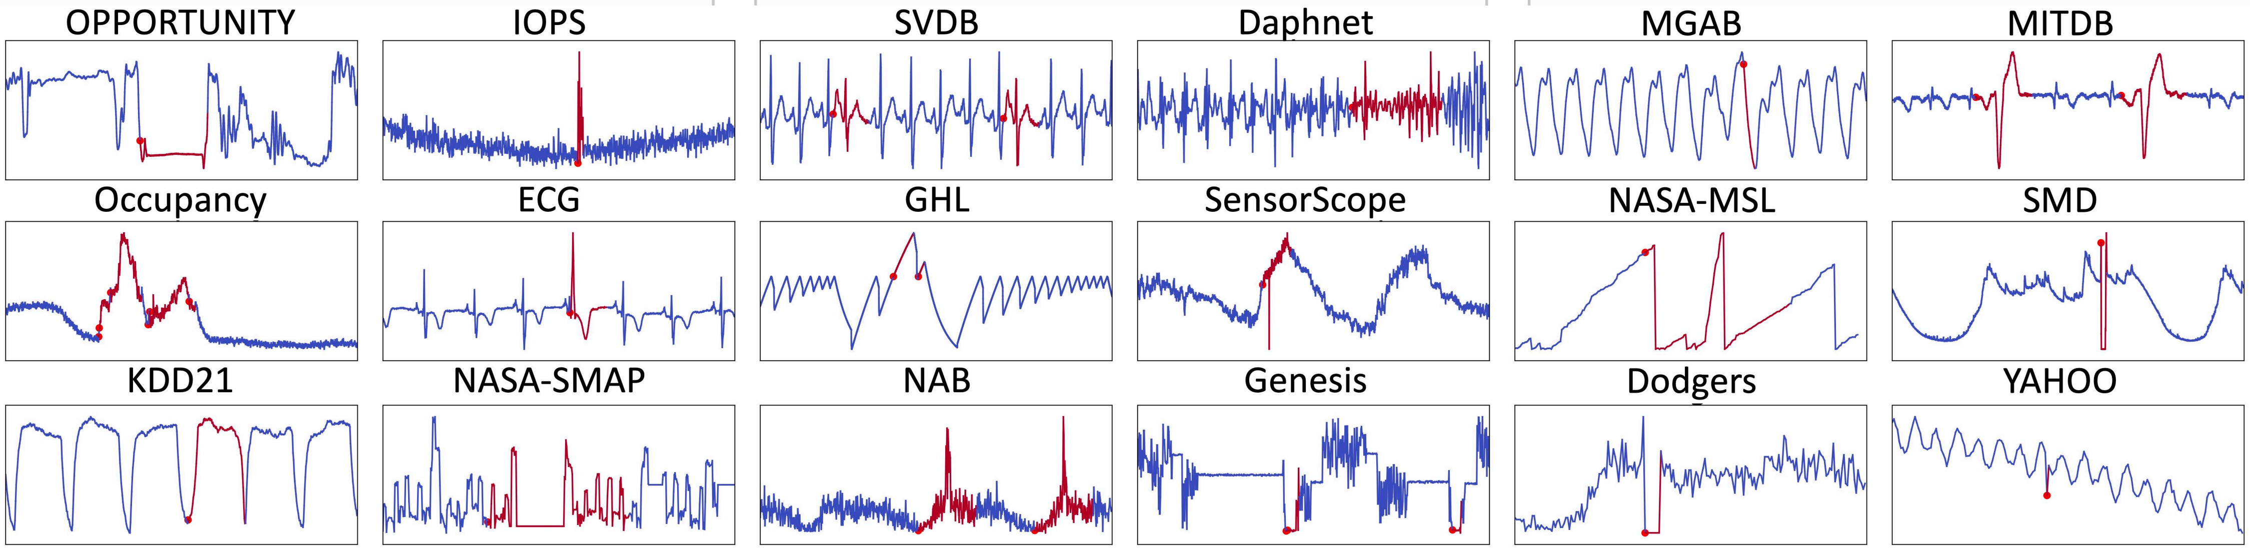

only 1 labeled anomaly

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

contamination, size or label 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

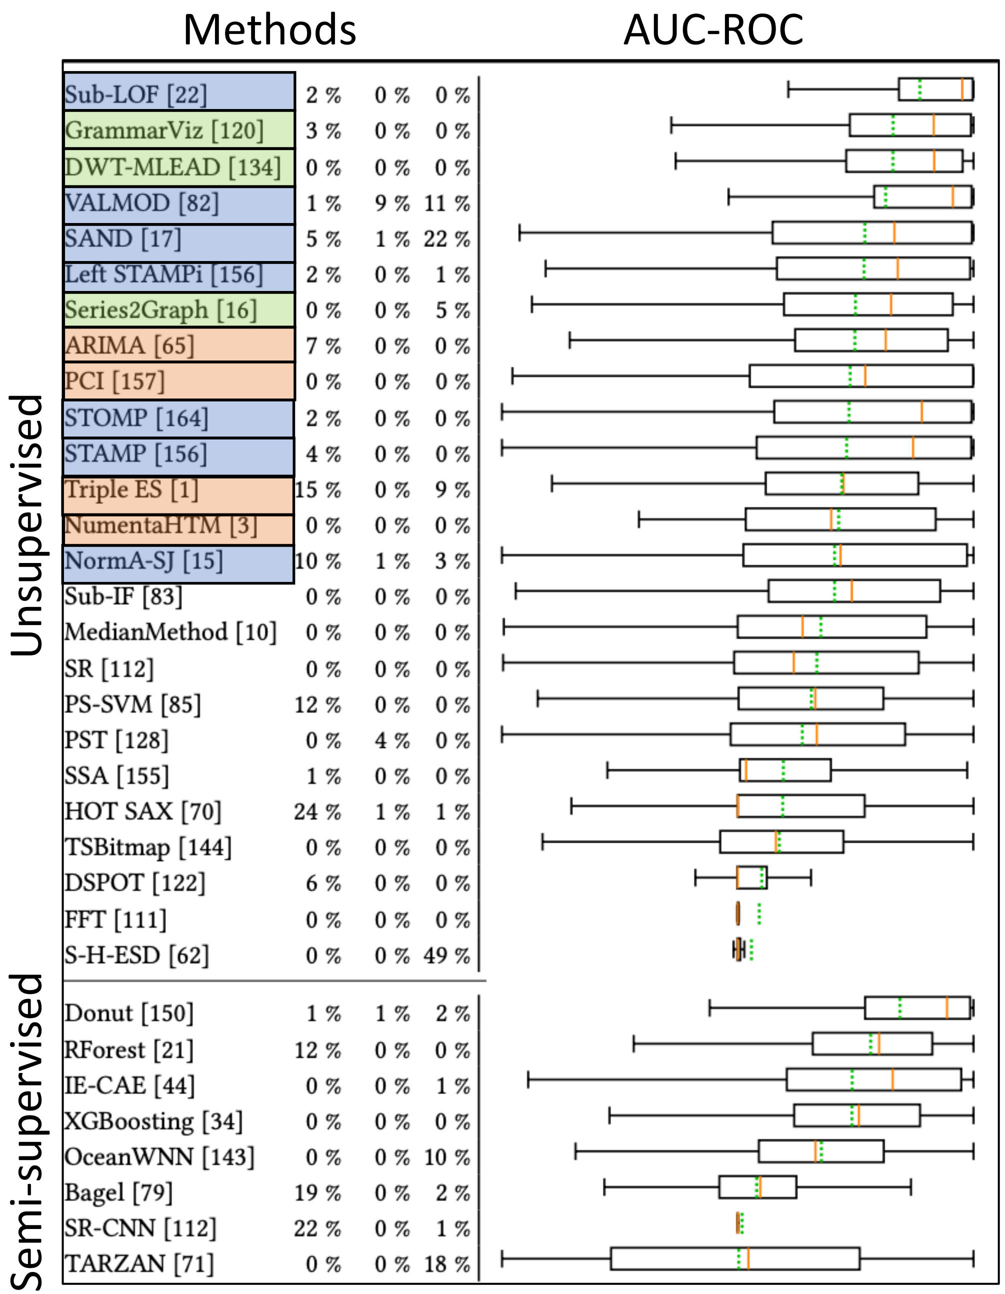

[5] Sebastian Schmidl, Phillip Wenig, and Thorsten Papenbrock. 2022. Anomaly detection in time series: a comprehensive evaluation. Proc. VLDB Endow. 15, 9 (May 2022), 1779–1797.

Anomaly Detection methods:
Experimental evaluation

Observations on 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

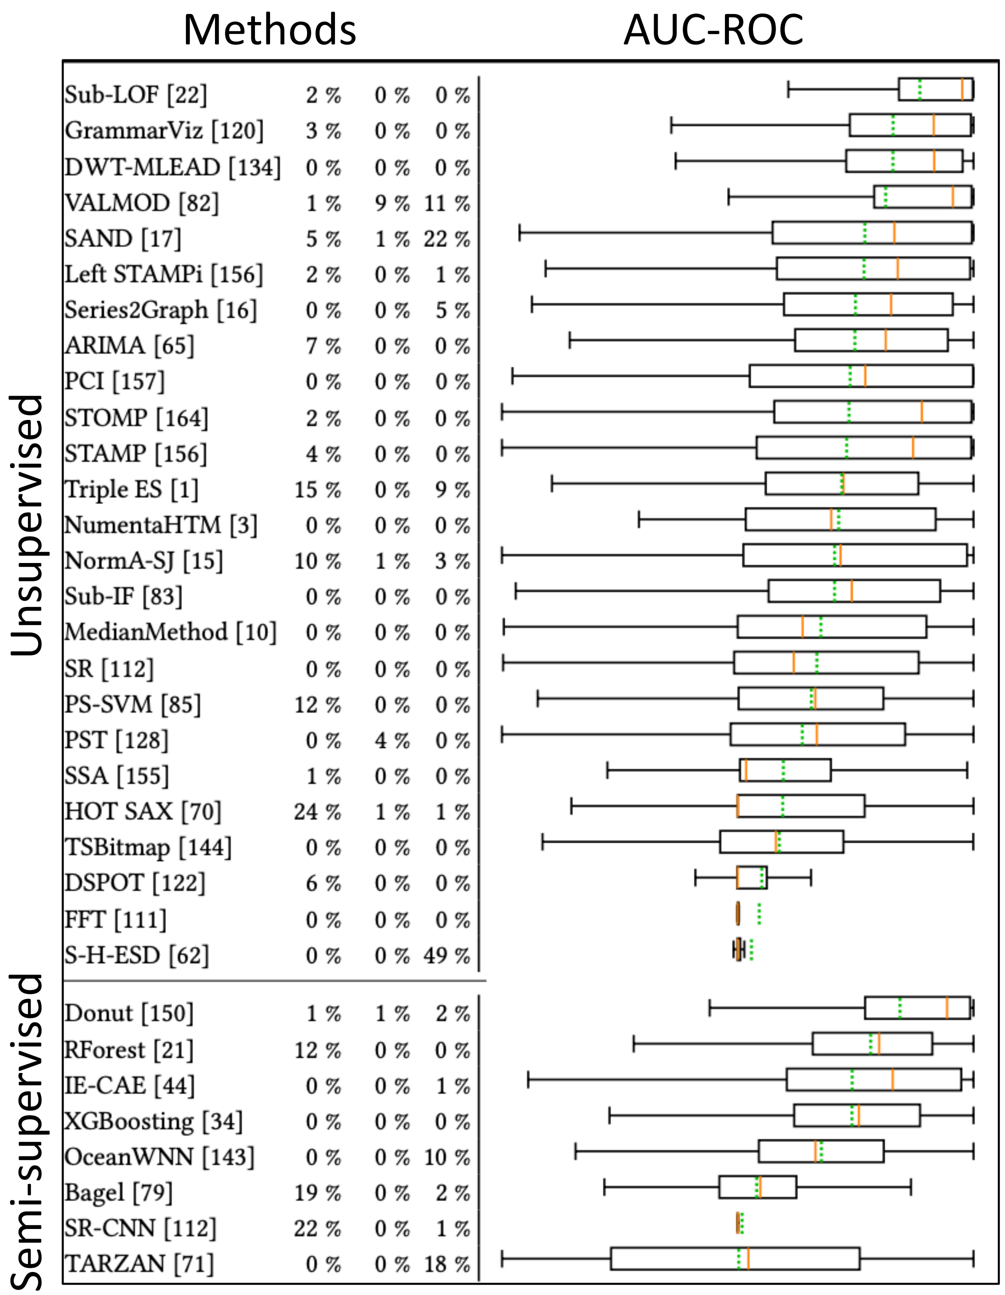

[5] Sebastian Schmidl, Phillip Wenig, and Thorsten Papenbrock. 2022. Anomaly detection in time series: a comprehensive evaluation. Proc. VLDB Endow. 15, 9 (May 2022), 1779–1797.

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

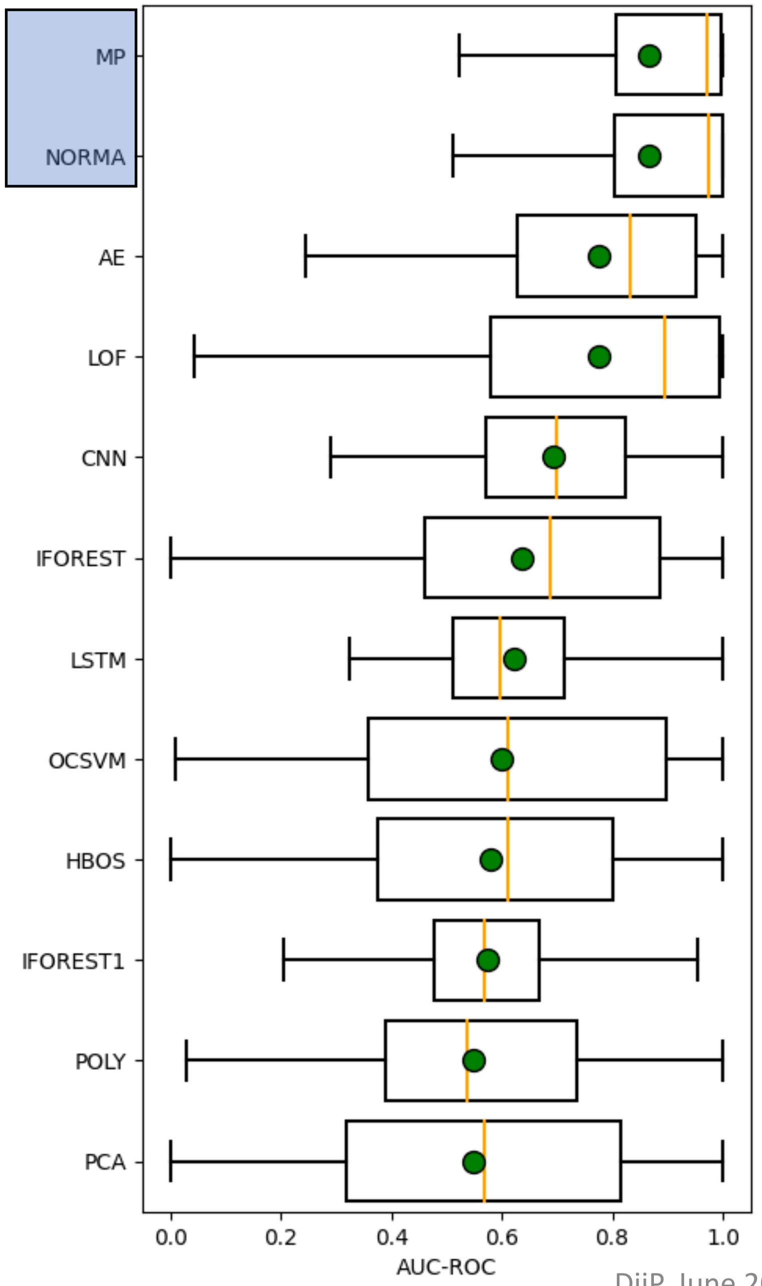

[18] R. Wu and E. Keogh, "Current Time Series Anomaly Detection Benchmarks are Flawed and are Creating the Illusion of Progress" in IEEE Transactions on Knowledge & Data Engineering, vol. 35, no. 03, pp. 2421-2429, 2023.

Anomaly Detection methods:
Experimental evaluation

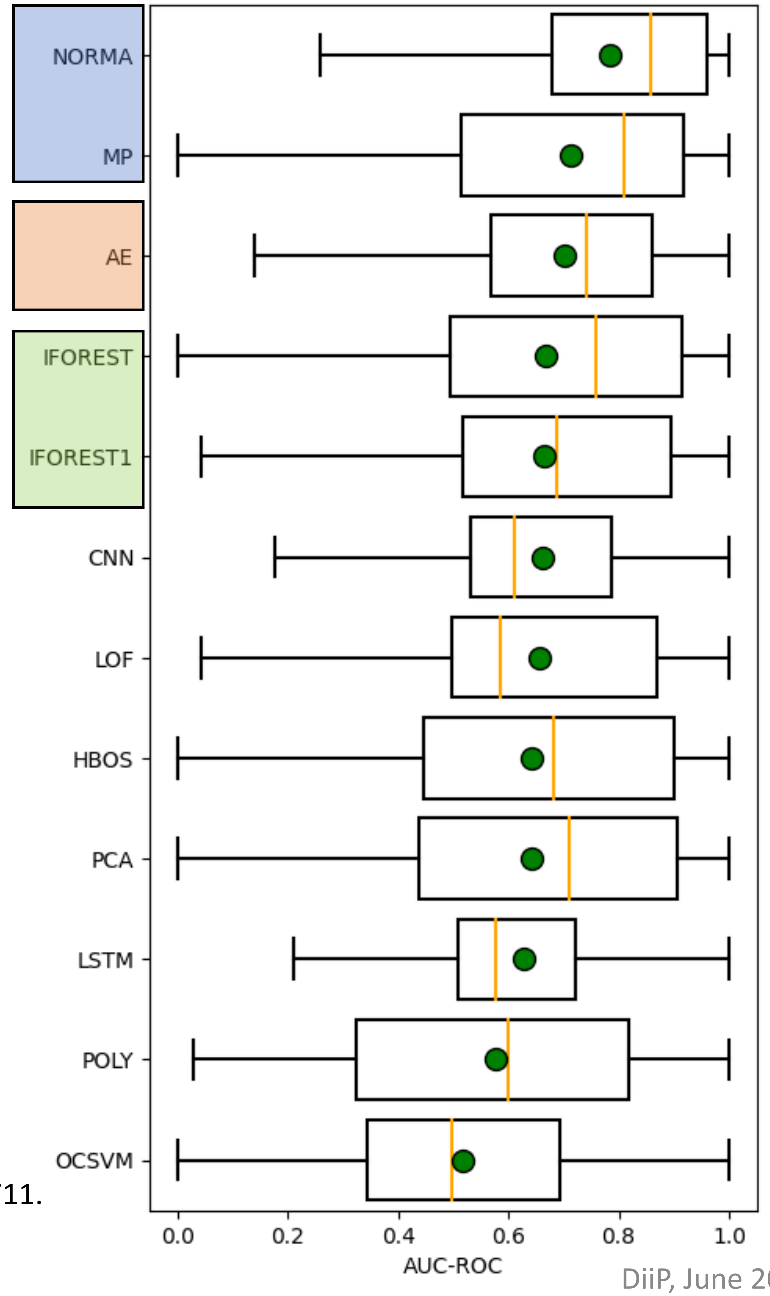

Observations on TSB-UAD [19]:

- Distance-based methods have a better accuracy (AUC-ROC) than forecasting-based methods.
- Isolation Forest (Tree-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.

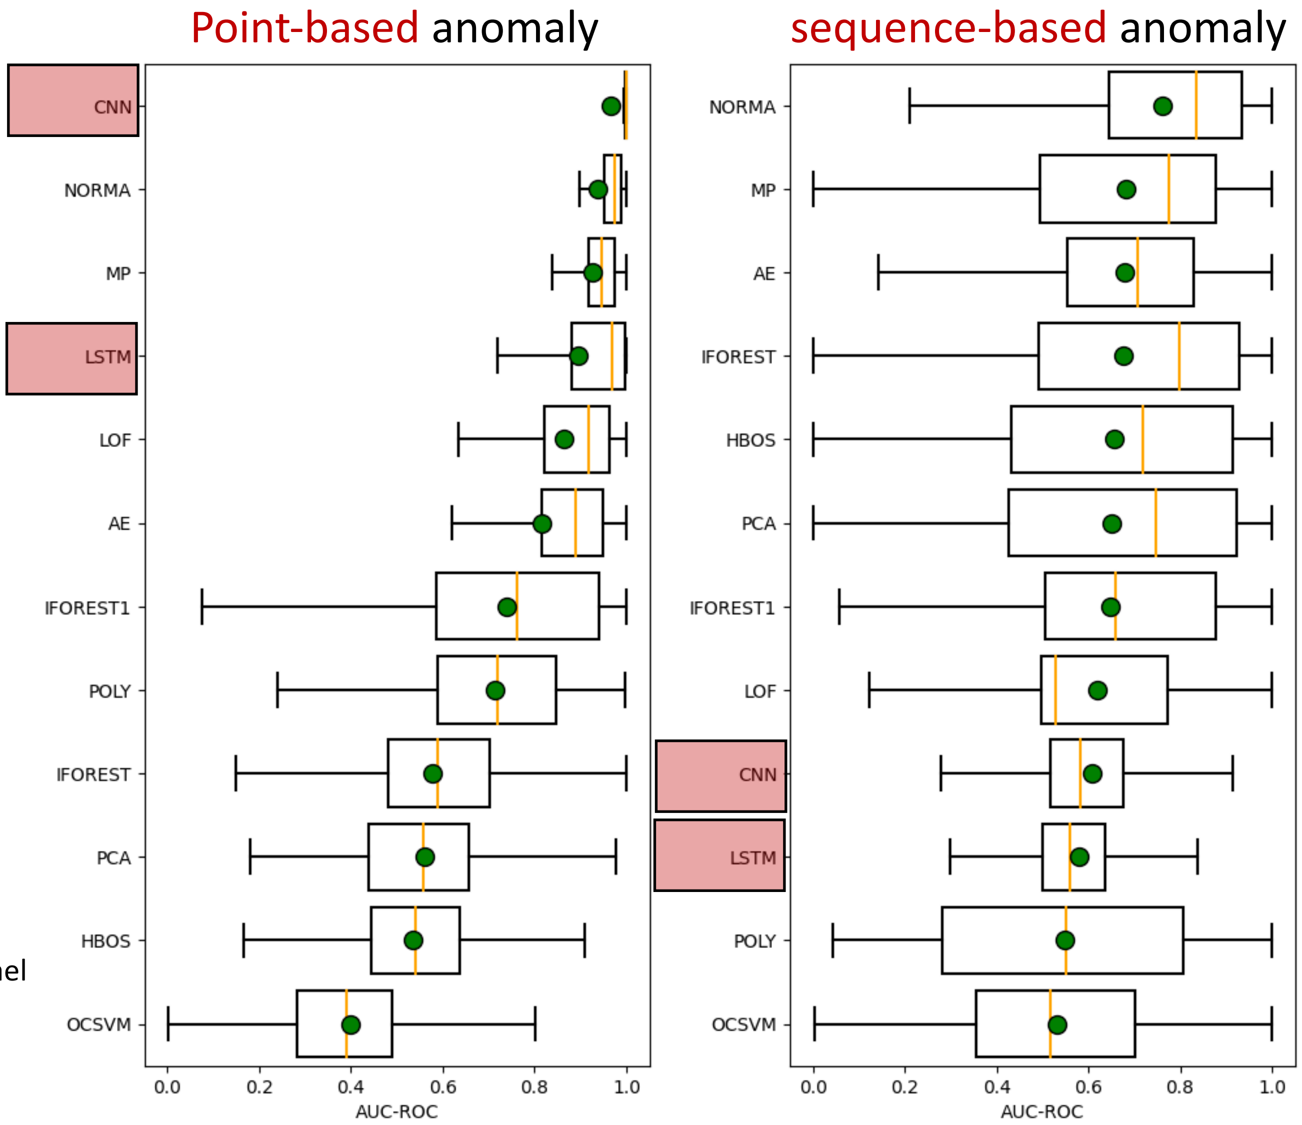

[19] John Paparrizos, Yuhao Kang, Paul Boniol, Ruey S. Tsay, Themis Palpanas, and Michael J. Franklin. 2022. TSB-UAD: an end-to-end benchmark suite for univariate time-series anomaly detection. Proc. VLDB Endow. 15, 8 (April 2022), 1697–1711.

Anomaly Detection methods:
Experimental evaluation

Observations on TSB-UAD [19]:

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

Ratio>0.1

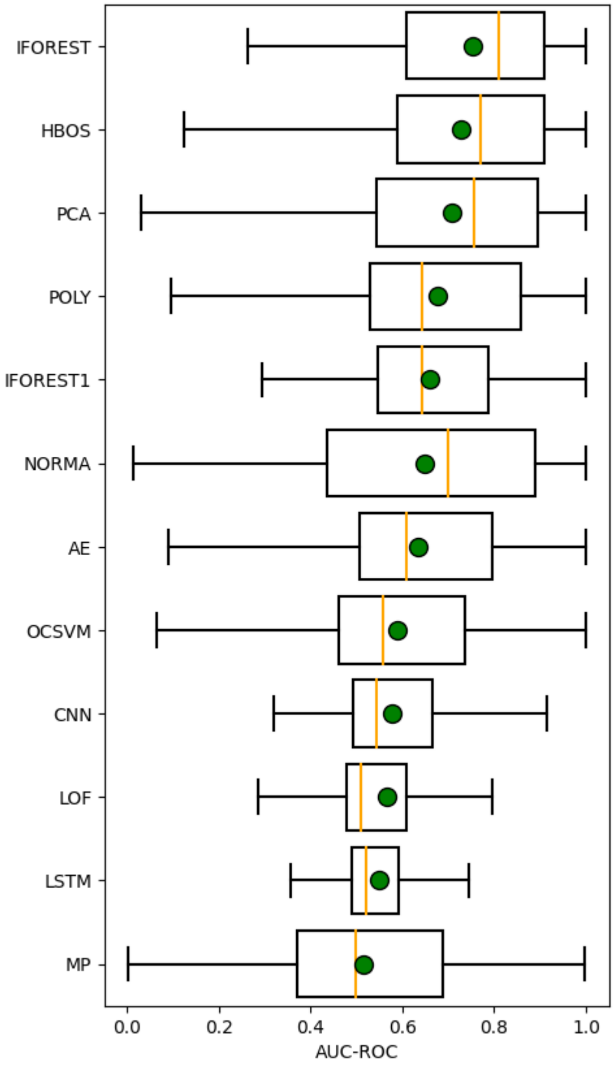

Ratio<0.001

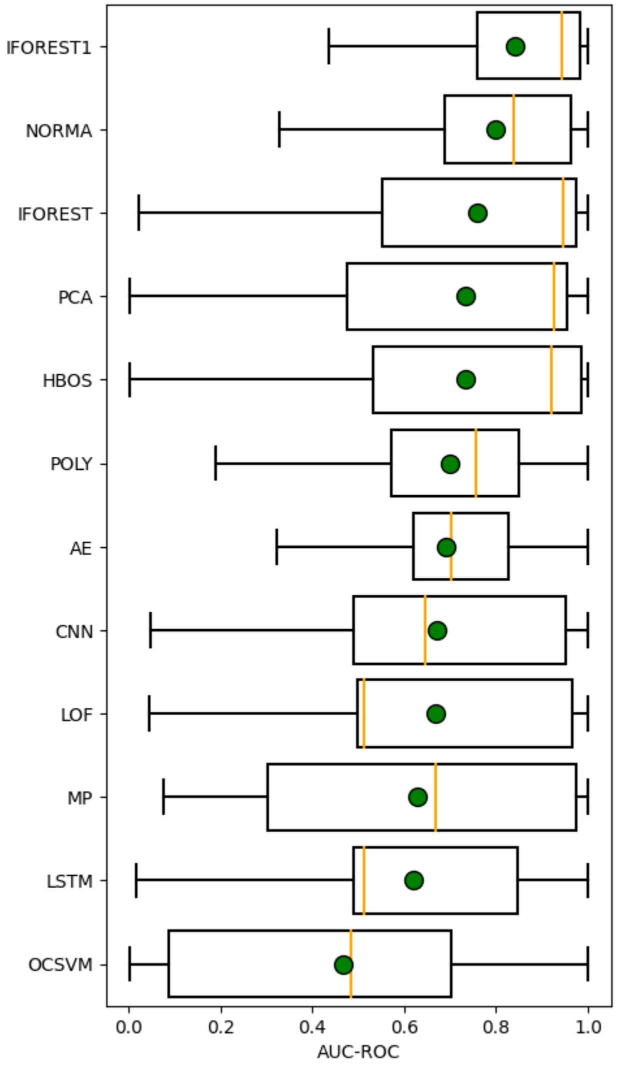

[19] John Paparrizos, Yuhao Kang, Paul Boniol, Ruey S. Tsay, Themis Palpanas, and Michael J. Franklin. 2022. TSB-UAD: an end-to-end benchmark suite for univariate time-series anomaly detection. Proc. VLDB Endow. 15, 8 (April 2022), 1697–1711.

Anomaly Detection methods: *Experimental evaluation*

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

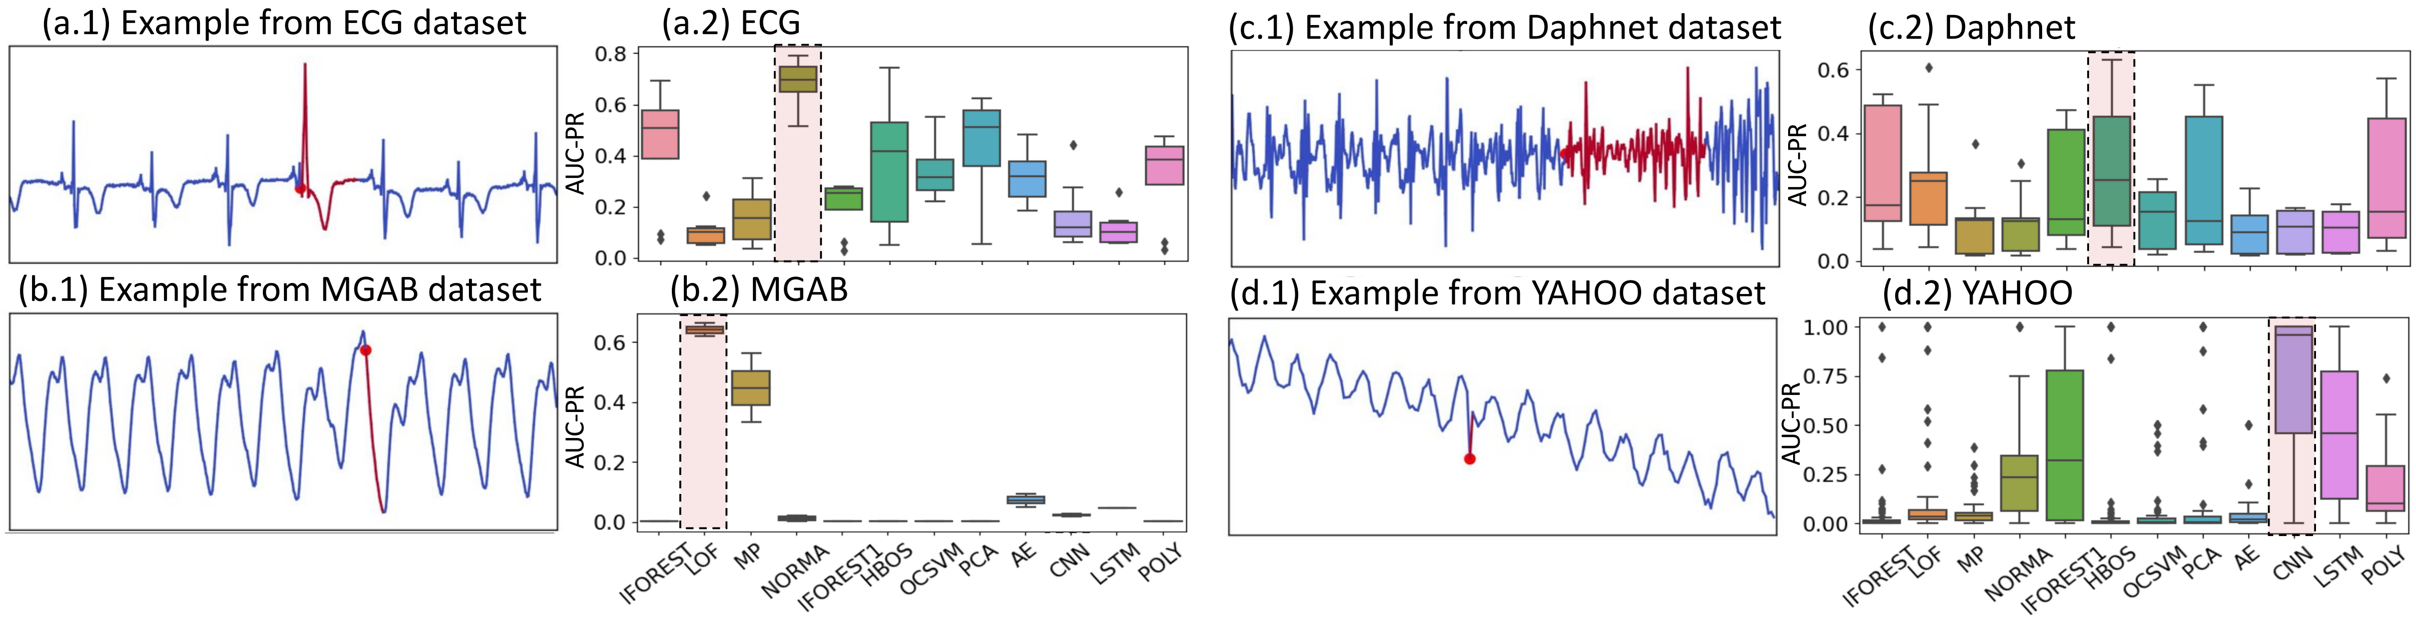

There is no overall winner.

[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.

Perspectives and challenges

Conclusion and Open Problems

If you are interested in anomaly detection in time series…

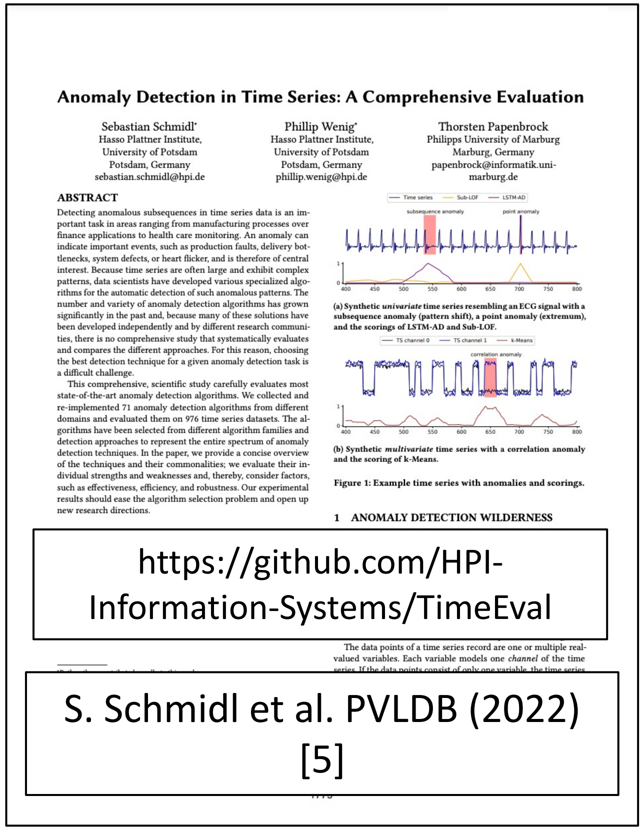

https://github.com/HPI-Information-Systems/TimeEval

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

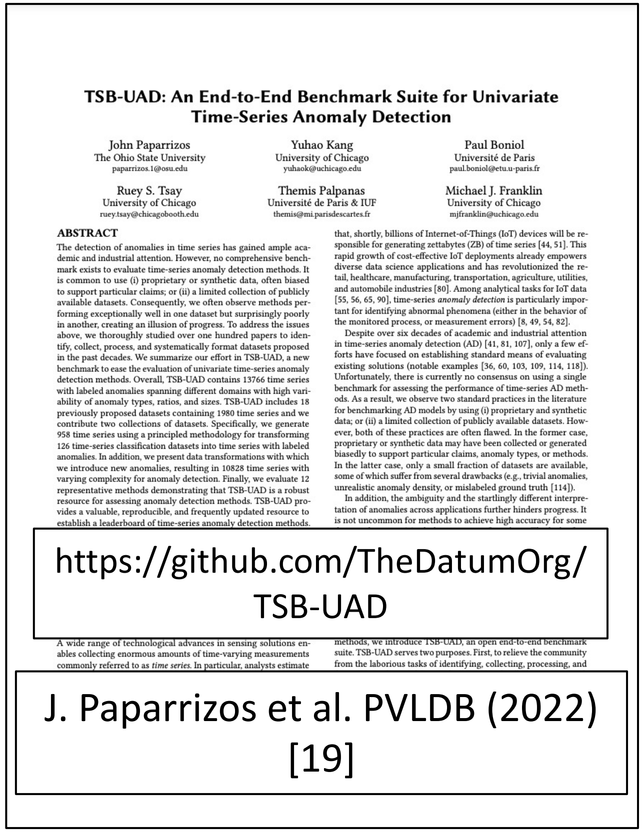

https://github.com/TheDatumOrg/TSB-UAD

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

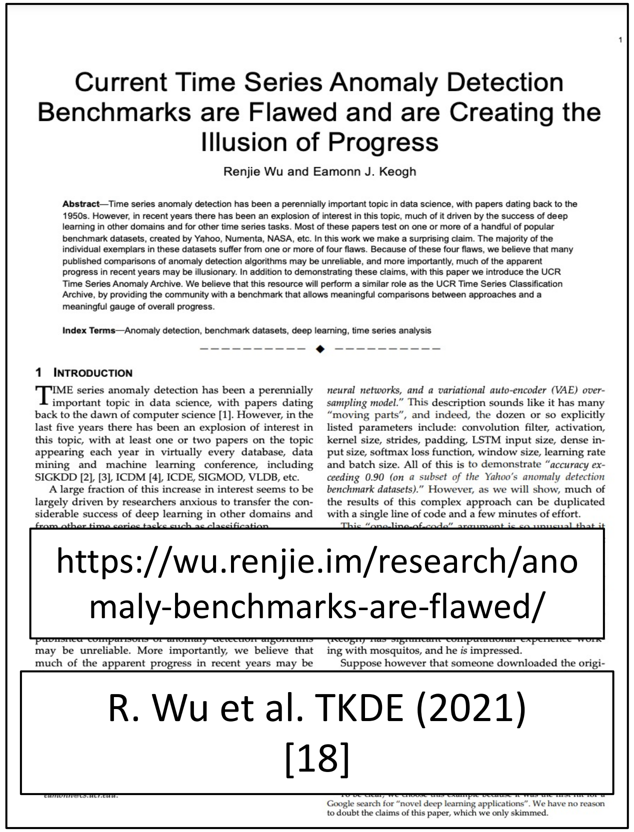

https://wu.renjie.im/research/anomaly-benchmarks-are-flawed/

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

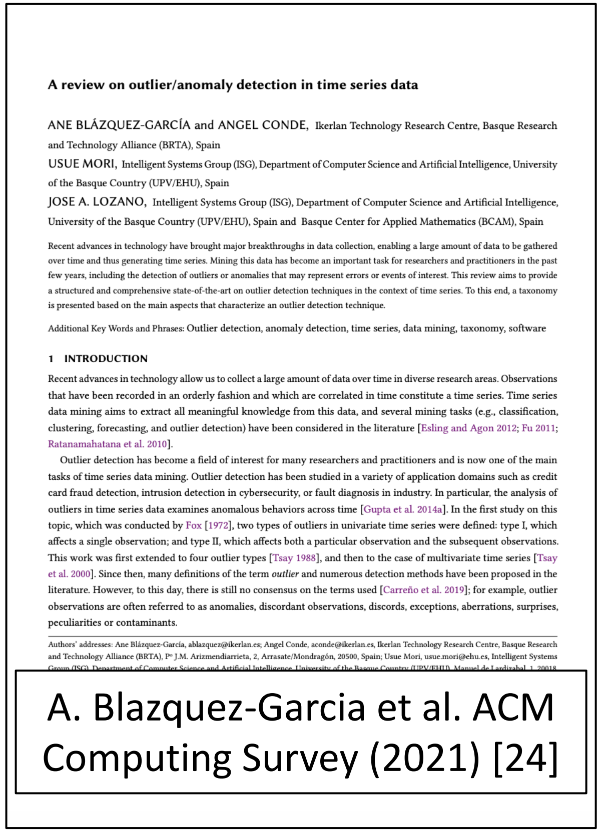

A. Blazquez-Garcia et al. ACM Computing Survey (2021) [24]

Conclusion and Open Problems

If you are interested in anomaly dete

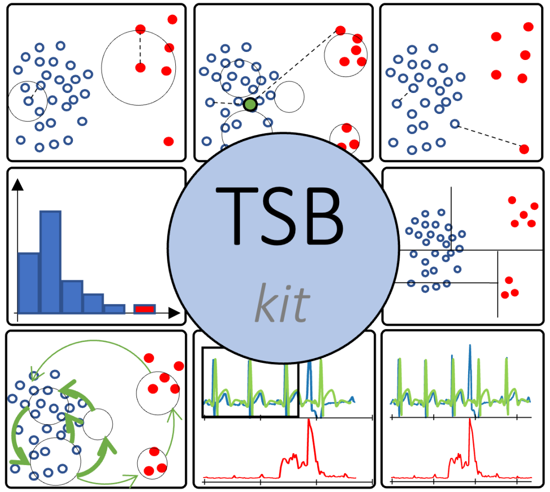

PyPI v0.0.5, Python 3.8

Pip install tsb-kit

GitHub

Documentation

Anomaly Detection in Time Series: A Comprehensive Evaluation

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

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

TSB-UAD: An End-

https://githu

J. Paparrizo

maly Detection
are Creating the
ress

research/ano
re-flawed/

E (2021)

A review on outlier/anomaly detection in time series data

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

Conclusion and Open Problems

Context-aware Unsupervised Anomaly Detection

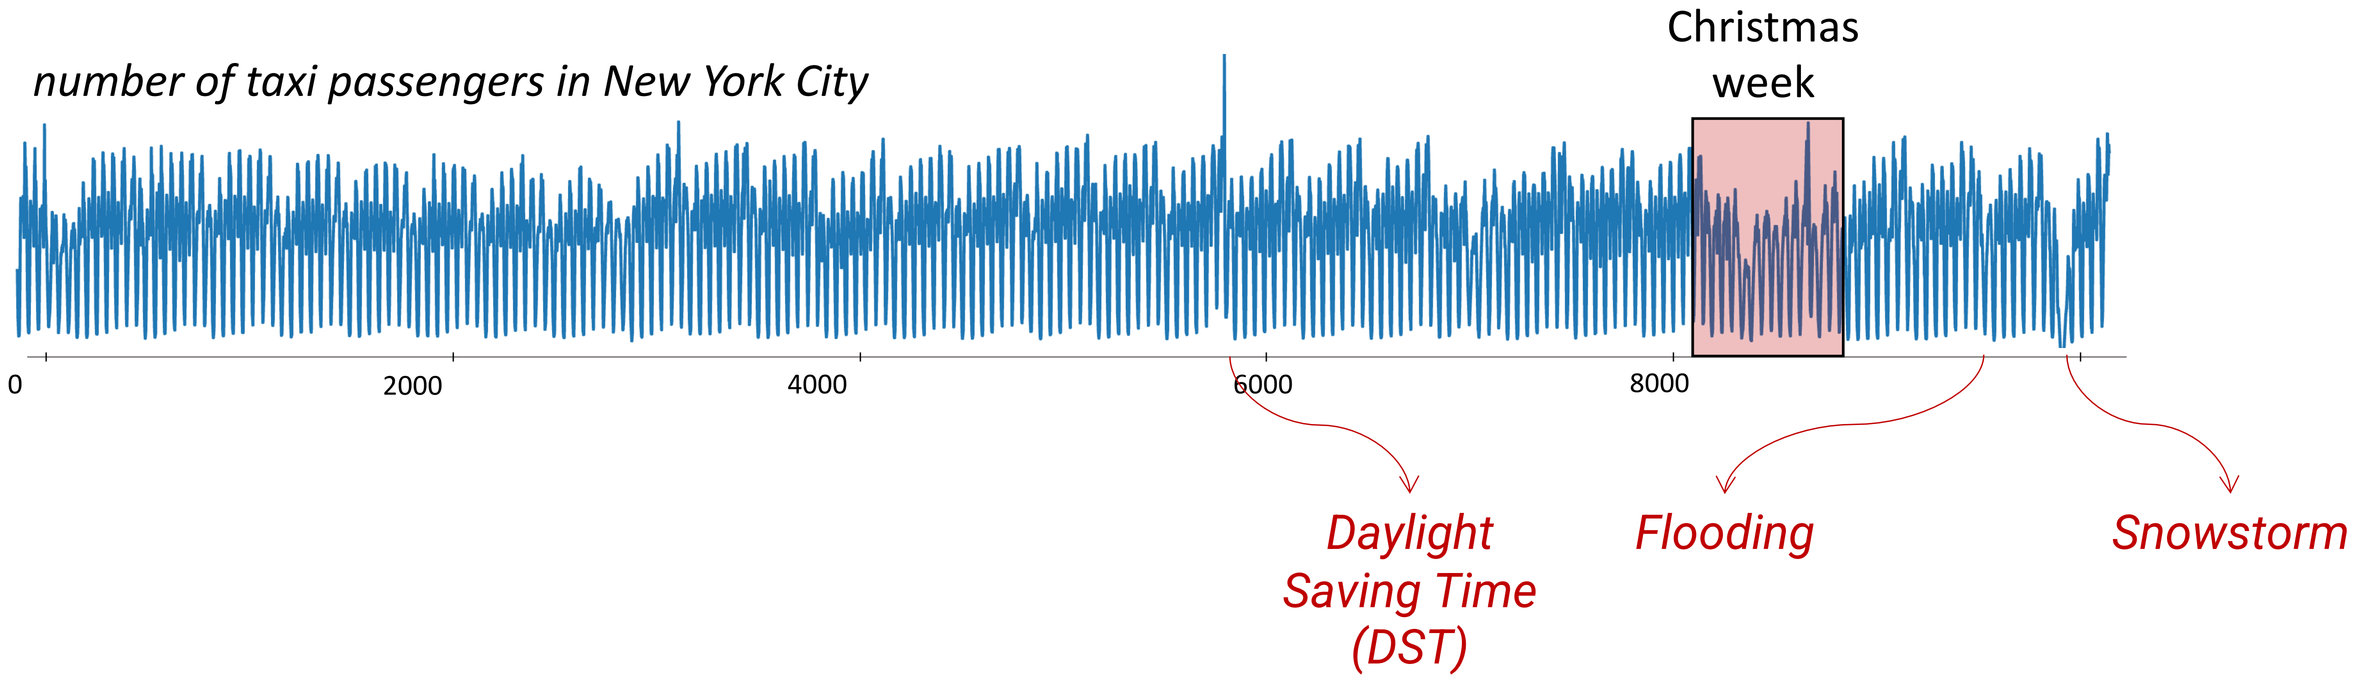

Conclusion and Open Problems

Evaluating Anomaly Detection

$(ex1)\ Example\ on\ IOPS$

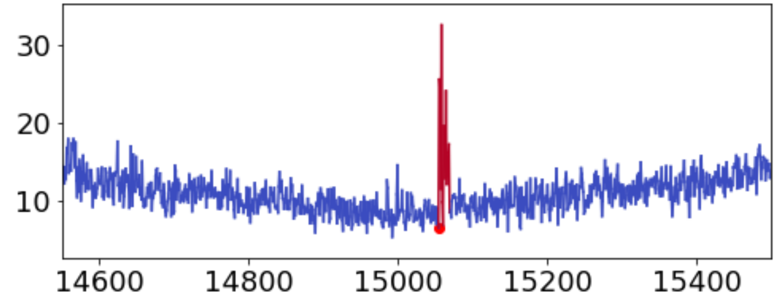

$(ex2)\ Example\ on\ SensorScope$

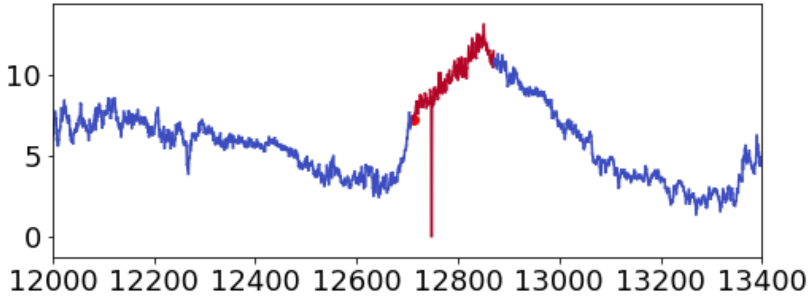

$(ex3)\ Example\ on\ NAB$

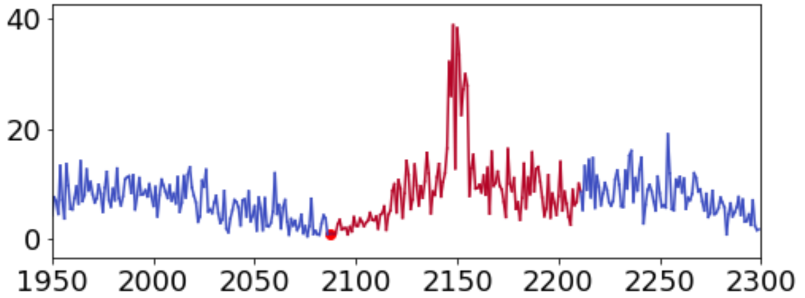

What is the problem here?

Conclusion and Open Problems

Threshold-based Evaluation Measures:

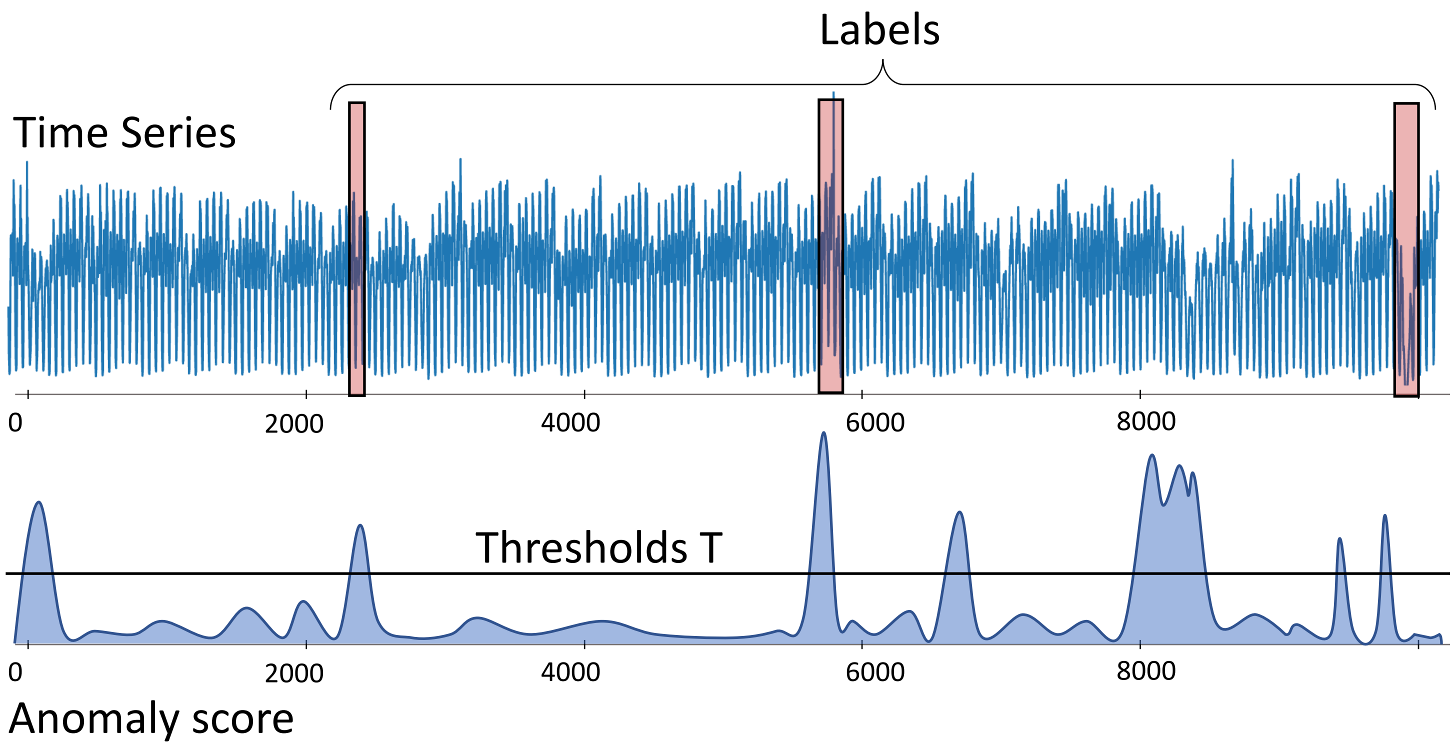

Conclusion and Open Problems

Threshold-based Evaluation Measures:

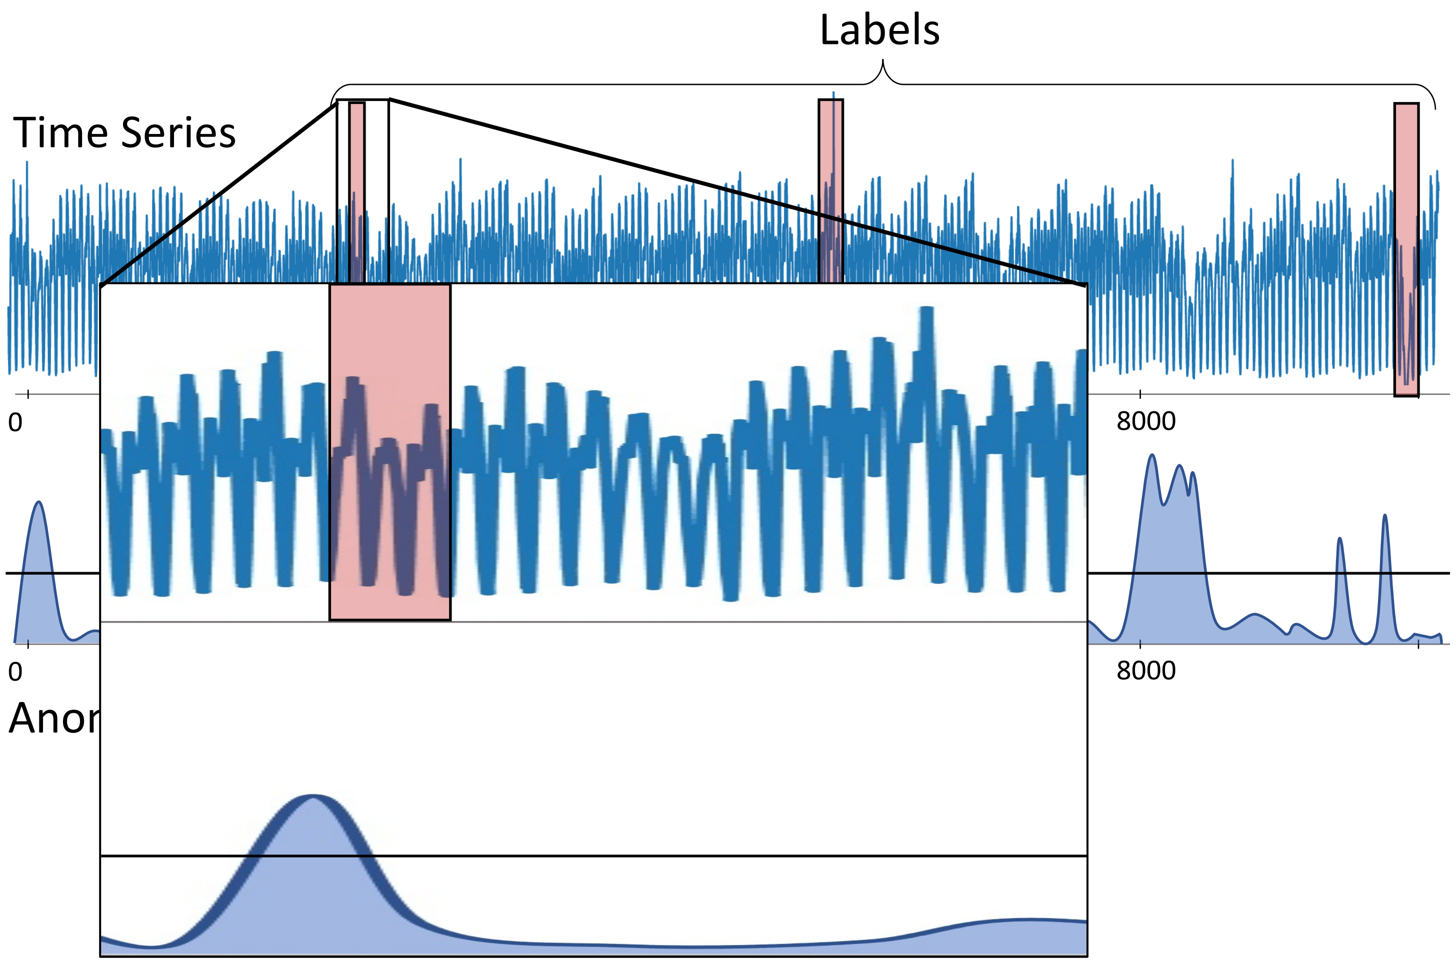

Conclusion and Open Problems

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: $\frac{(1+\beta^2)*Precision}{\beta^2*Precision+Recall}$

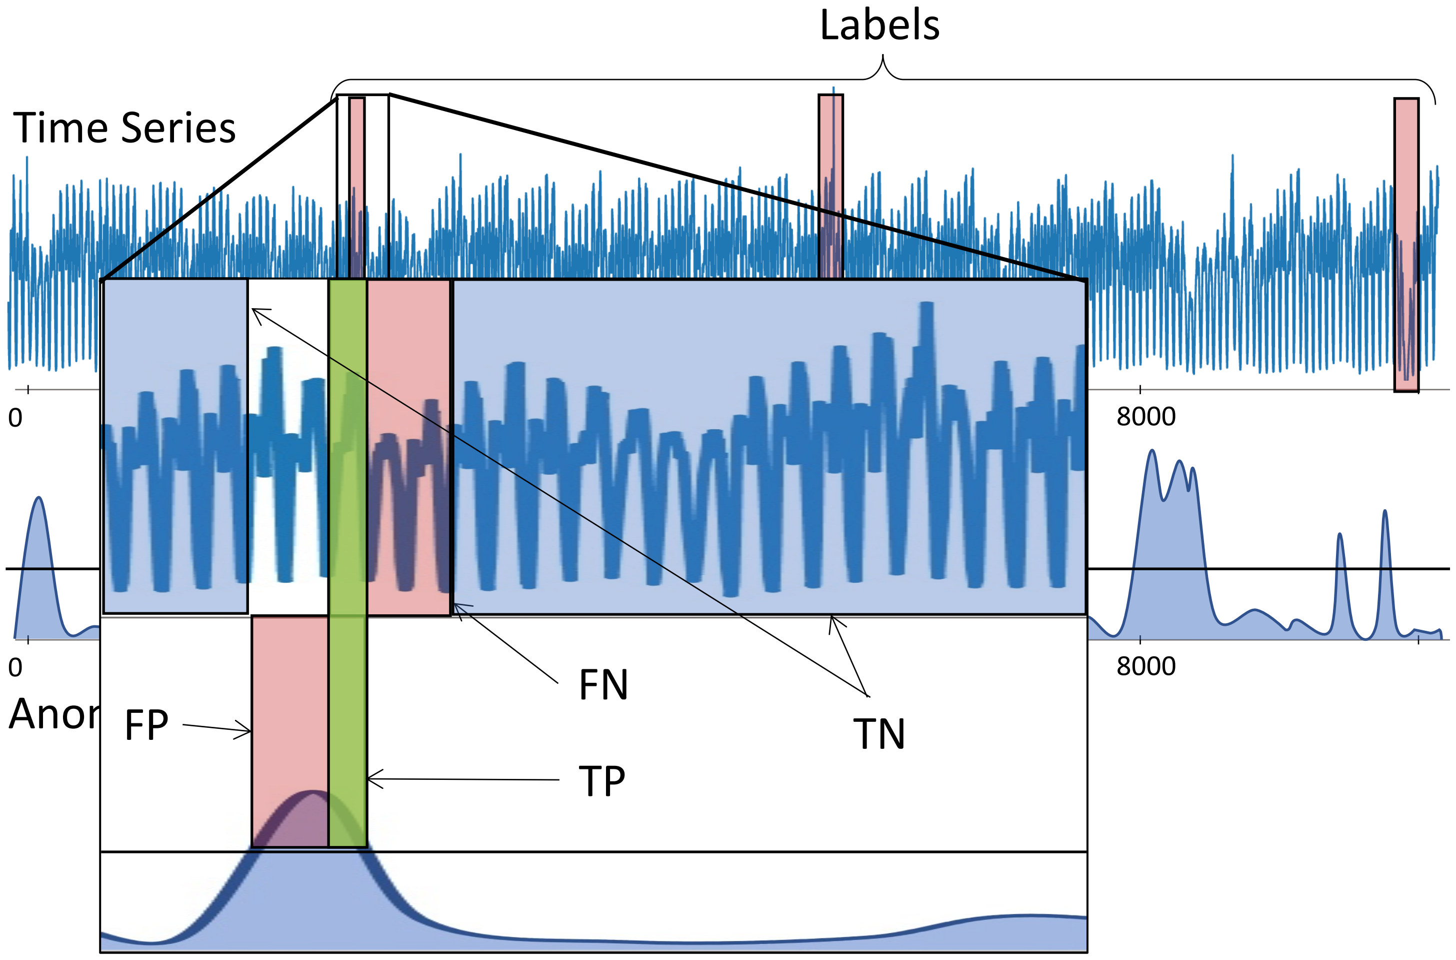

Conclusion and Open Problems

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:
  - Methods that produce misaligned anomaly scores.
  - Different Labeling strategies between domains and applications

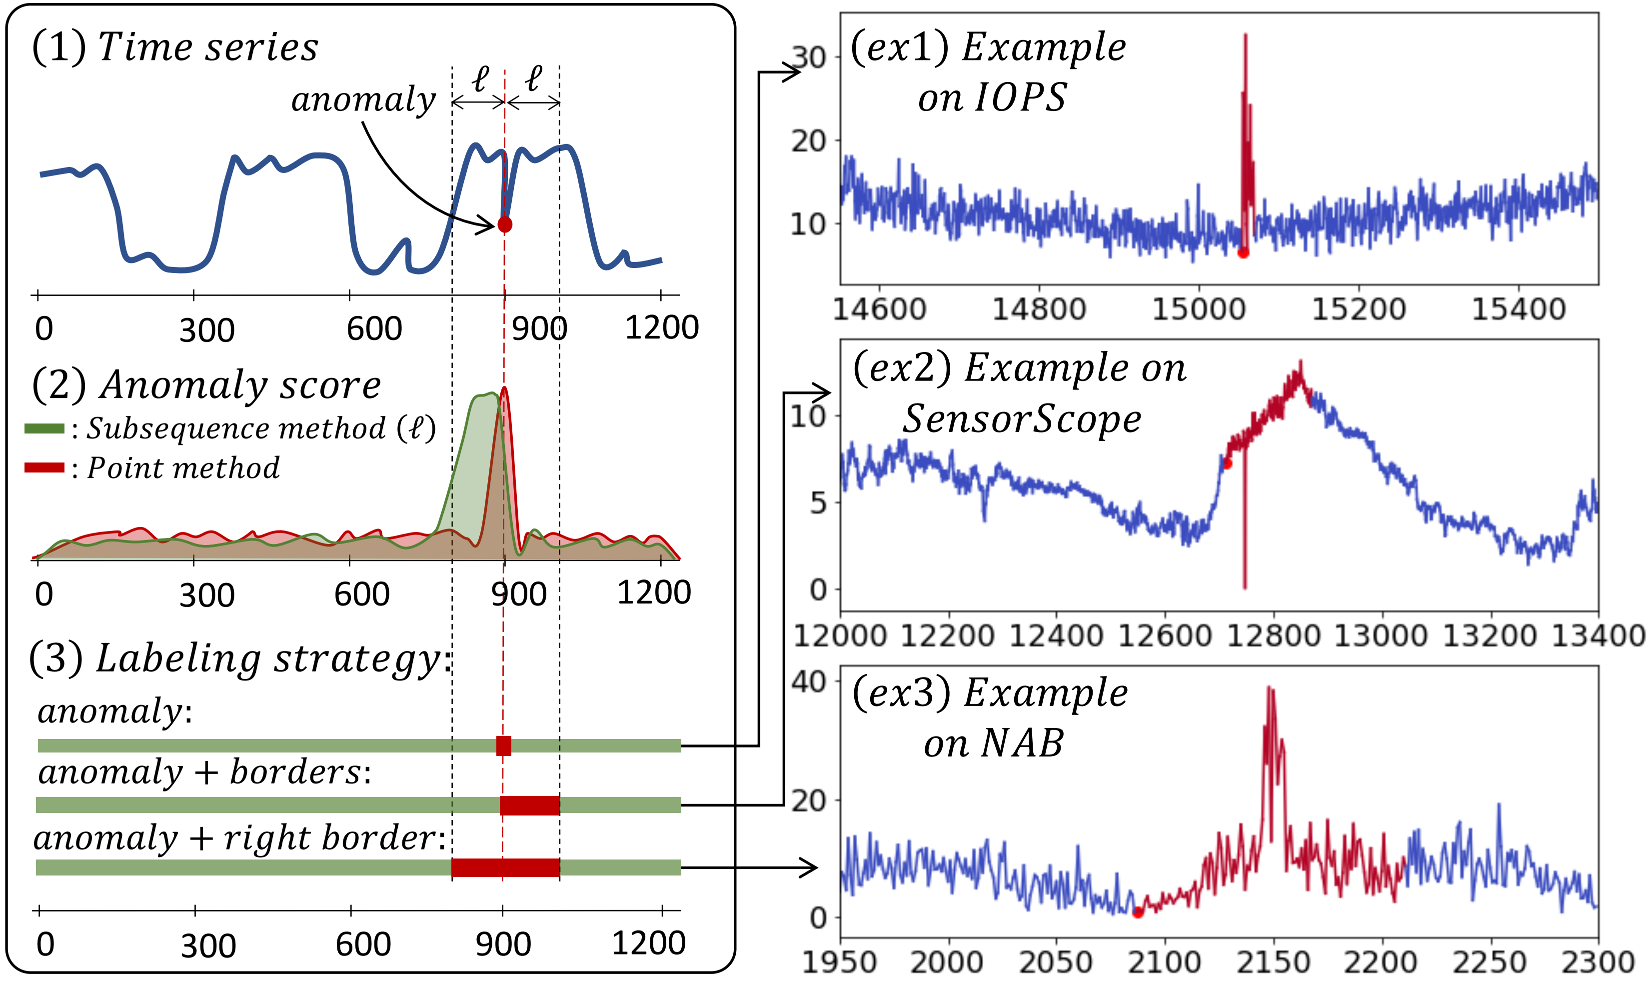

Conclusion and Open Problems

If you are interested in evaluation measures for anomaly detection…

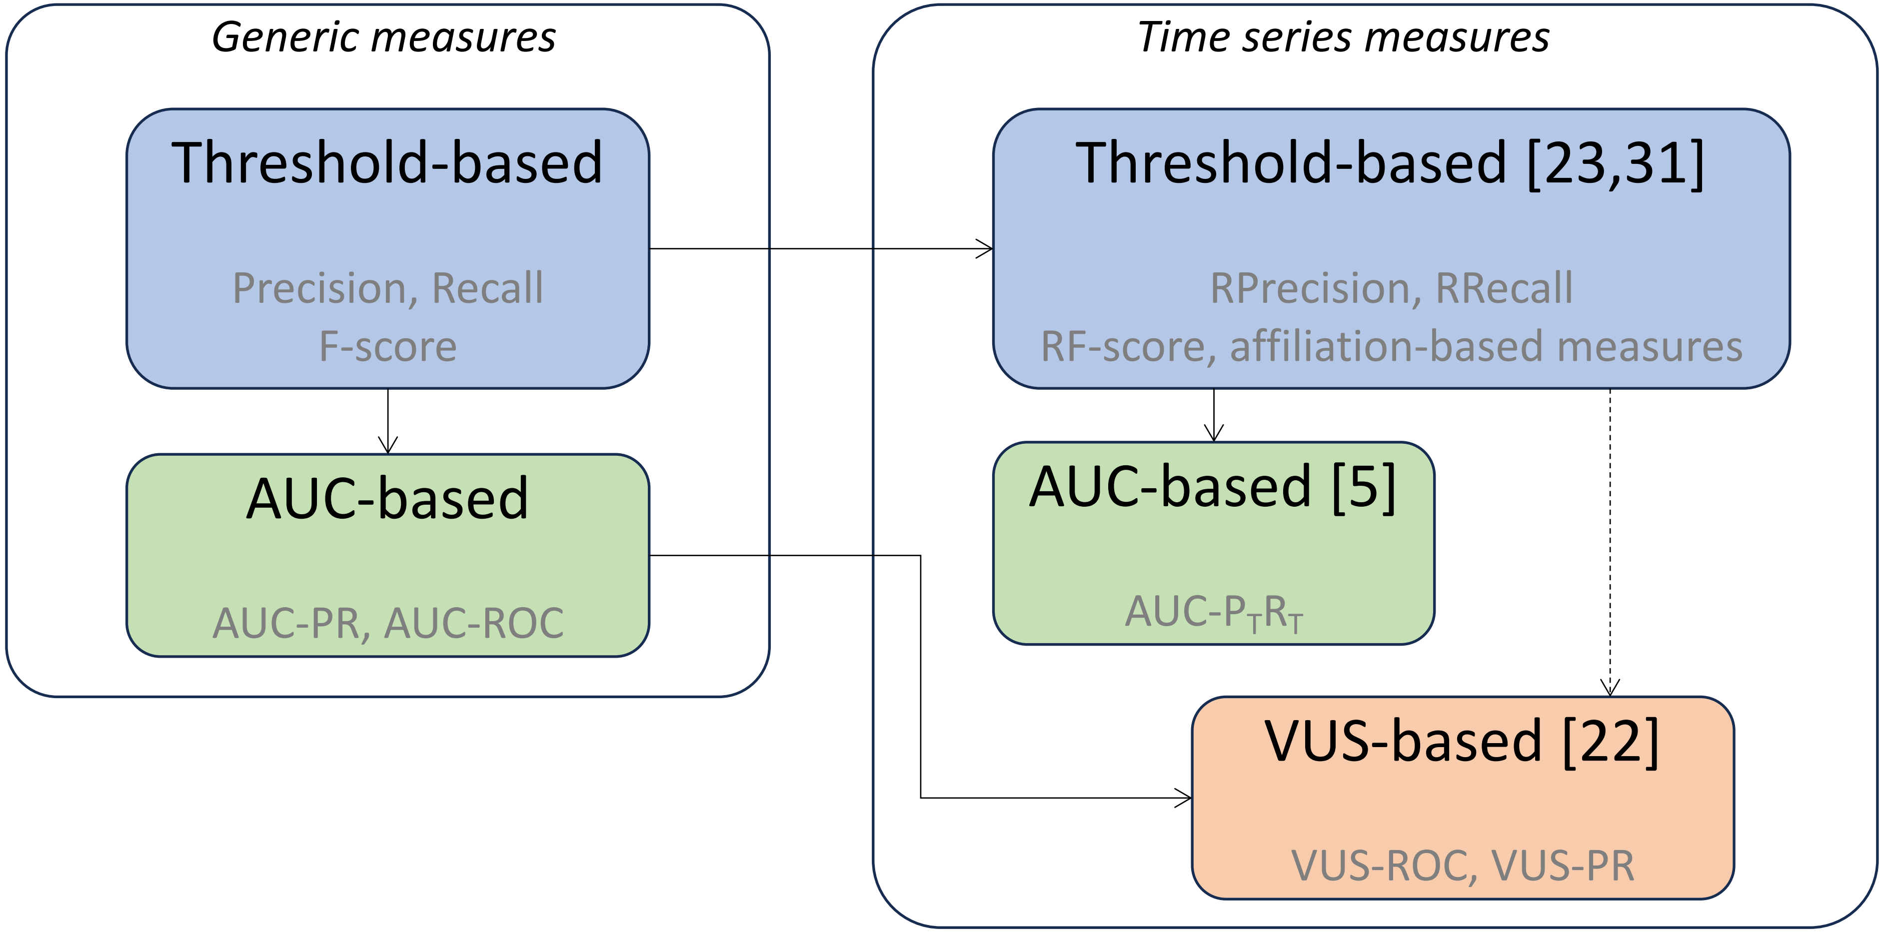

Conclusion and Open Problems

If you are interested in evaluation measures for anomaly detection…

Precision and Recall for Time Series

Nesime Tatbul *, Intel Labs and MIT, tatbul@csail.mit.edu
Tae Jun Lee *, Microsoft, tae_jun_lee@alumni.brown.edu
Stan Zdonik, Brown University, sbz@cs.brown.edu
Mejbah Alam, Intel Labs, mejbah.alam@intel.com
Justin Gottschlich, Intel Labs, justin.gottschlich@intel.com

https://arxiv.org/abs/1803.03639

N. Tatbul et al. NeurIPS 2018 [23]

Volume Under the Surface: A New Accuracy Evaluation Measure for Time-Series Anomaly Detection

John Paparrizos, The Ohio State University
Paul Boniol, Université Paris Cité
Themis Palpanas, Université Paris Cité; IUF
Ruey S. Tsay, University of Chicago
Aaron Elmore, University of Chicago
Michael J. Franklin, University of Chicago

https://www.vldb.org/pvldb/vol15/p2774-paparrizos.pdf

J. Paparrizos et al. PVLDB 2022 [22]

Local Evaluation of Time Series Anomaly Detection Algorithms

Alexis Huet, Huawei Technologies Co., Ltd., France, alexis.huet@huawei.com
Jose Manuel Navarro, Huawei Technologies Co., Ltd., France, jose.manuel.navarro@huawei.com
Dario Rossi, Huawei Technologies Co., Ltd., France, dario.rossi@huawei.com

https://arxiv.org/abs/2206.13167

A. Huet et al. KDD 2022 [31]

Navigating the Metric Maze: A Taxonomy of Evaluation Metrics for Anomaly Detection in Time Series

Sondre Sørbø[1], sondre.sorbo@sintef.no
Massimiliano Ruocco[1,2], massimiliano.ruocco@sintef.no
[1]SINTEF Digital, Trondheim, Norway
[2]Norwegian University of Science and Technology, Trondheim, Norway

https://arxiv.org/abs/2303.01272

S. Sørbø et al. DAMI 2024 [29]

Conclusion and Open Problems

Model selection for anomaly detection

Methods ranking changes
significantly between datasets [19]

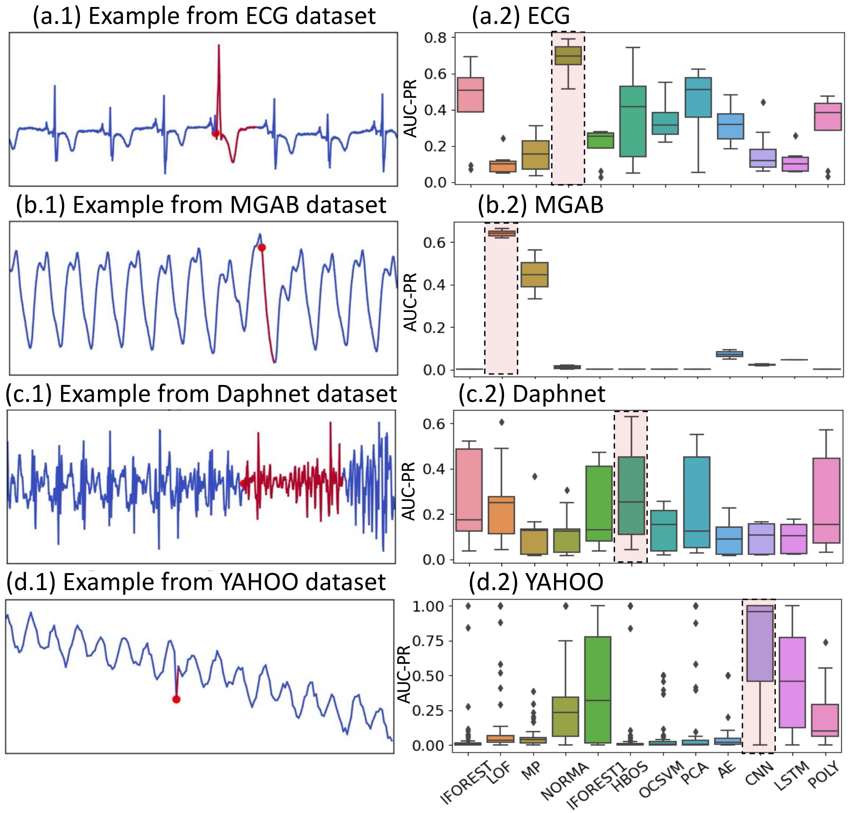

Results over TSB-UAD

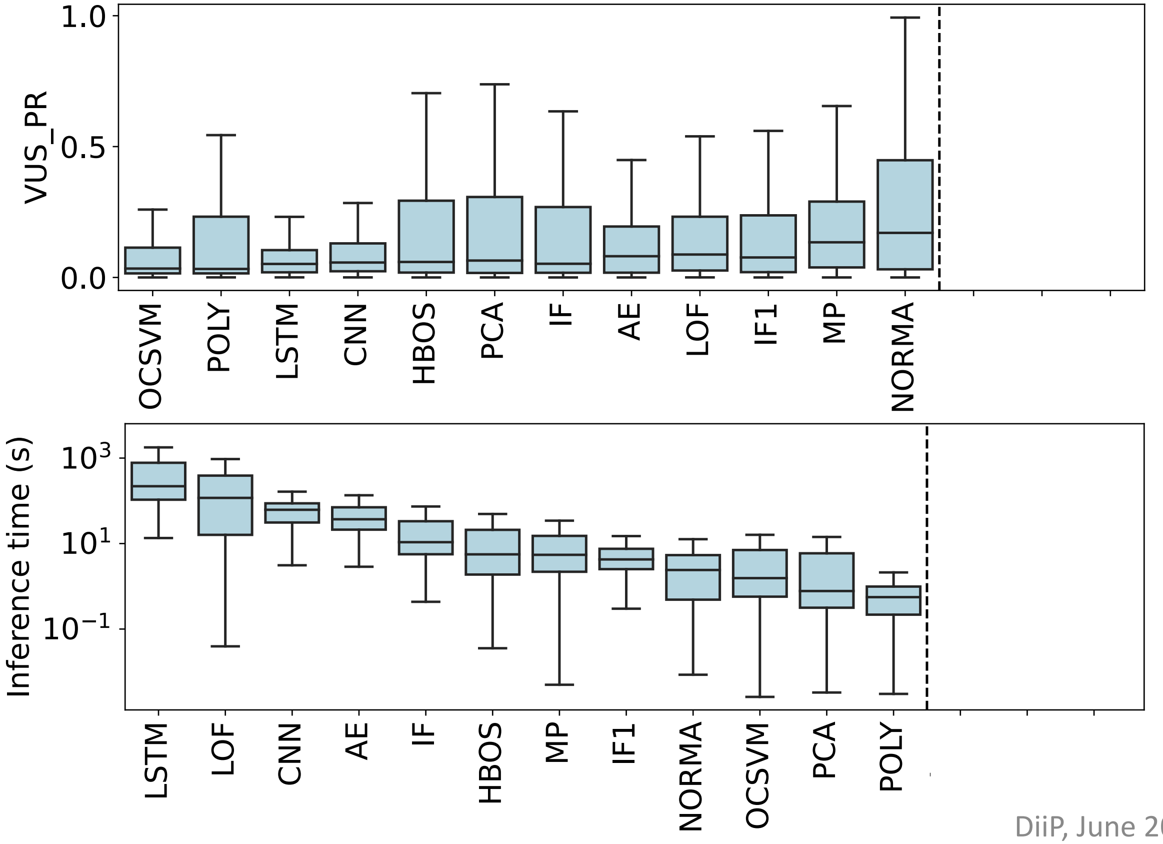

Conclusion and Open Problems

Model selection for anomaly detection

Methods ranking changes significantly between datasets [19]

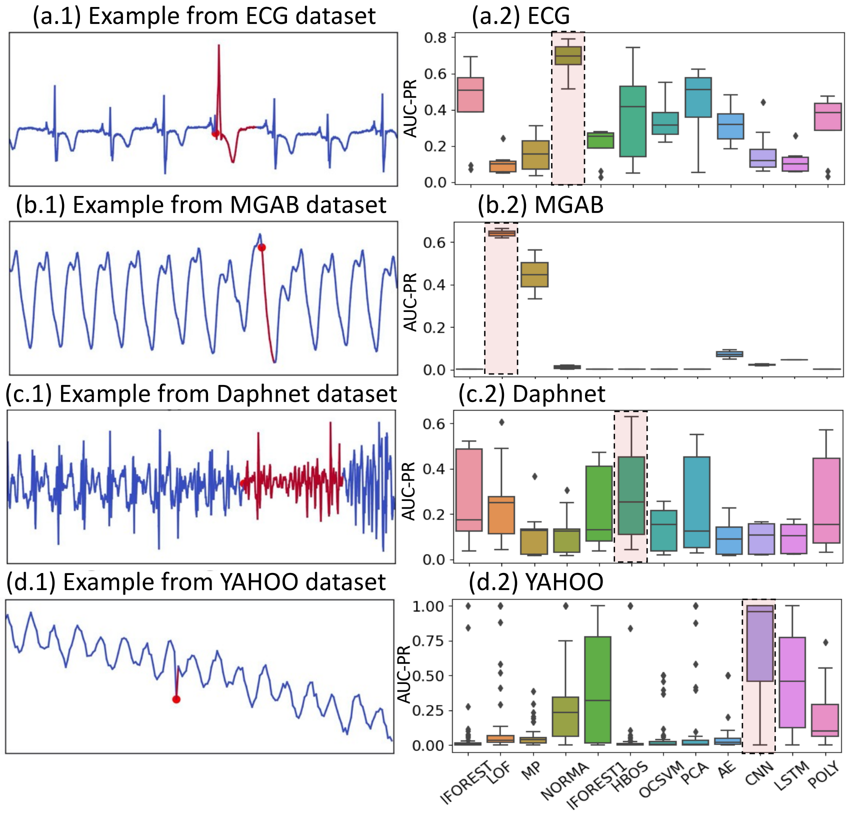

Can *Ensembling* methods solve the problem?

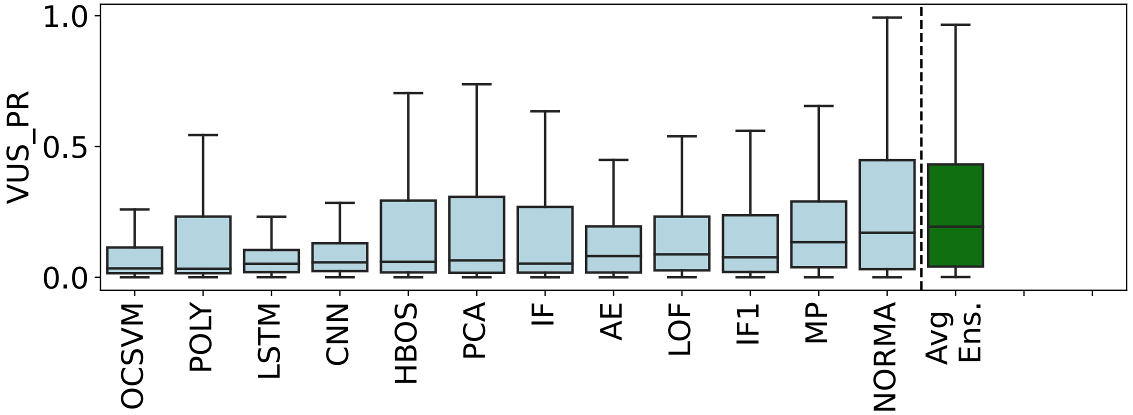

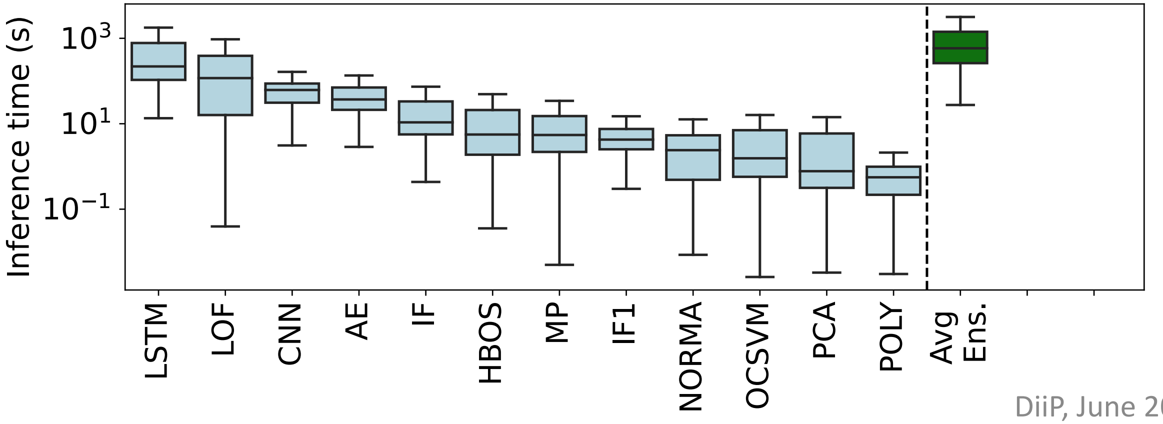

Conclusion and Open Problems

Model selection for anomaly detection

Methods ranking changes significantly between datasets [19]

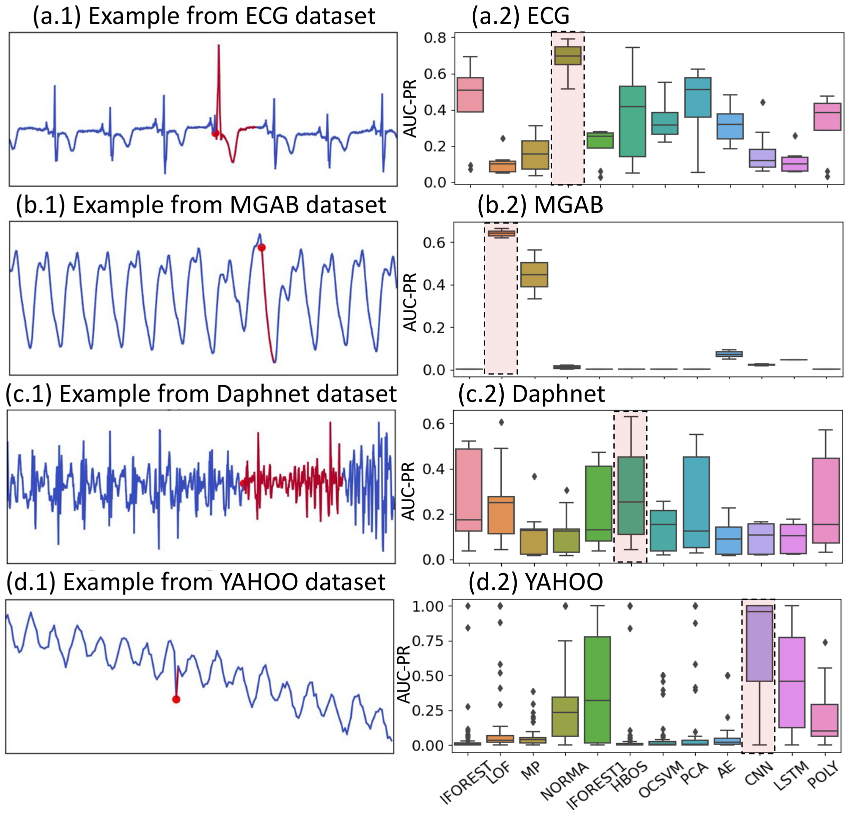

Can *automatic model selection solve the problem*?

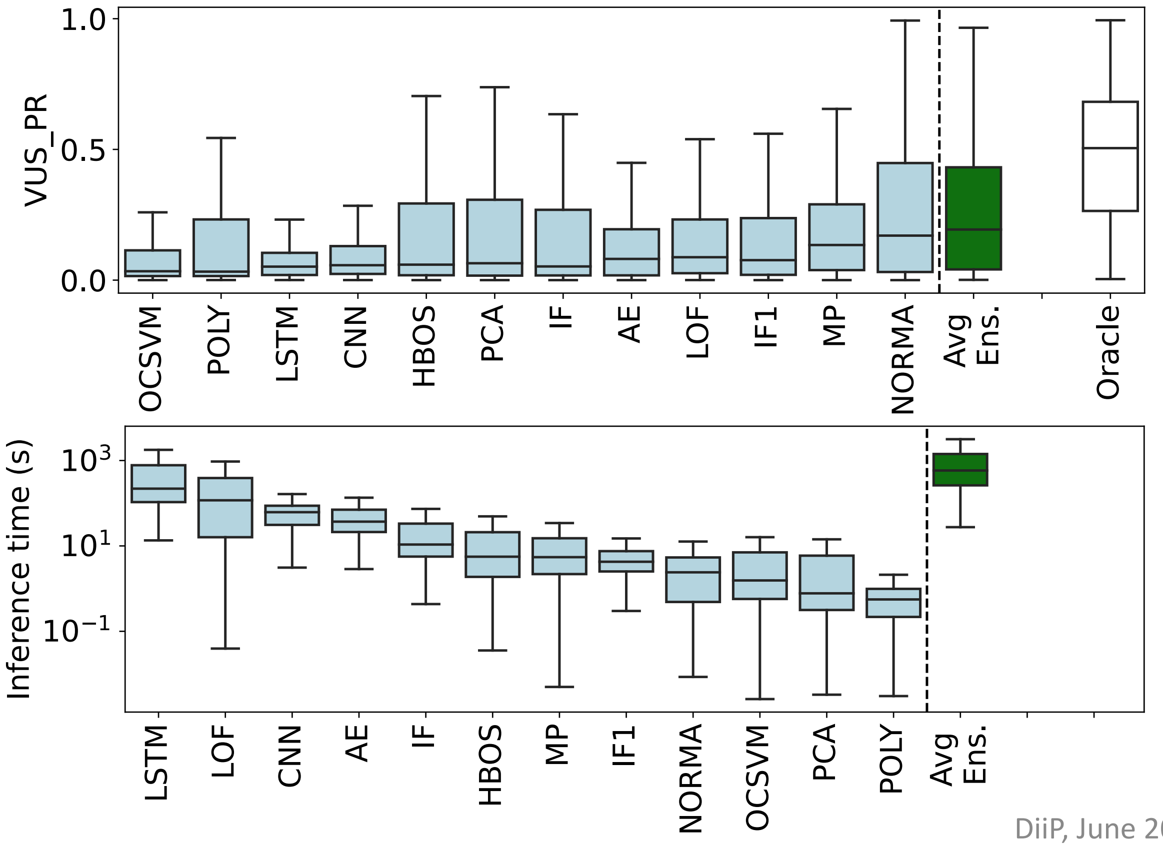

Conclusion and Open Problems

Model selection for anomaly detection

Methods ranking changes significantly between datasets [19]

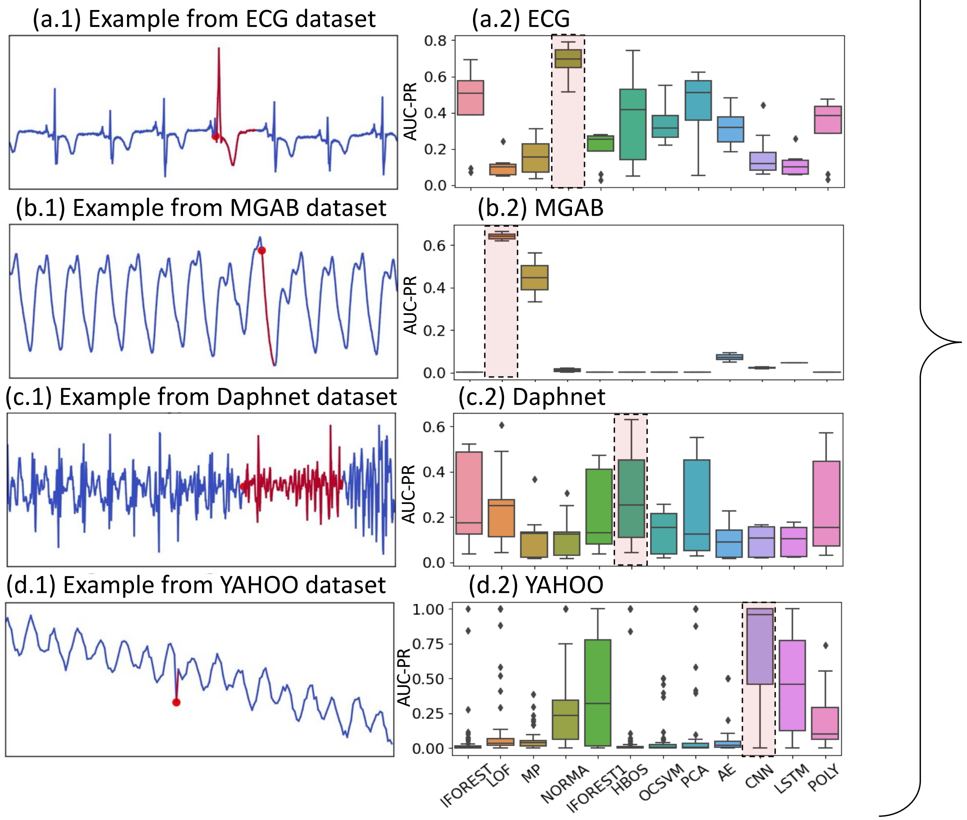

Can *automatic model selection solve the problem*?

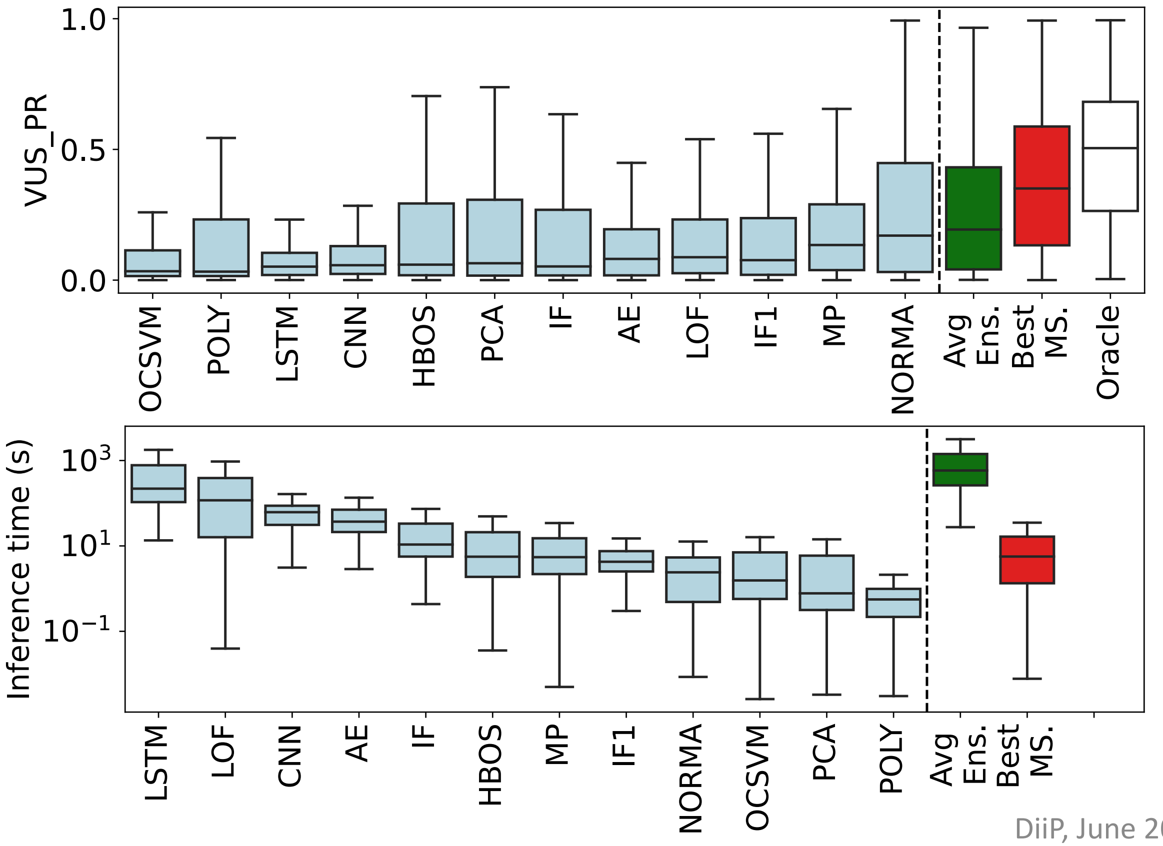

Conclusion and Open Problems

Model selection for anomaly detection

Methods ranking changes significantly between datasets [19]

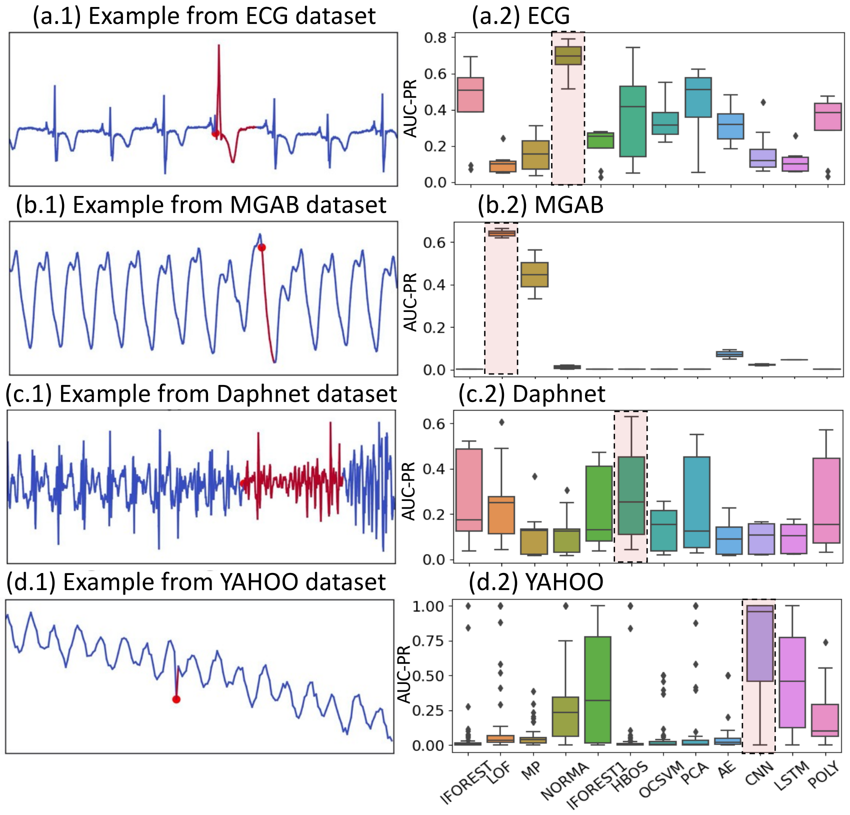

Can *automatic model selection solve the problem*?

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 classication 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 Prole 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. Kavantzas 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 lekki. 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).
[27] Paul Boniol, John Paparrizos, Yuhao Kang, Themis Palpanas, Ruey S. Tsay, Aaron J. Elmore, and Michael J. Franklin. 2022. **Theseus: navigating the labyrinth of time-series anomaly detection.** Proc. VLDB Endow. 15, 12 (August 2022), 3702–3705.
[28] Paul Boniol, Themis Palpanas, Mohammed Meftah, and Emmanuel Remy. 2020. **GraphAn: graph-based subsequence anomaly detection.** Proc. VLDB Endow. 13, 12 (August 2020), 2941–2944.
[29] Sørbø, S., Ruocco, M. **Navigating the metric maze: a taxonomy of evaluation metrics for anomaly detection in time series**. Data Min Knowl Disc 38, 1027–1068 (2024)
[30] 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.
[31] Alexis Huet, Jose Manuel Navarro, and Dario Rossi. 2022. **Local Evaluation of Time Series Anomaly Detection Algorithms.** In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22). Association for Computing Machinery, New York, NY, USA, 635–645.

Thank you for attending!

Any Questions?