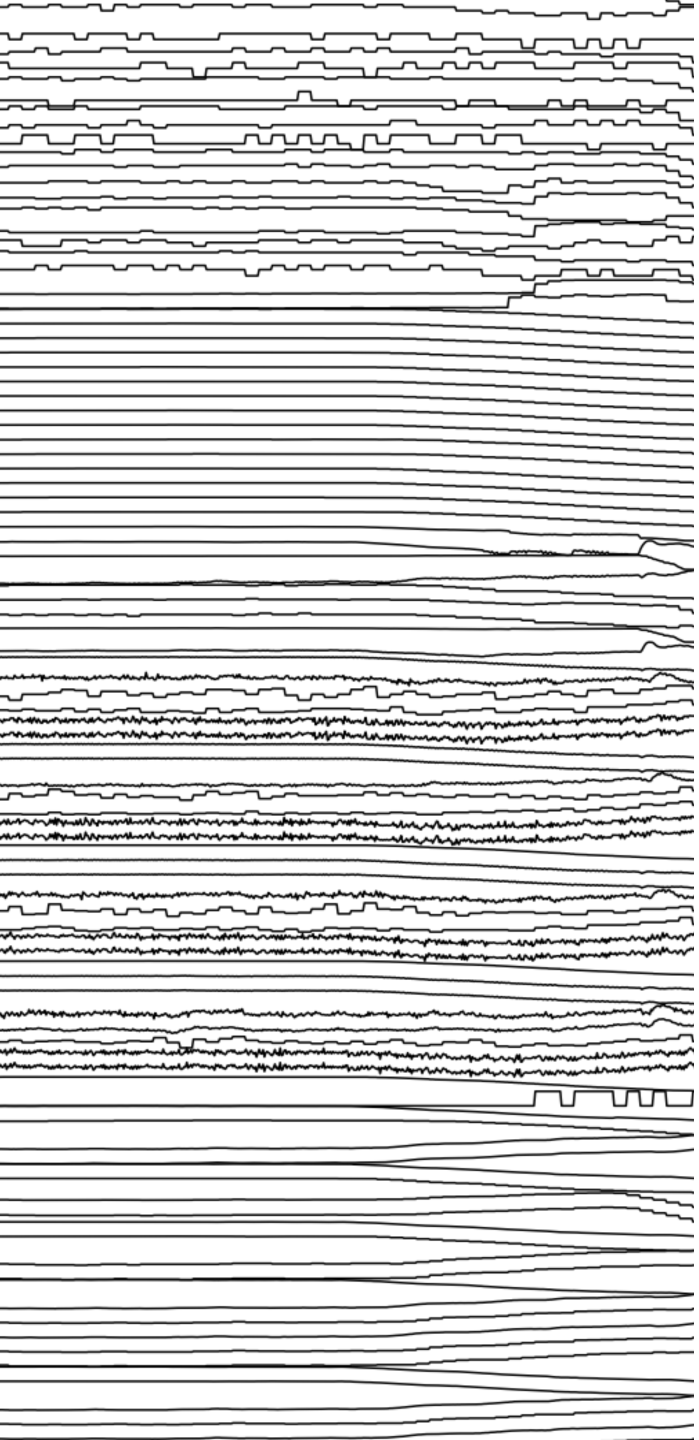

Anomaly Detection in Time Series

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

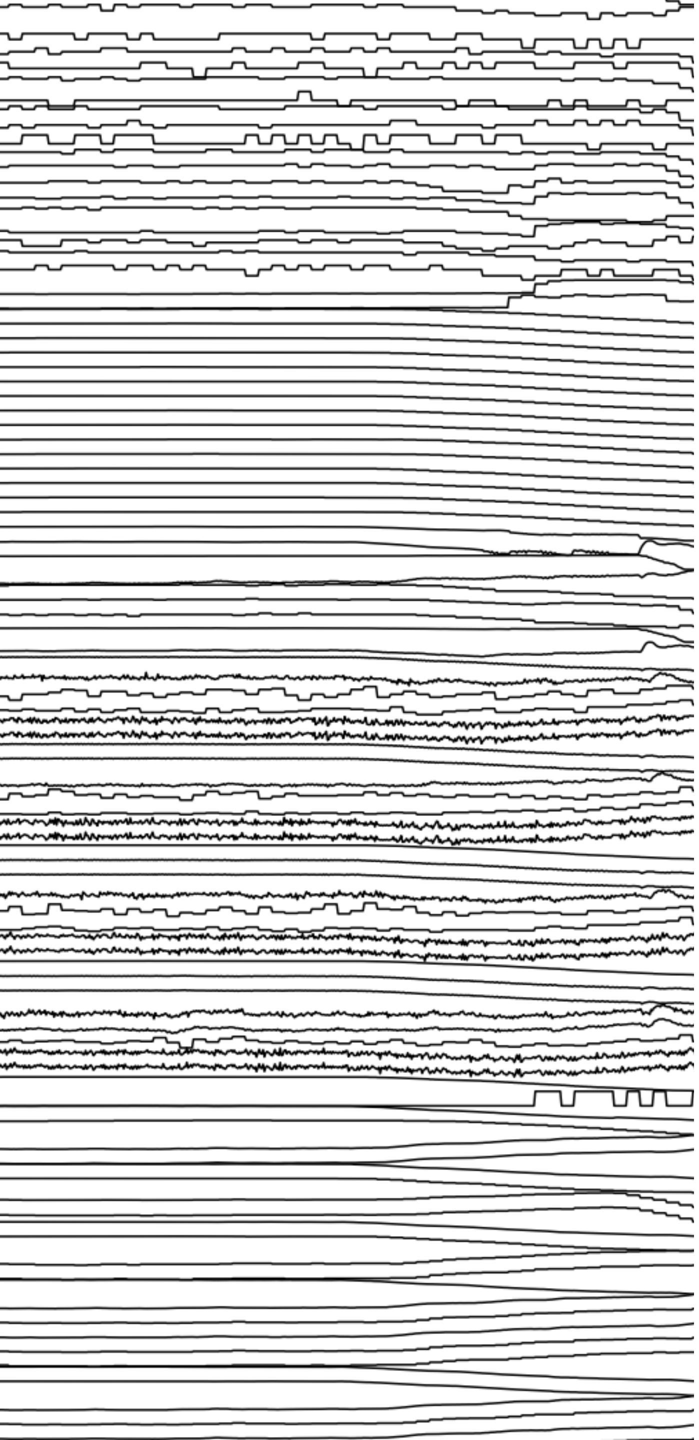

I. Introduction

What is a time series? What is an anomaly?

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

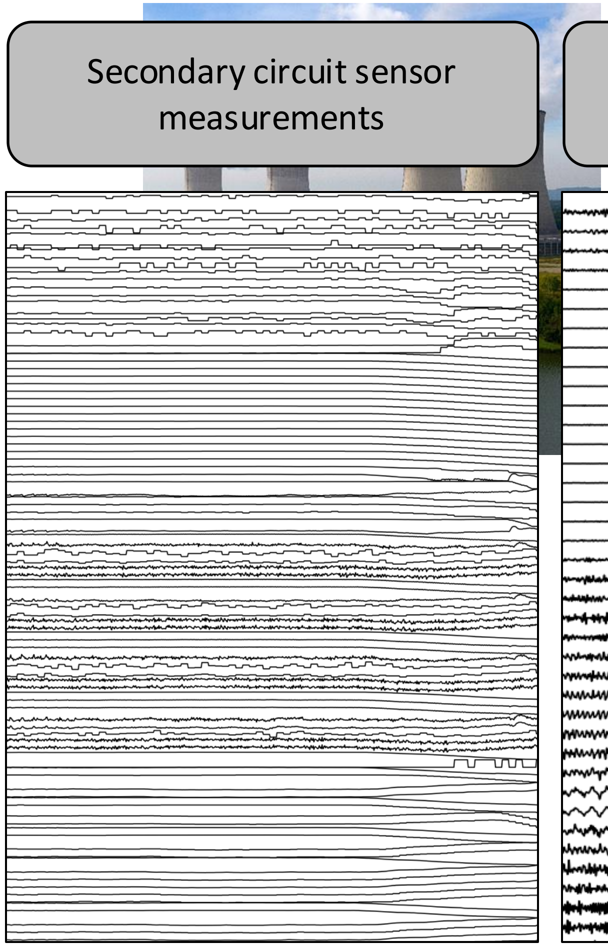

Astrophysics

Fiber-acoustic sensors in the VIRGO north building

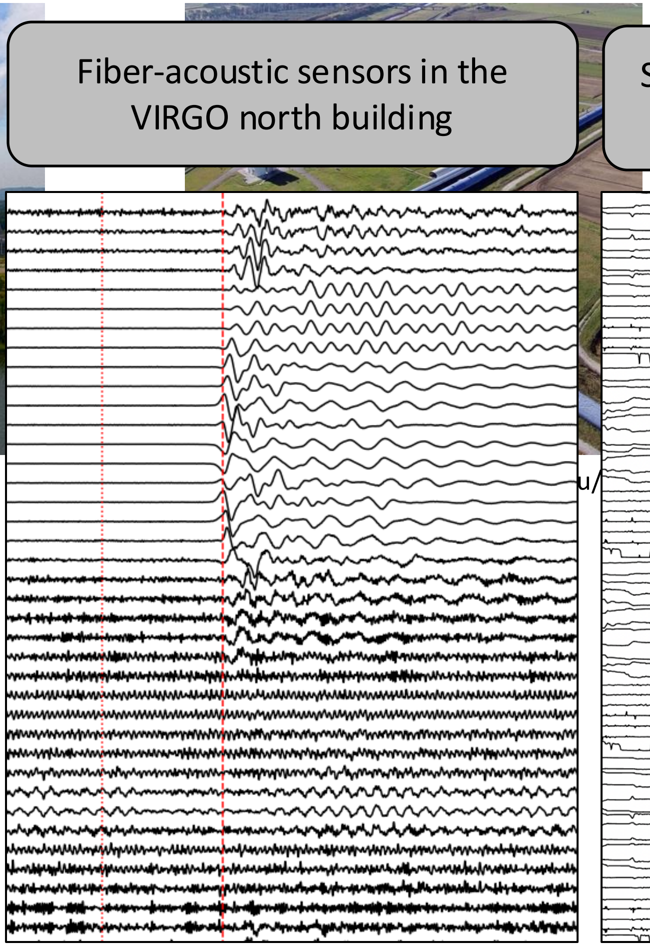

Medicine

Sensor measurements of the Da-Vinci surgery robot

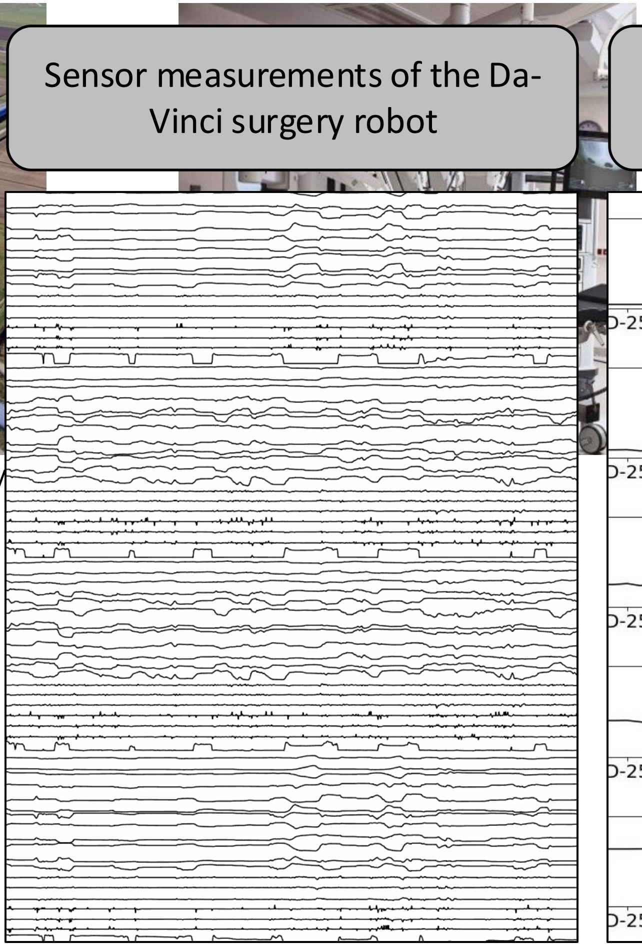

Volcanology

Sensor measurements on le Piton de la Fournaise

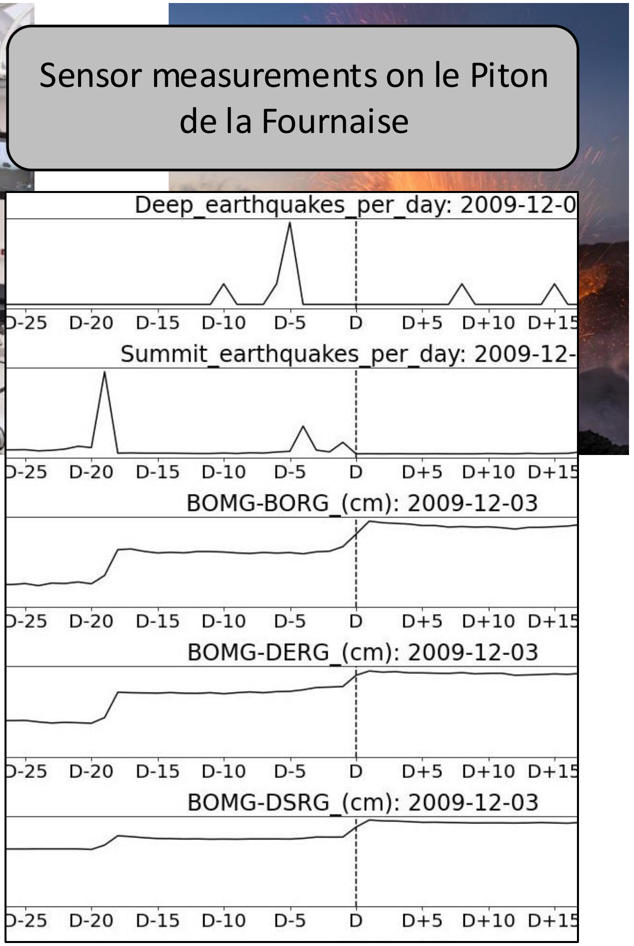

Introduction: *with Important Challenges*

Energy Production

Secondary circuit sensor measurements

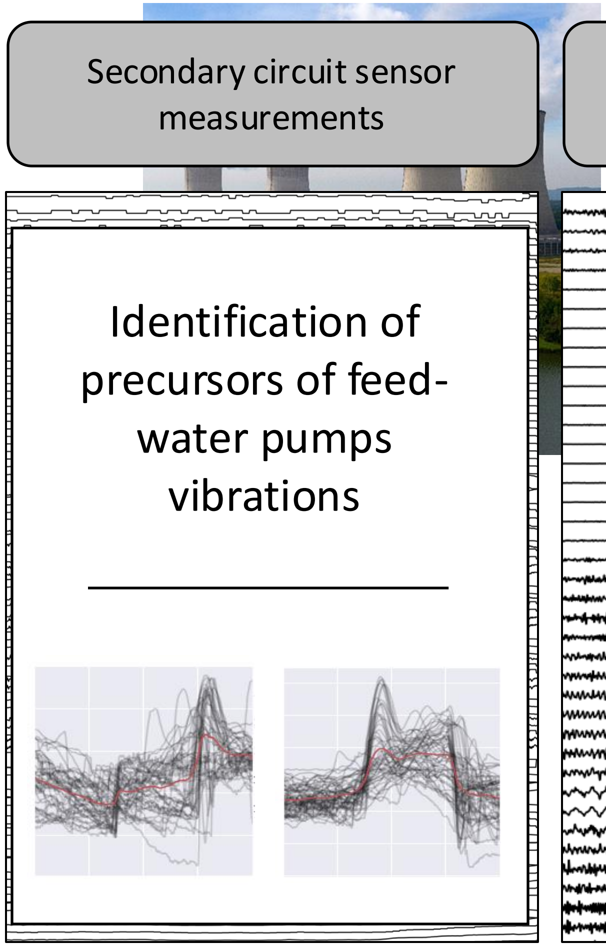

Identification of precursors of feed-water pumps vibrations

Astrophysics

Fiber-acoustic sensors in the VIRGO north building

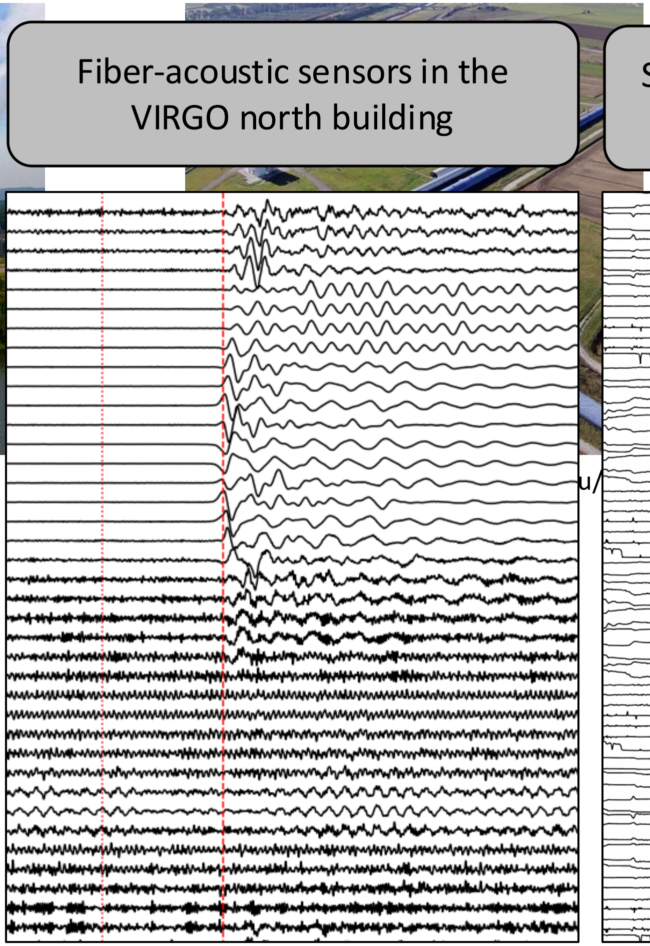

Medicine

Sensor measurements of the Da-Vinci surgery robot

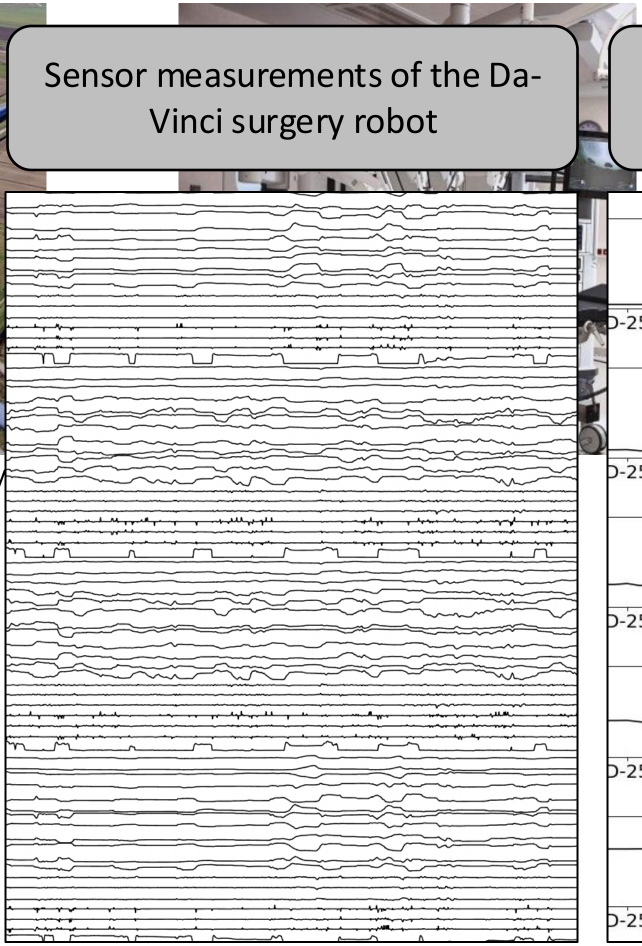

Volcanology

Sensor measurements on le Piton de la Fournaise

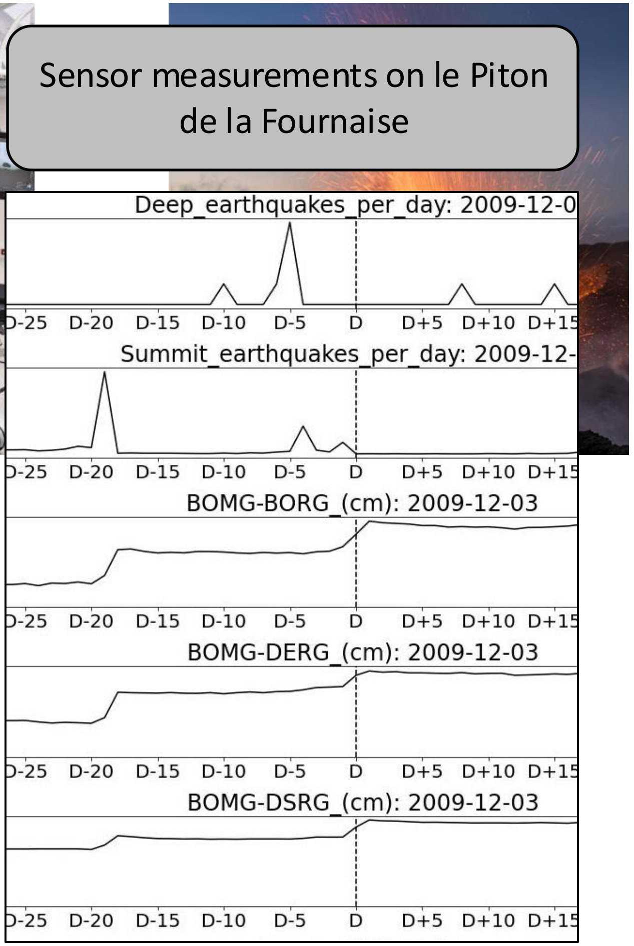

Introduction: *with Important Challenges*

Energy Production

Secondary circuit sensor measurements

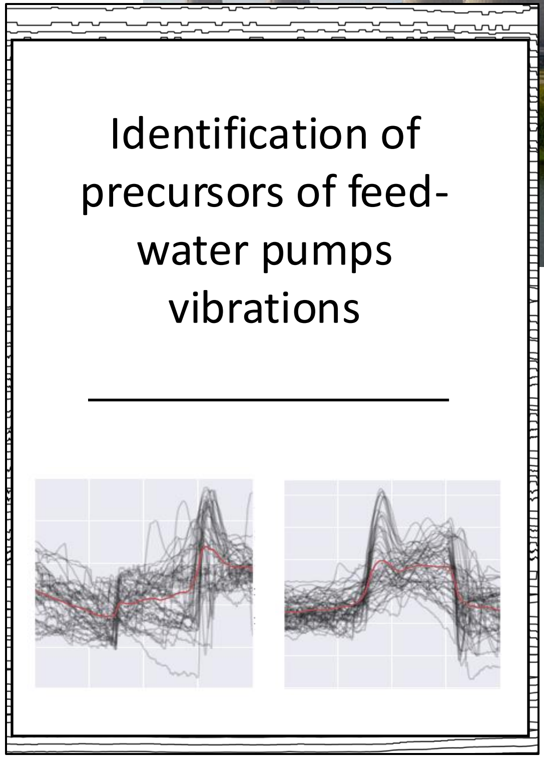

Identification of precursors of feed-water pumps vibrations

Astrophysics

Fiber-acoustic sensors in the VIRGO north building

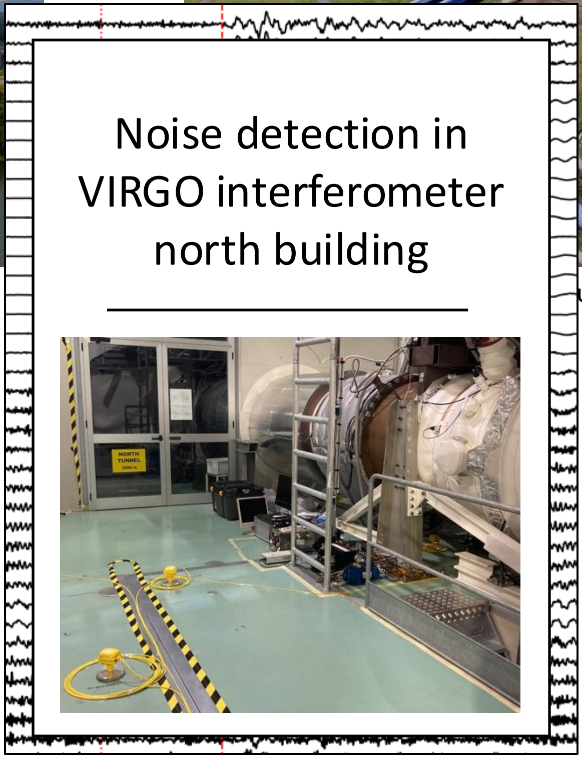

Noise detection in VIRGO interferometer north building

Medicine

Sensor measurements of the Da-Vinci surgery robot

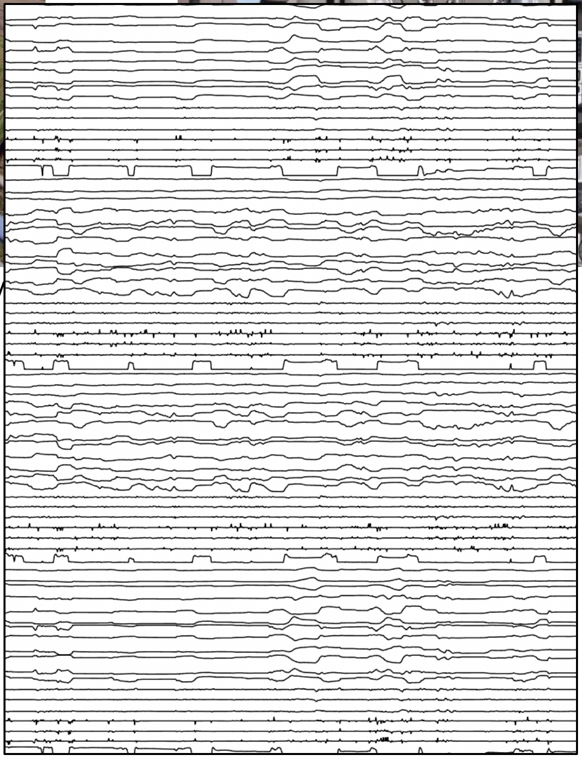

Volcanology

Sensor measurements on le Piton de la Fournaise

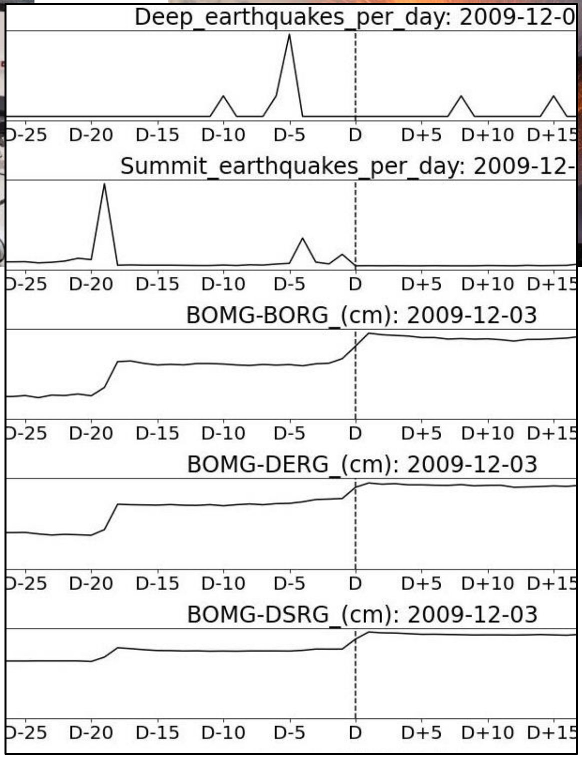

Introduction: *with Important Challenges*

Energy Production

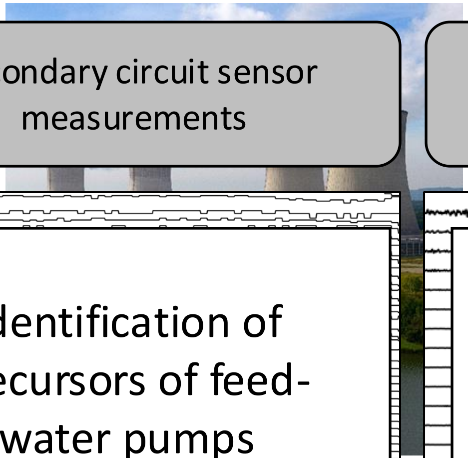

Secondary circuit sensor measurements

Identification of precursors of feed-water pumps vibrations

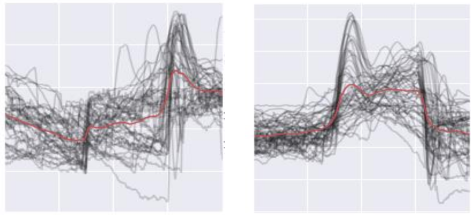

Astrophysics

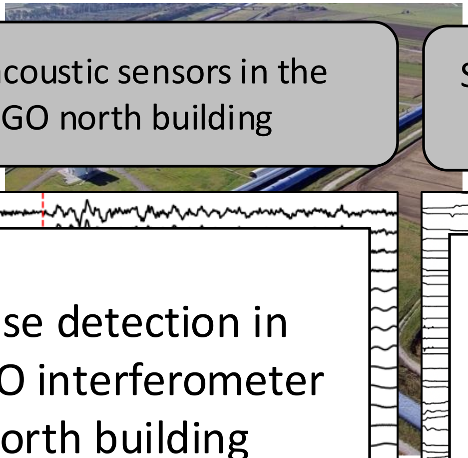

Fiber-acoustic sensors in the VIRGO north building

Noise detection in VIRGO interferometer north building

Medicine

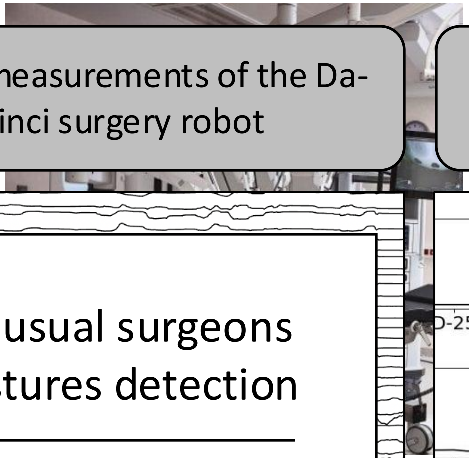

Sensor measurements of the Da-Vinci surgery robot

Unusual surgeons gestures detection

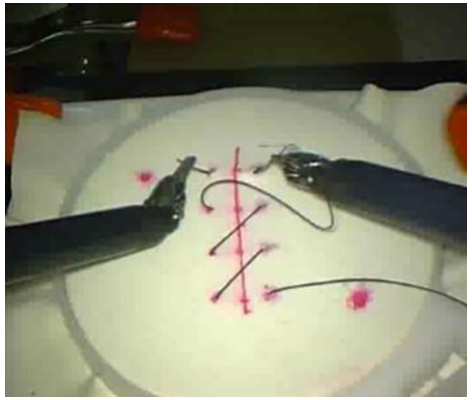

Volcanology

Sensor measurements on le Piton de la Fournaise

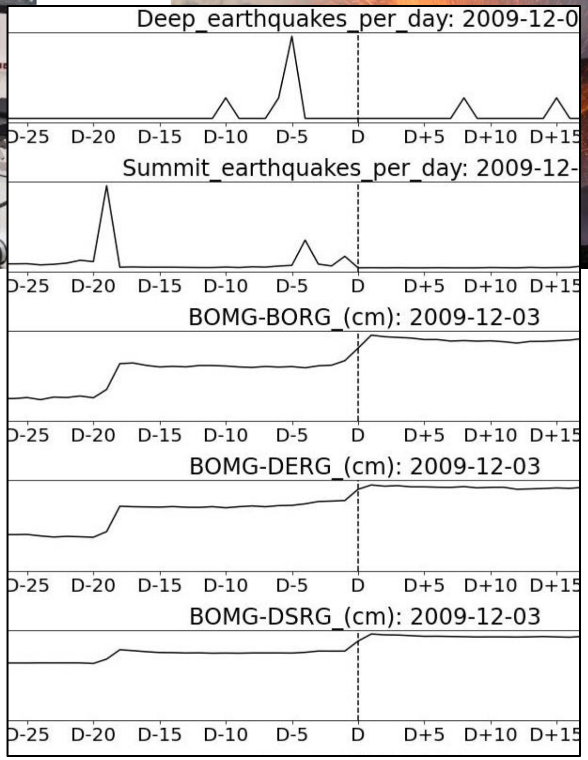

Introduction: *with Important Challenges*

Energy Production

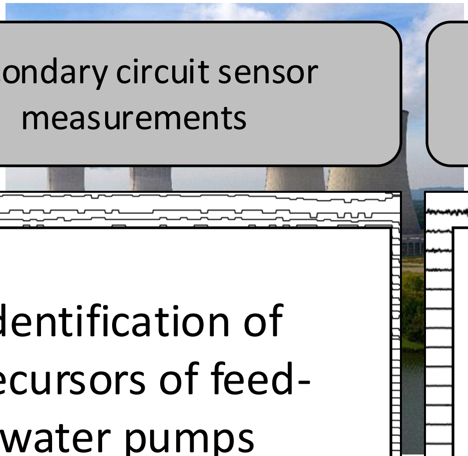

Secondary circuit sensor measurements

Identification of precursors of feed-water pumps vibrations

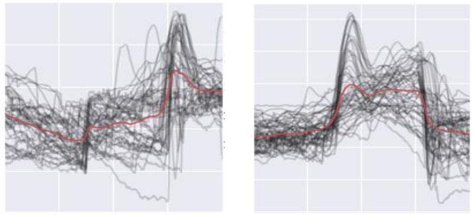

Astrophysics

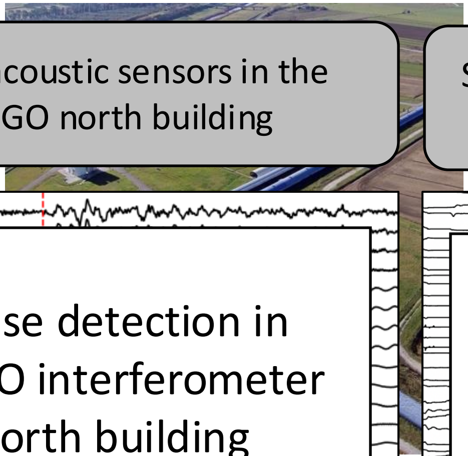

Fiber-acoustic sensors in the VIRGO north building

Noise detection in VIRGO interferometer north building

Medicine

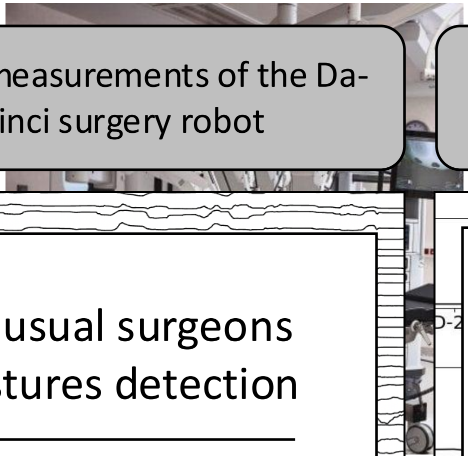

Sensor measurements of the Da-Vinci surgery robot

Unusual surgeons gestures detection

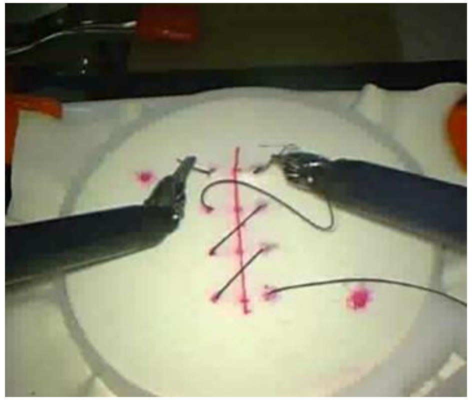

Volcanology

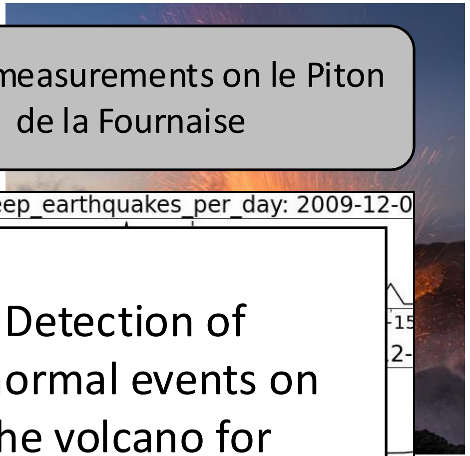

Sensor measurements on le Piton de la Fournaise

Detection of abnormal events on the volcano for predicting eruptions

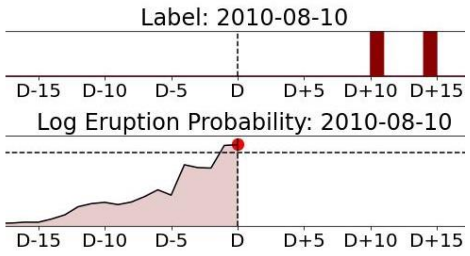

Introduction: *with Important Challenges*

Large-scale time series database

Energy Production

Edf.fr: tinyurl.com/yc7x5xje

Introduction: *with Important Challenges*

Large-scale time series database

Energy Production

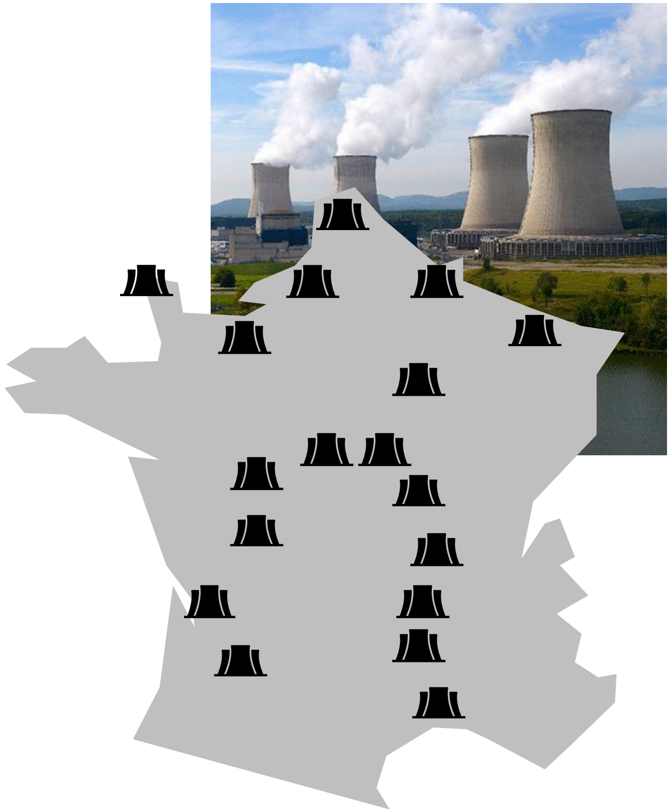

Edf.fr: tinyurl.com/yc7x5xje

Example of Nuclear production

- 58 nuclear power plants across France

Introduction: *with Important Challenges*

Large-scale time series database

Energy Production

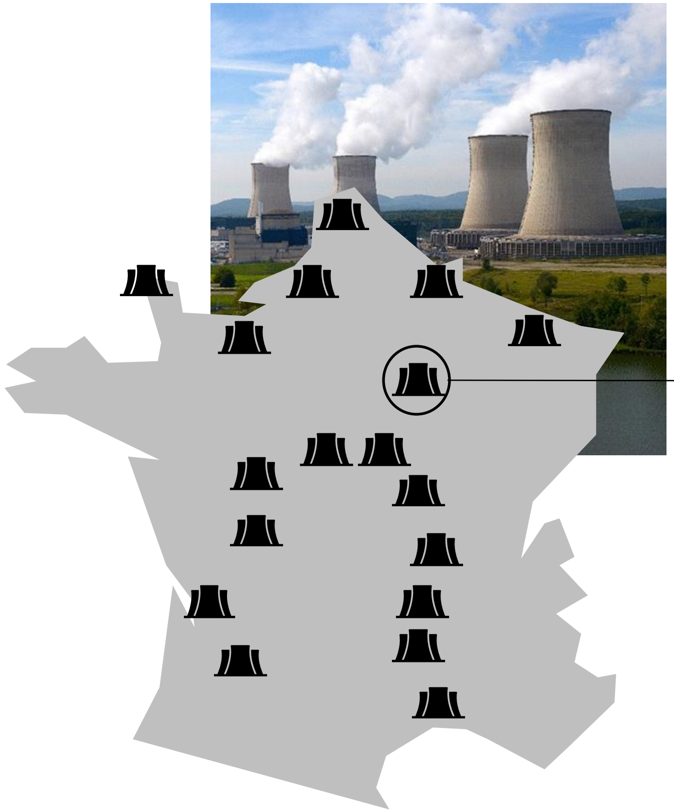

Edf.fr: tinyurl.com/yc7x5xje

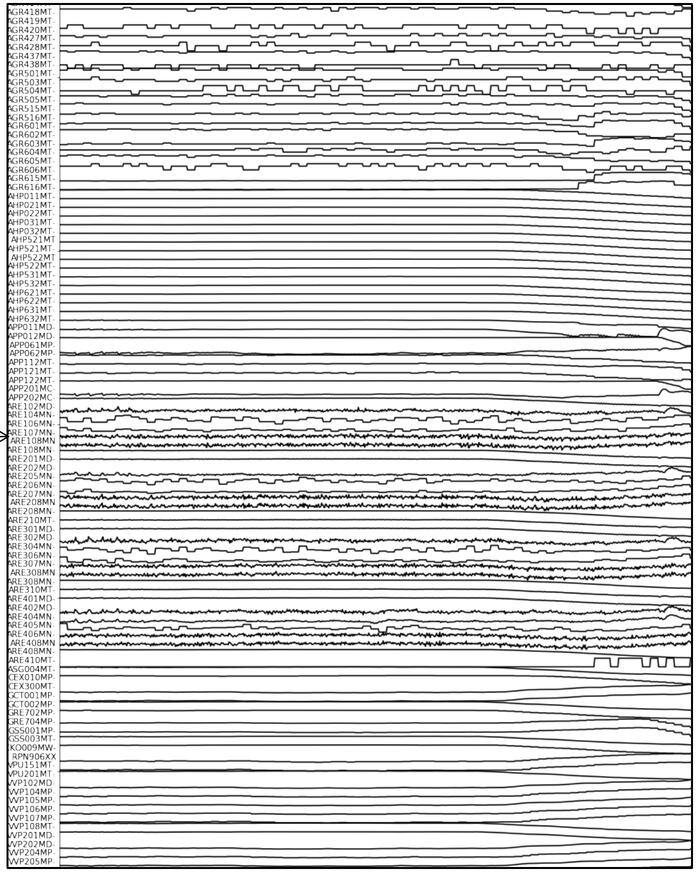

Example of Nuclear production

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

A total of 1.5 PetaBytes

Introduction: *with Important Challenges*

Large-scale time series database

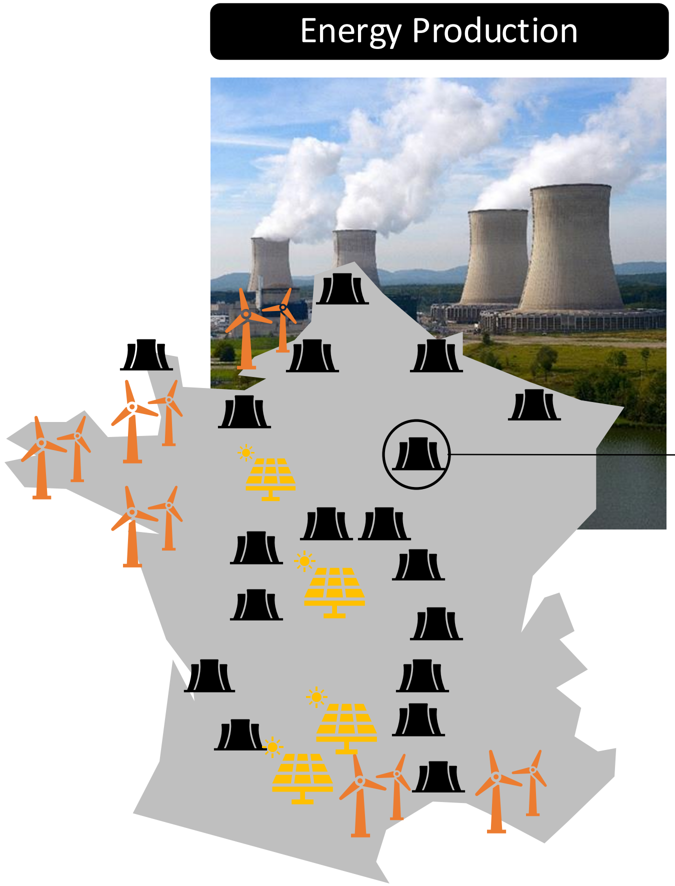

Edf.fr: tinyurl.com/yc7x5xje

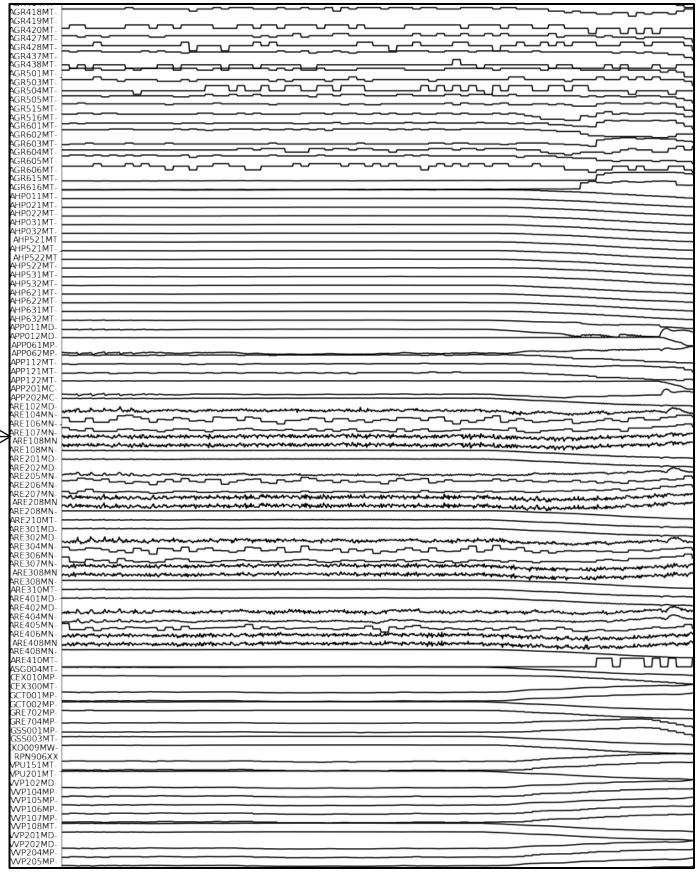

Example of Nuclear production

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

A total of 1.5 PetaBytes

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)*

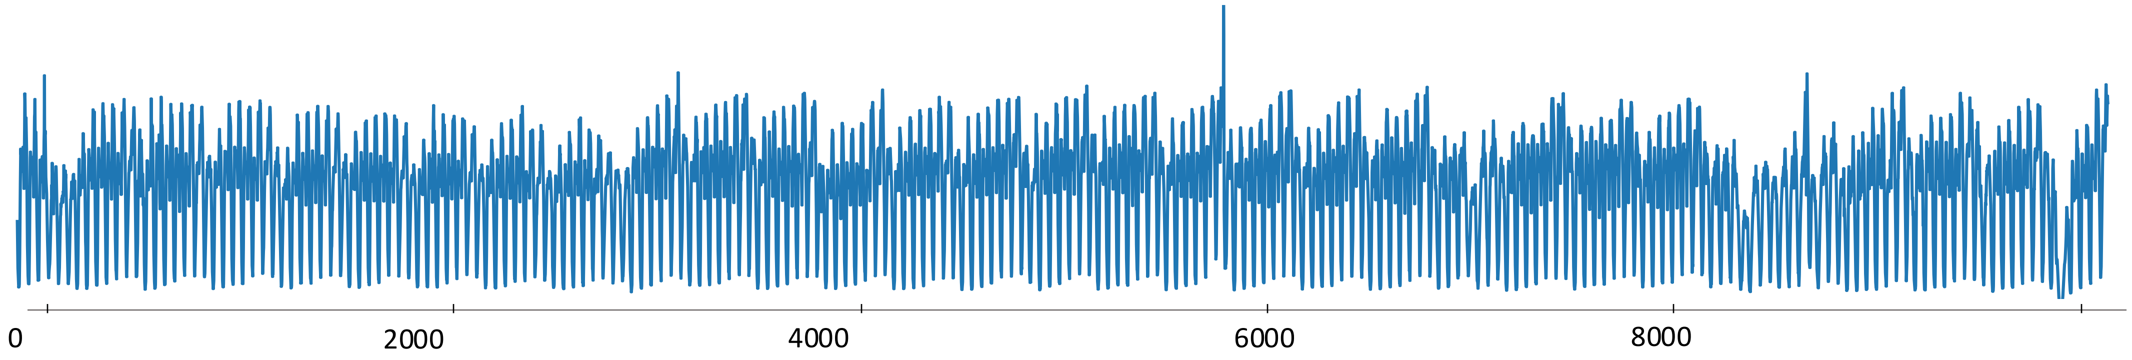

Introduction: *Anomaly Detection in Time Series*

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

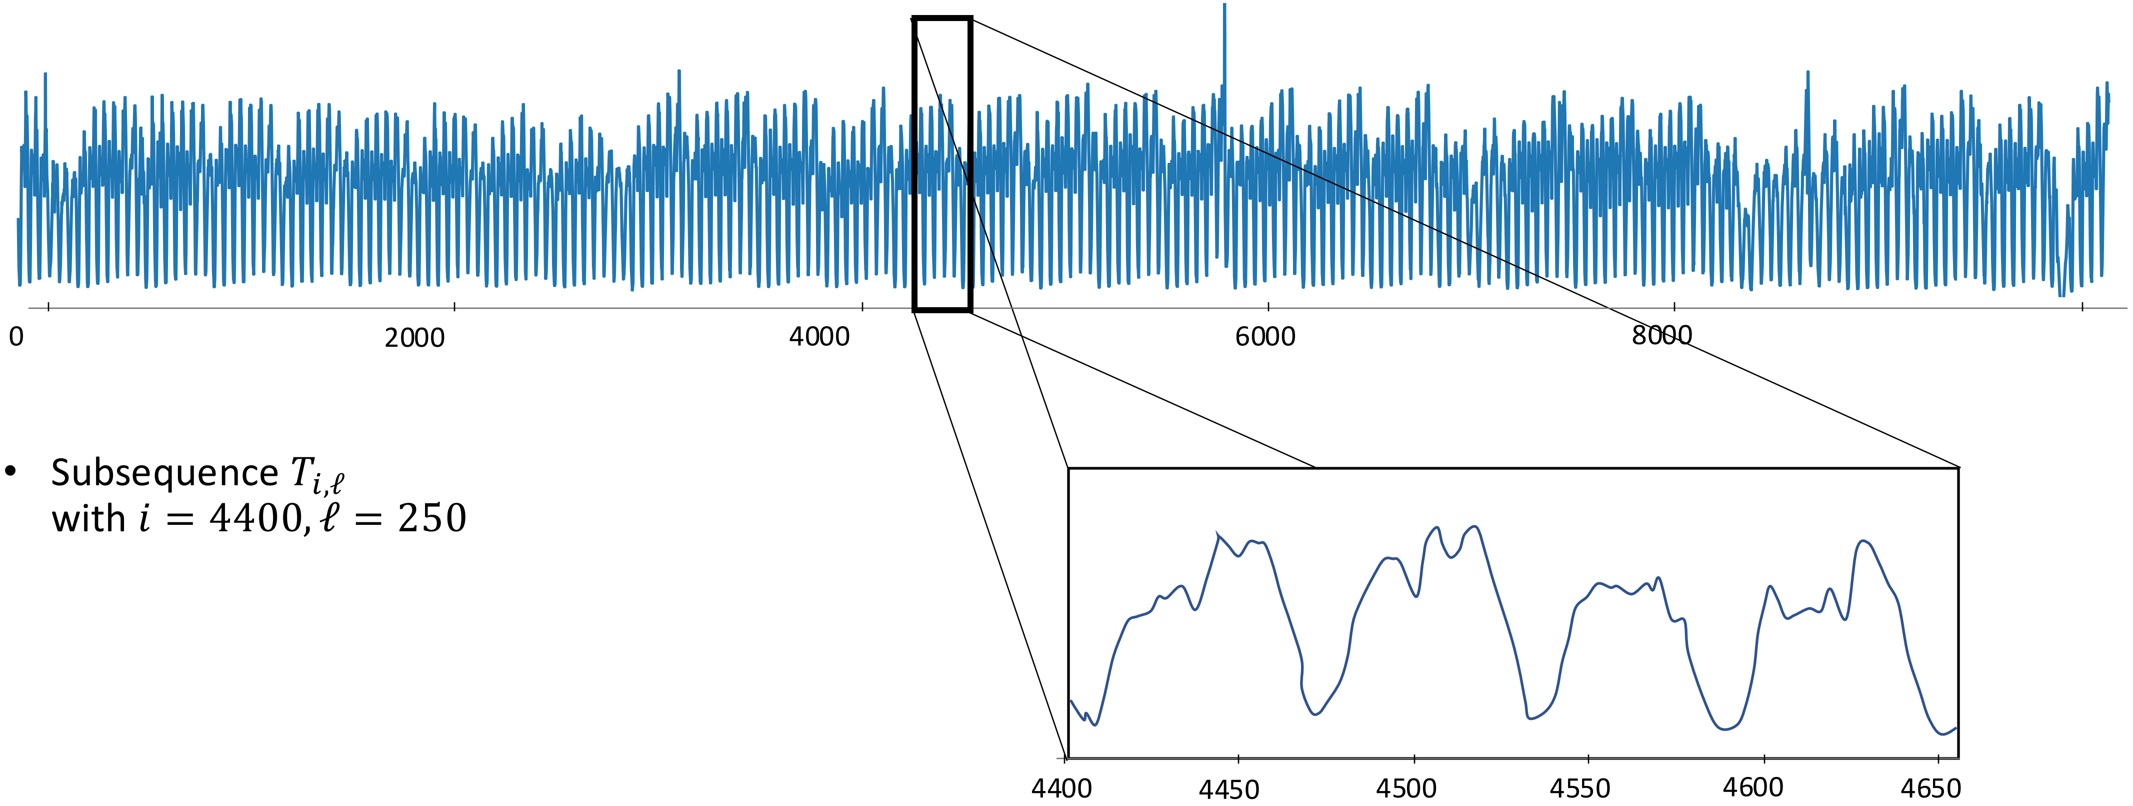

- Subsequence $T_{i,\ell}$
  with $i = 4400, \ell = 250$

Introduction: *Anomaly Detection in Time Series*

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

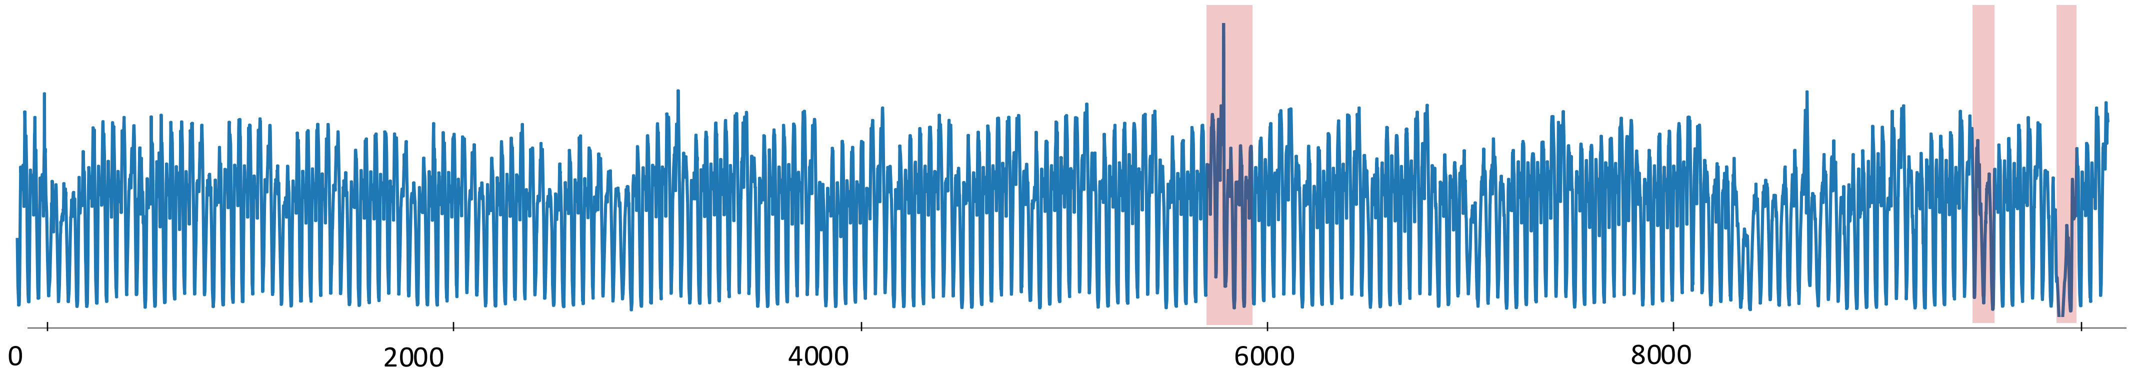

- *Anomaly: rare point or sequence (of a given length) potentially non-desired*

Introduction: *Anomaly Detection in Time Series*

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

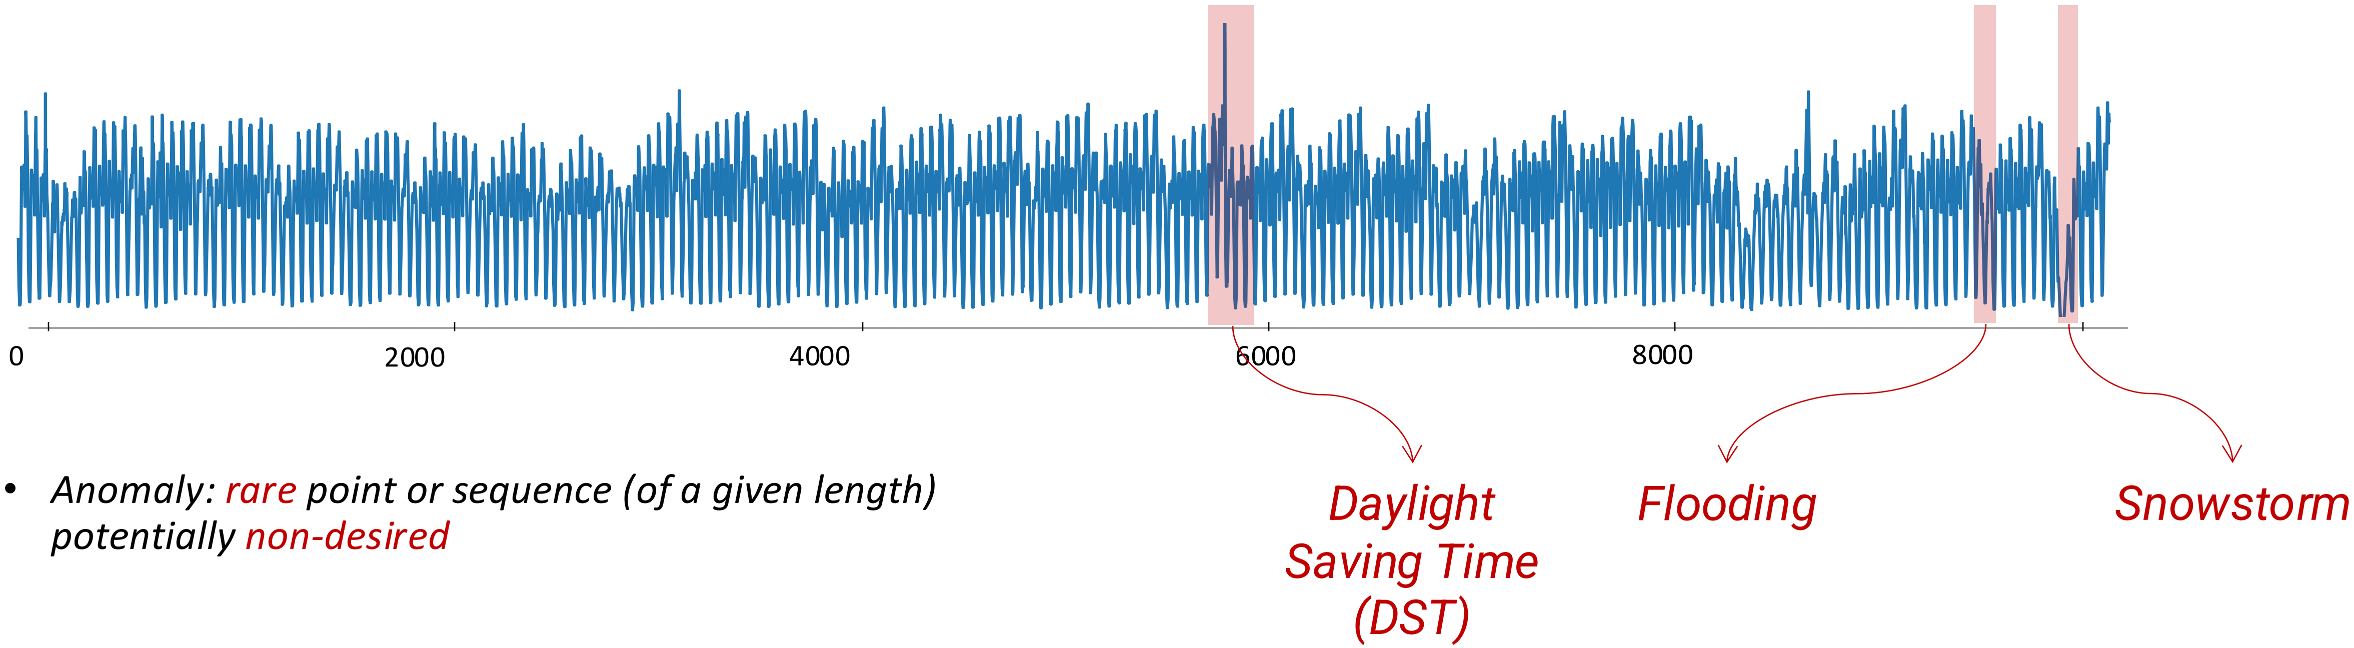

- *Anomaly: rare point or sequence (of a given length) potentially non-desired*

Introduction: *Anomaly Detection in Time Series*

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

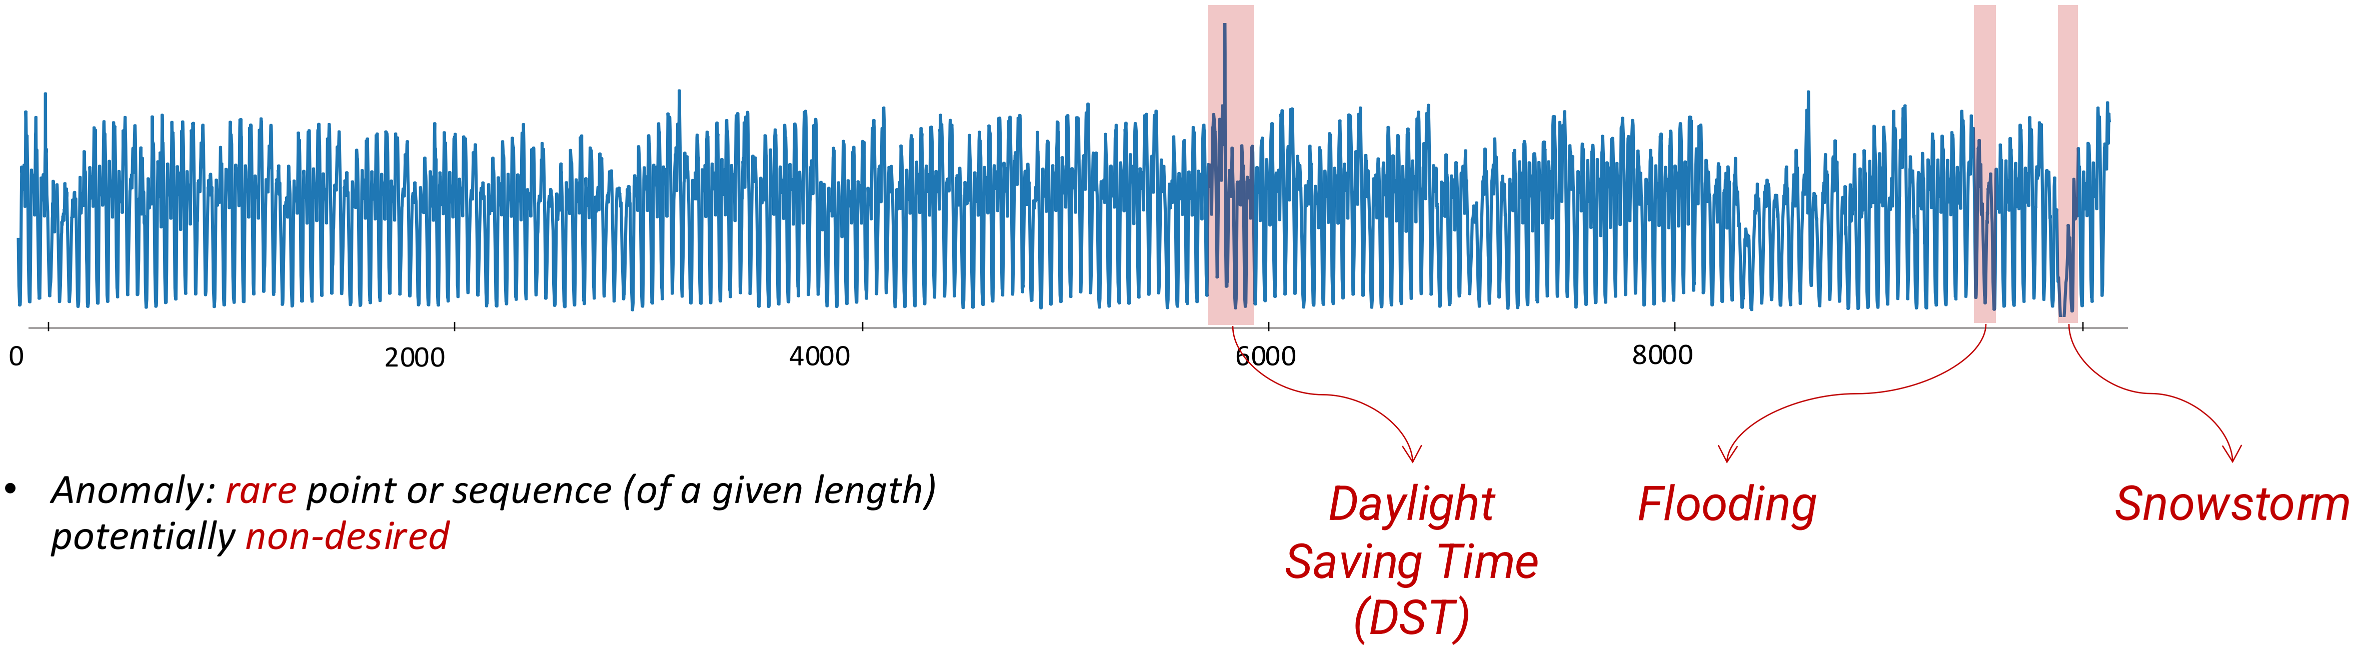

- *Anomaly: rare point or sequence (of a given length) potentially non-desired*

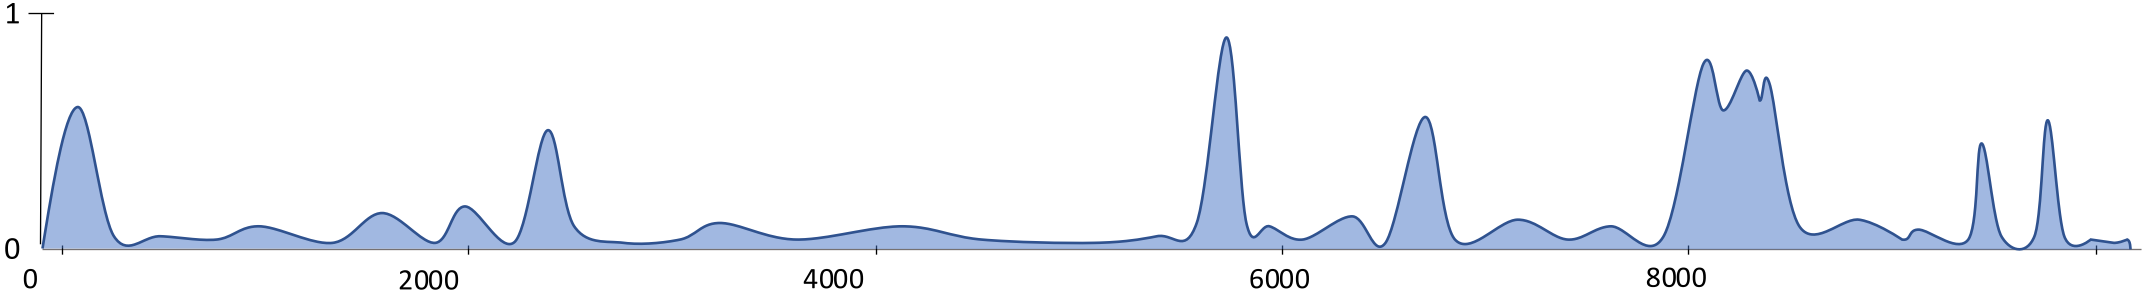

Anomaly score S_T

Introduction: *Anomaly Detection in Time Series*

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

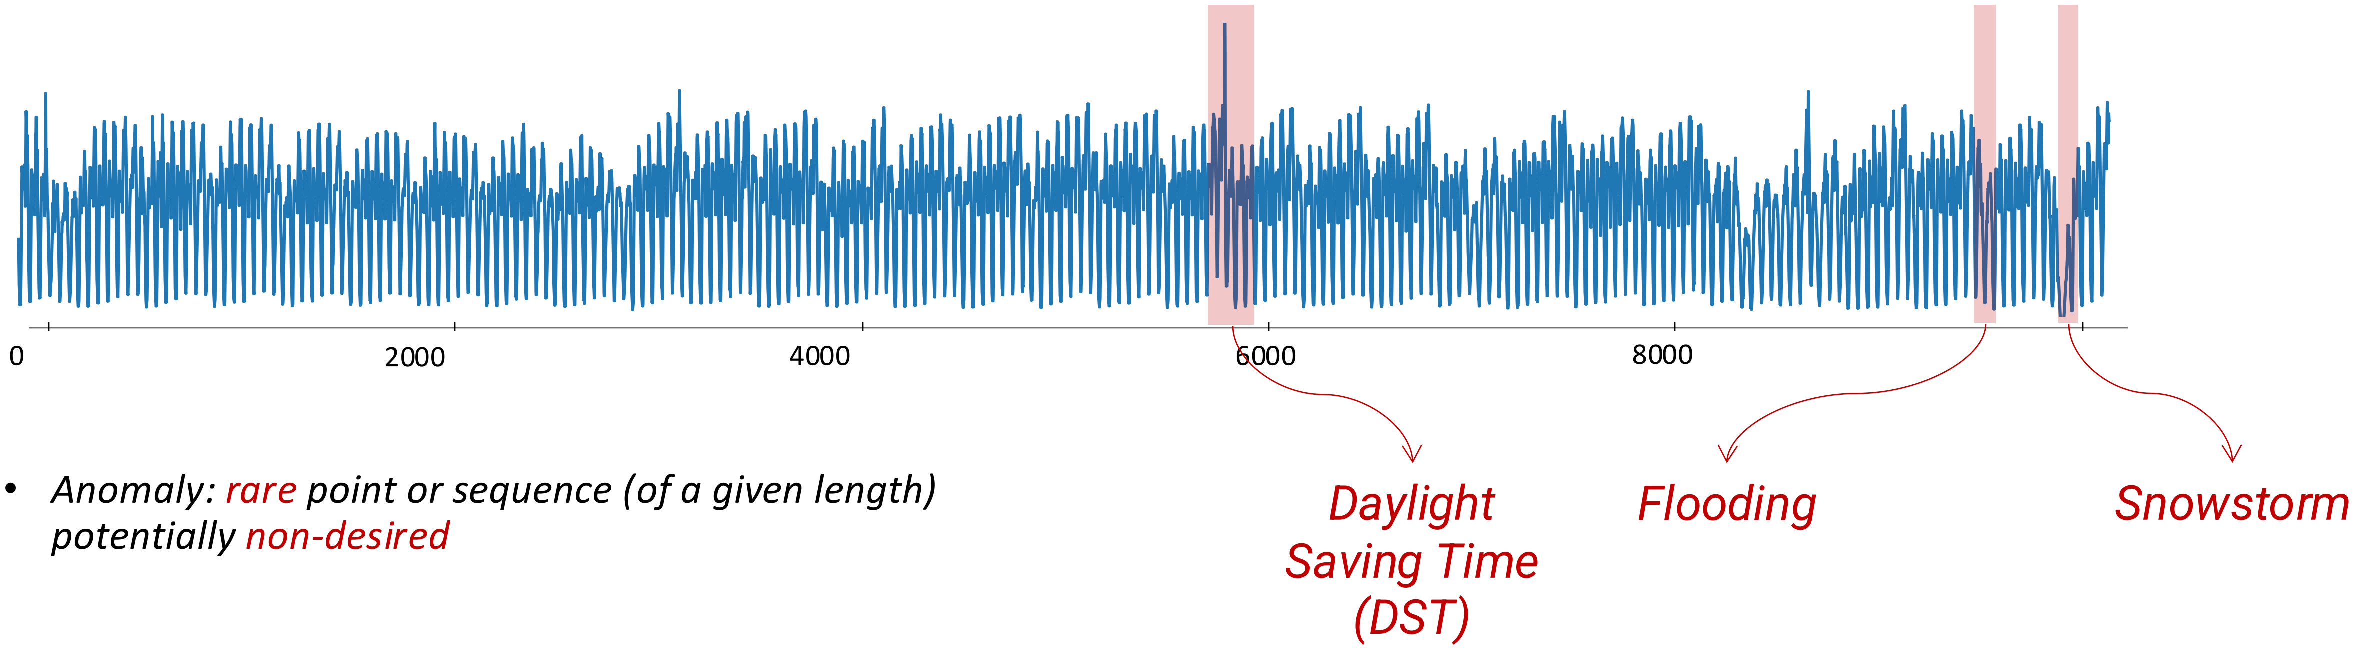

- *Anomaly: rare point or sequence (of a given length) potentially non-desired*

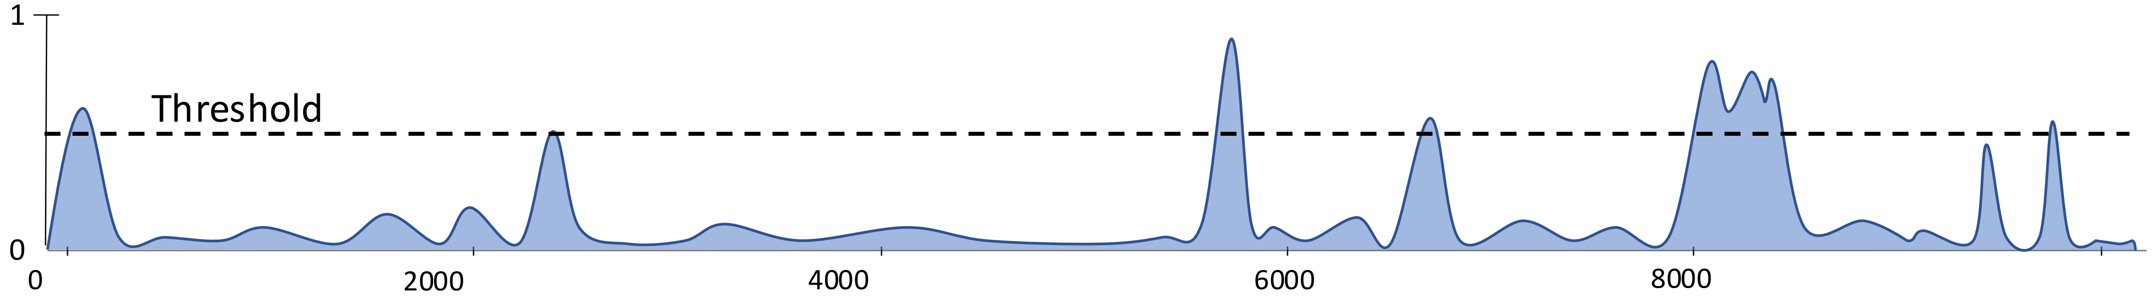

Anomaly score S_T

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*

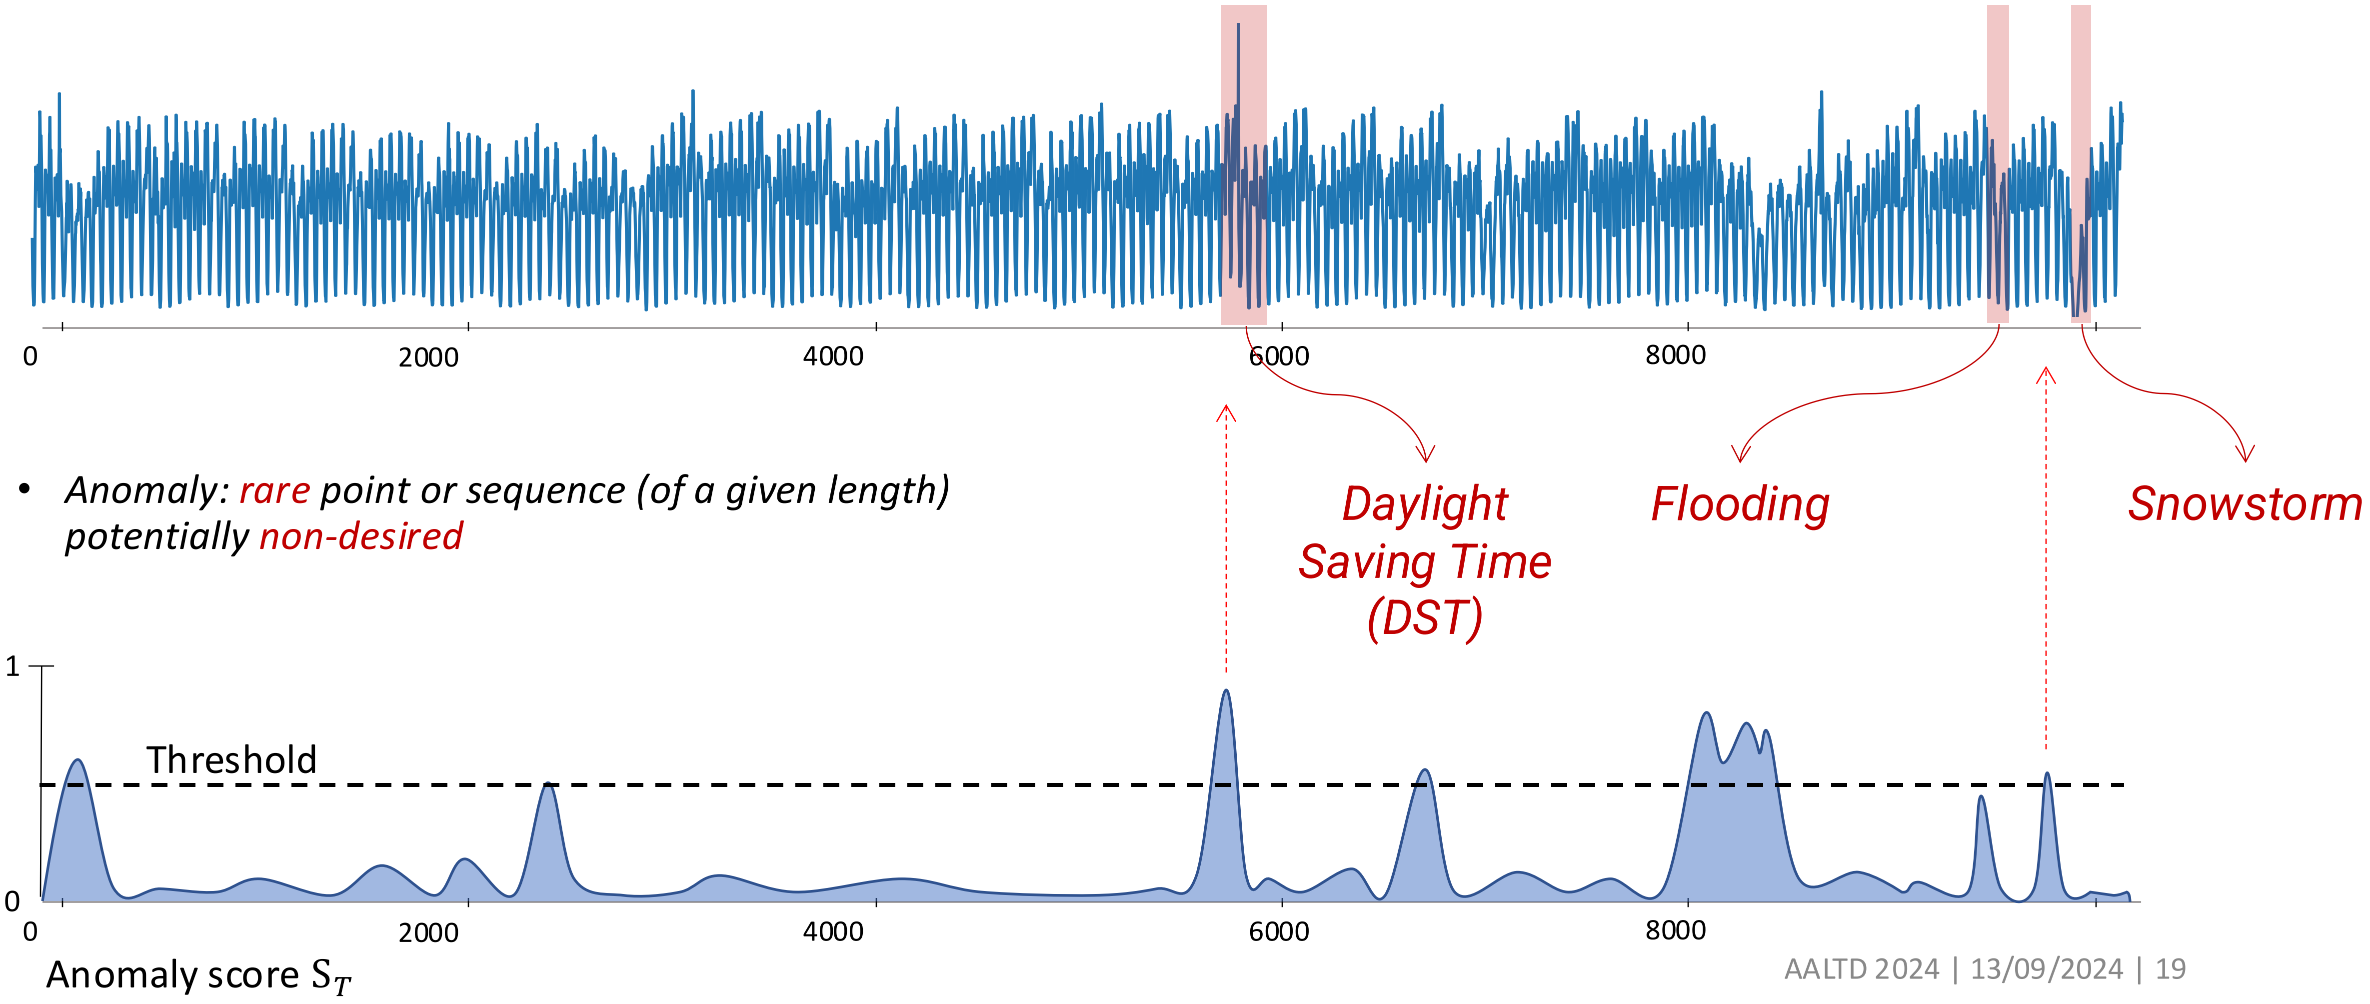

Introduction: *Type of anomalies*

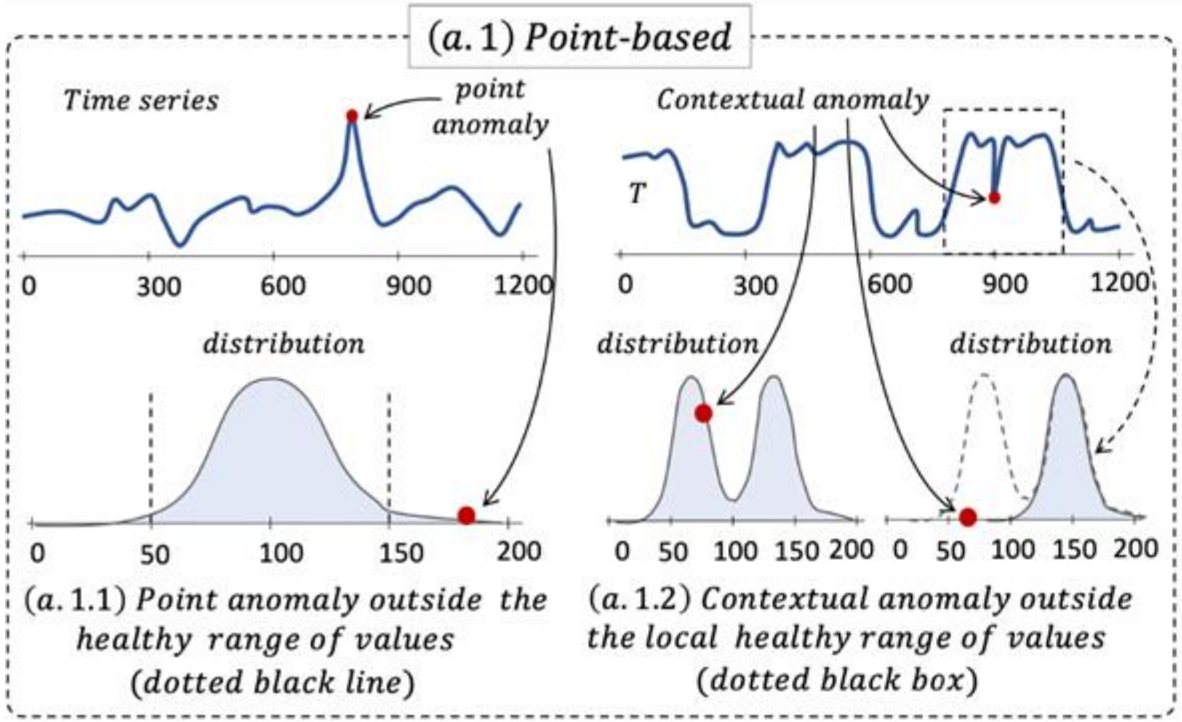

Introduction: *Type of anomalies*

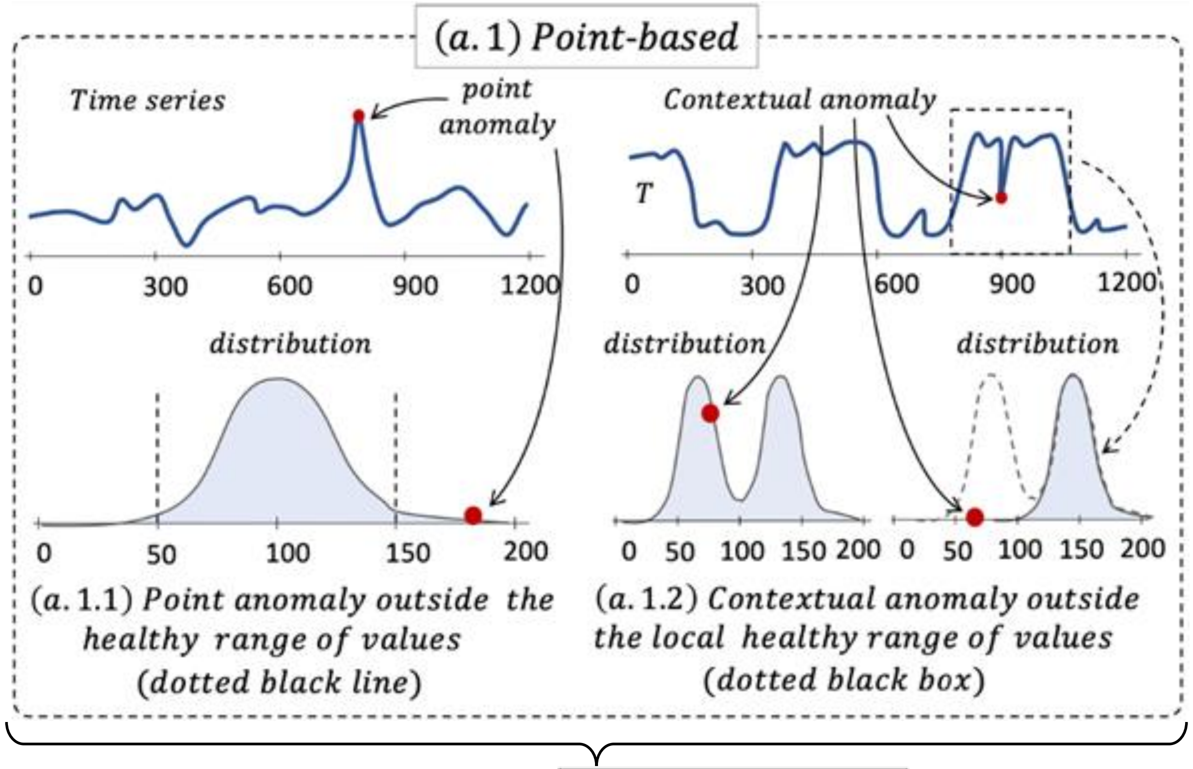

Example of point-based anomaly [1]

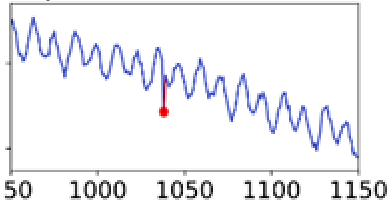

Introduction: *Type of anomalies*

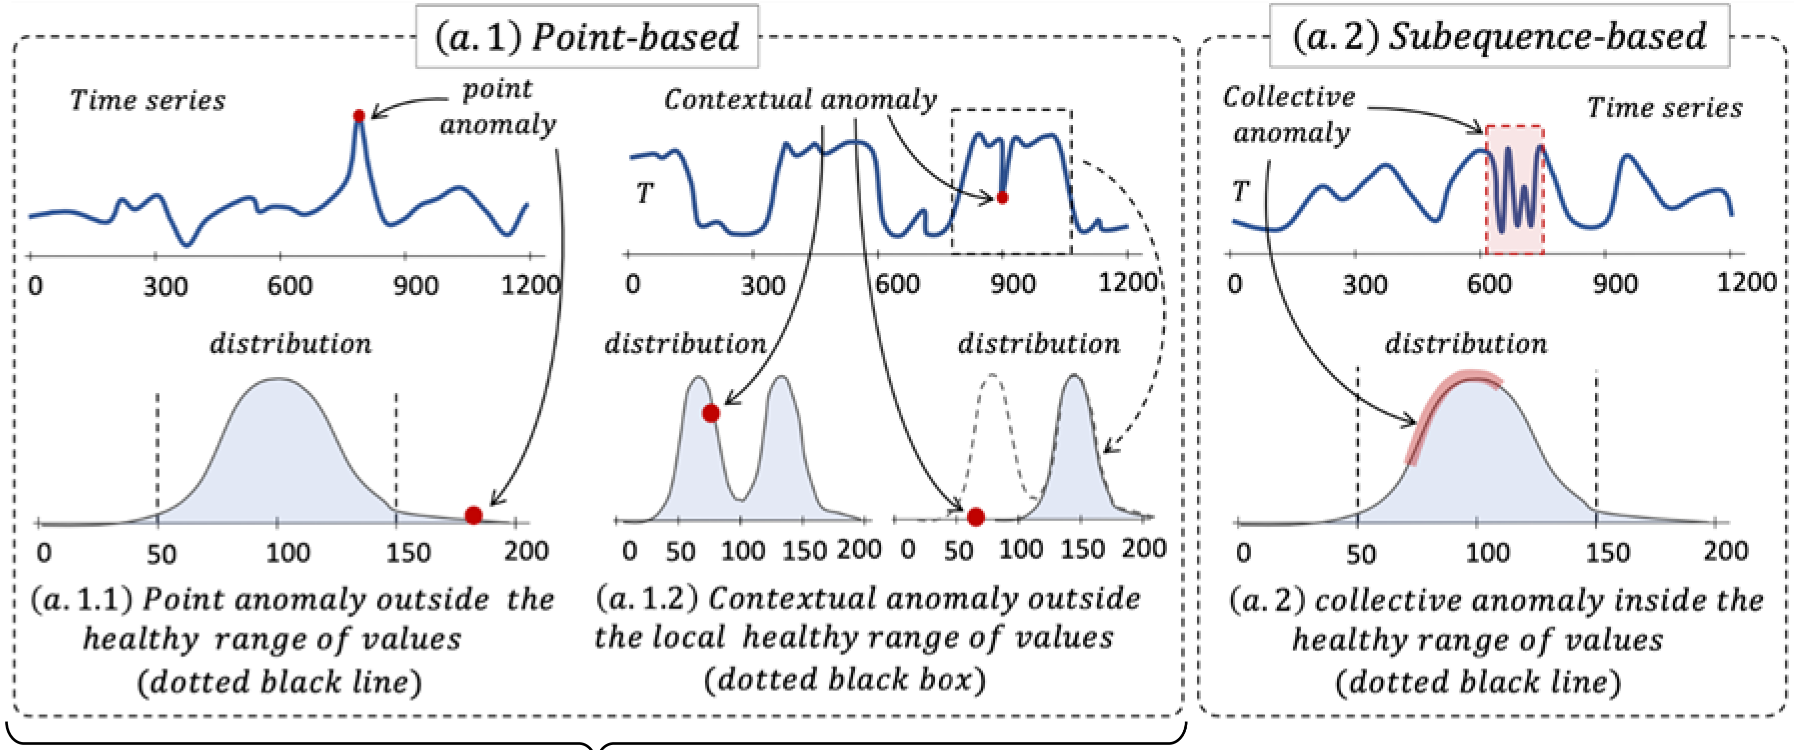

Example of point-based anomaly [1]

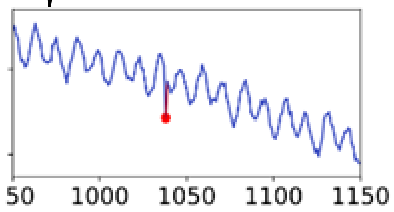

Introduction: *Type of anomalies*

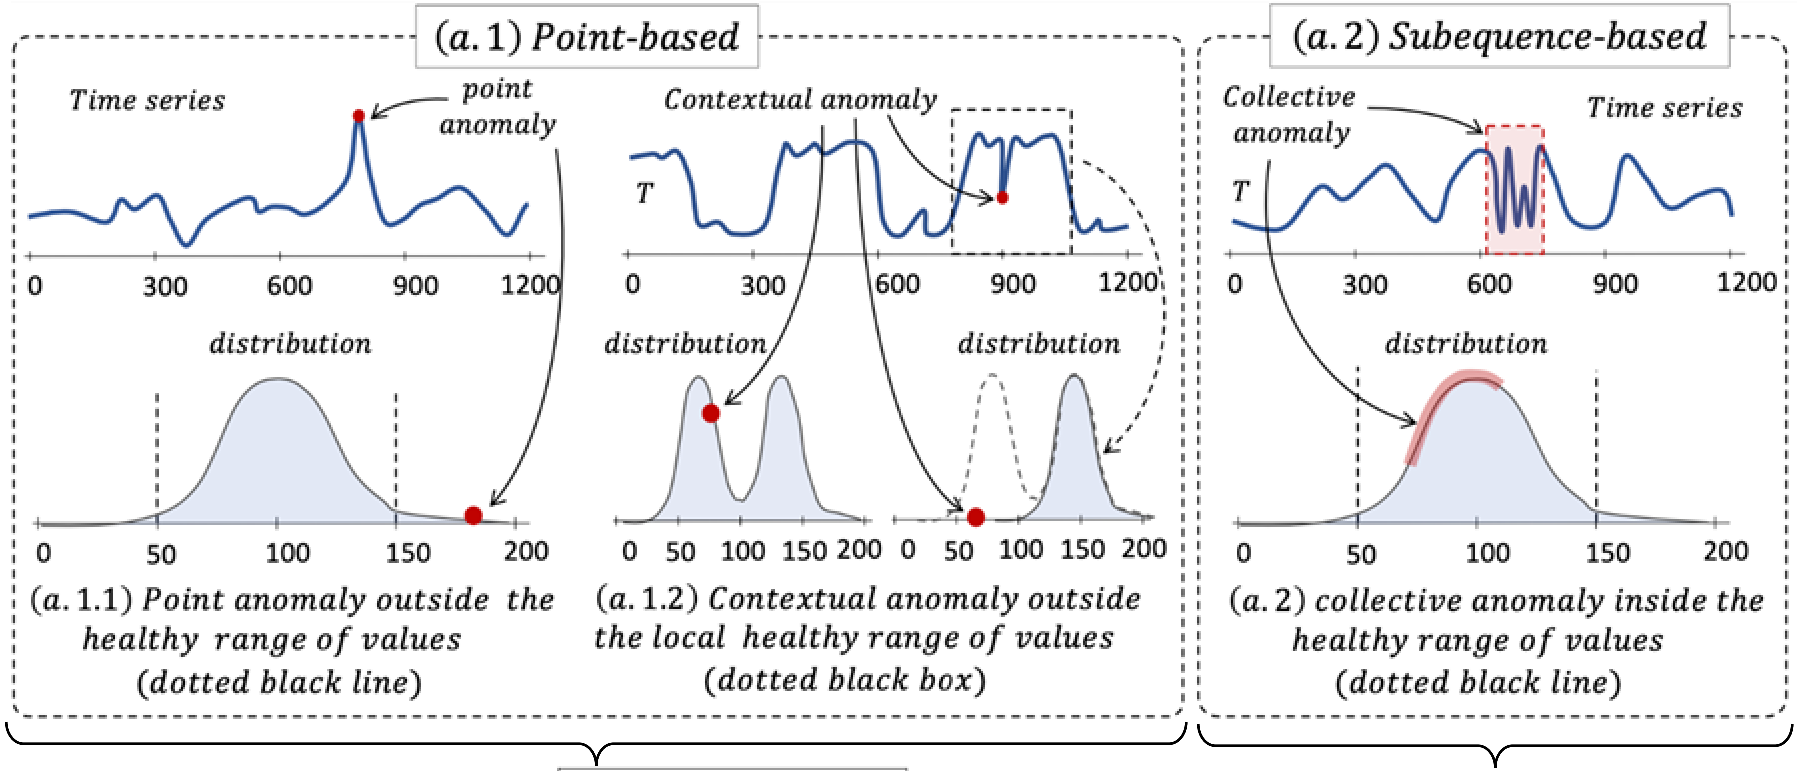

Example of point-based anomaly [1]

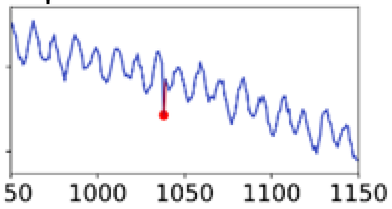

Example of subsequence-based anomaly [2]

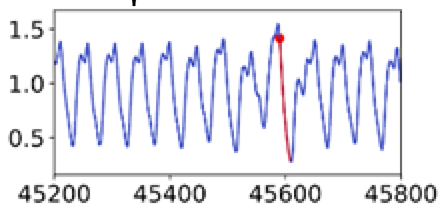

Introduction: *Type of anomalies*

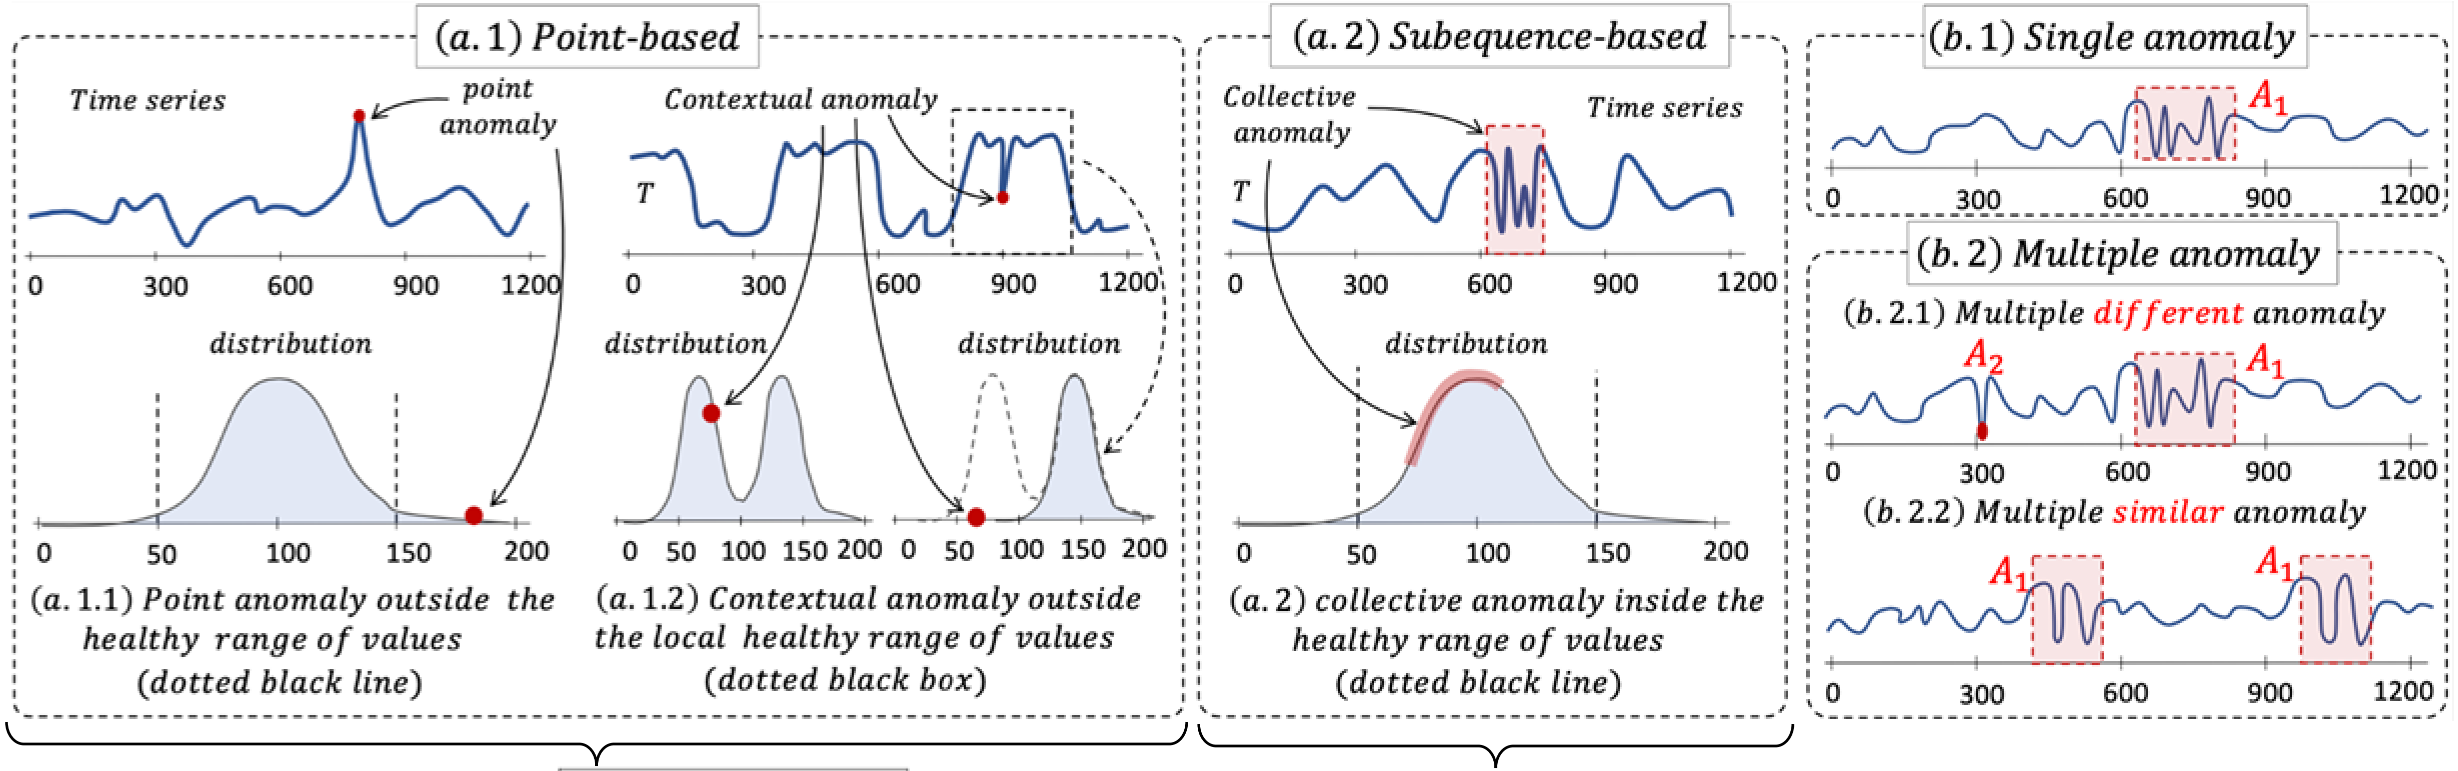

Example of point-based anomaly [1]

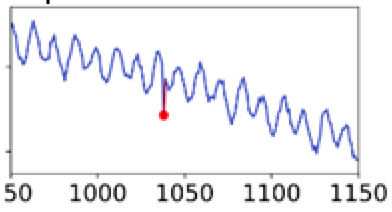

Example of subsequence-based anomaly [2]

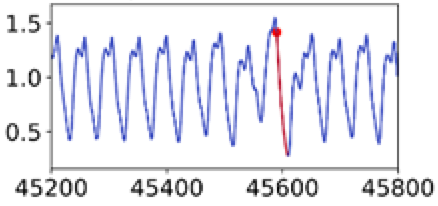

Introduction: *Type of anomalies*

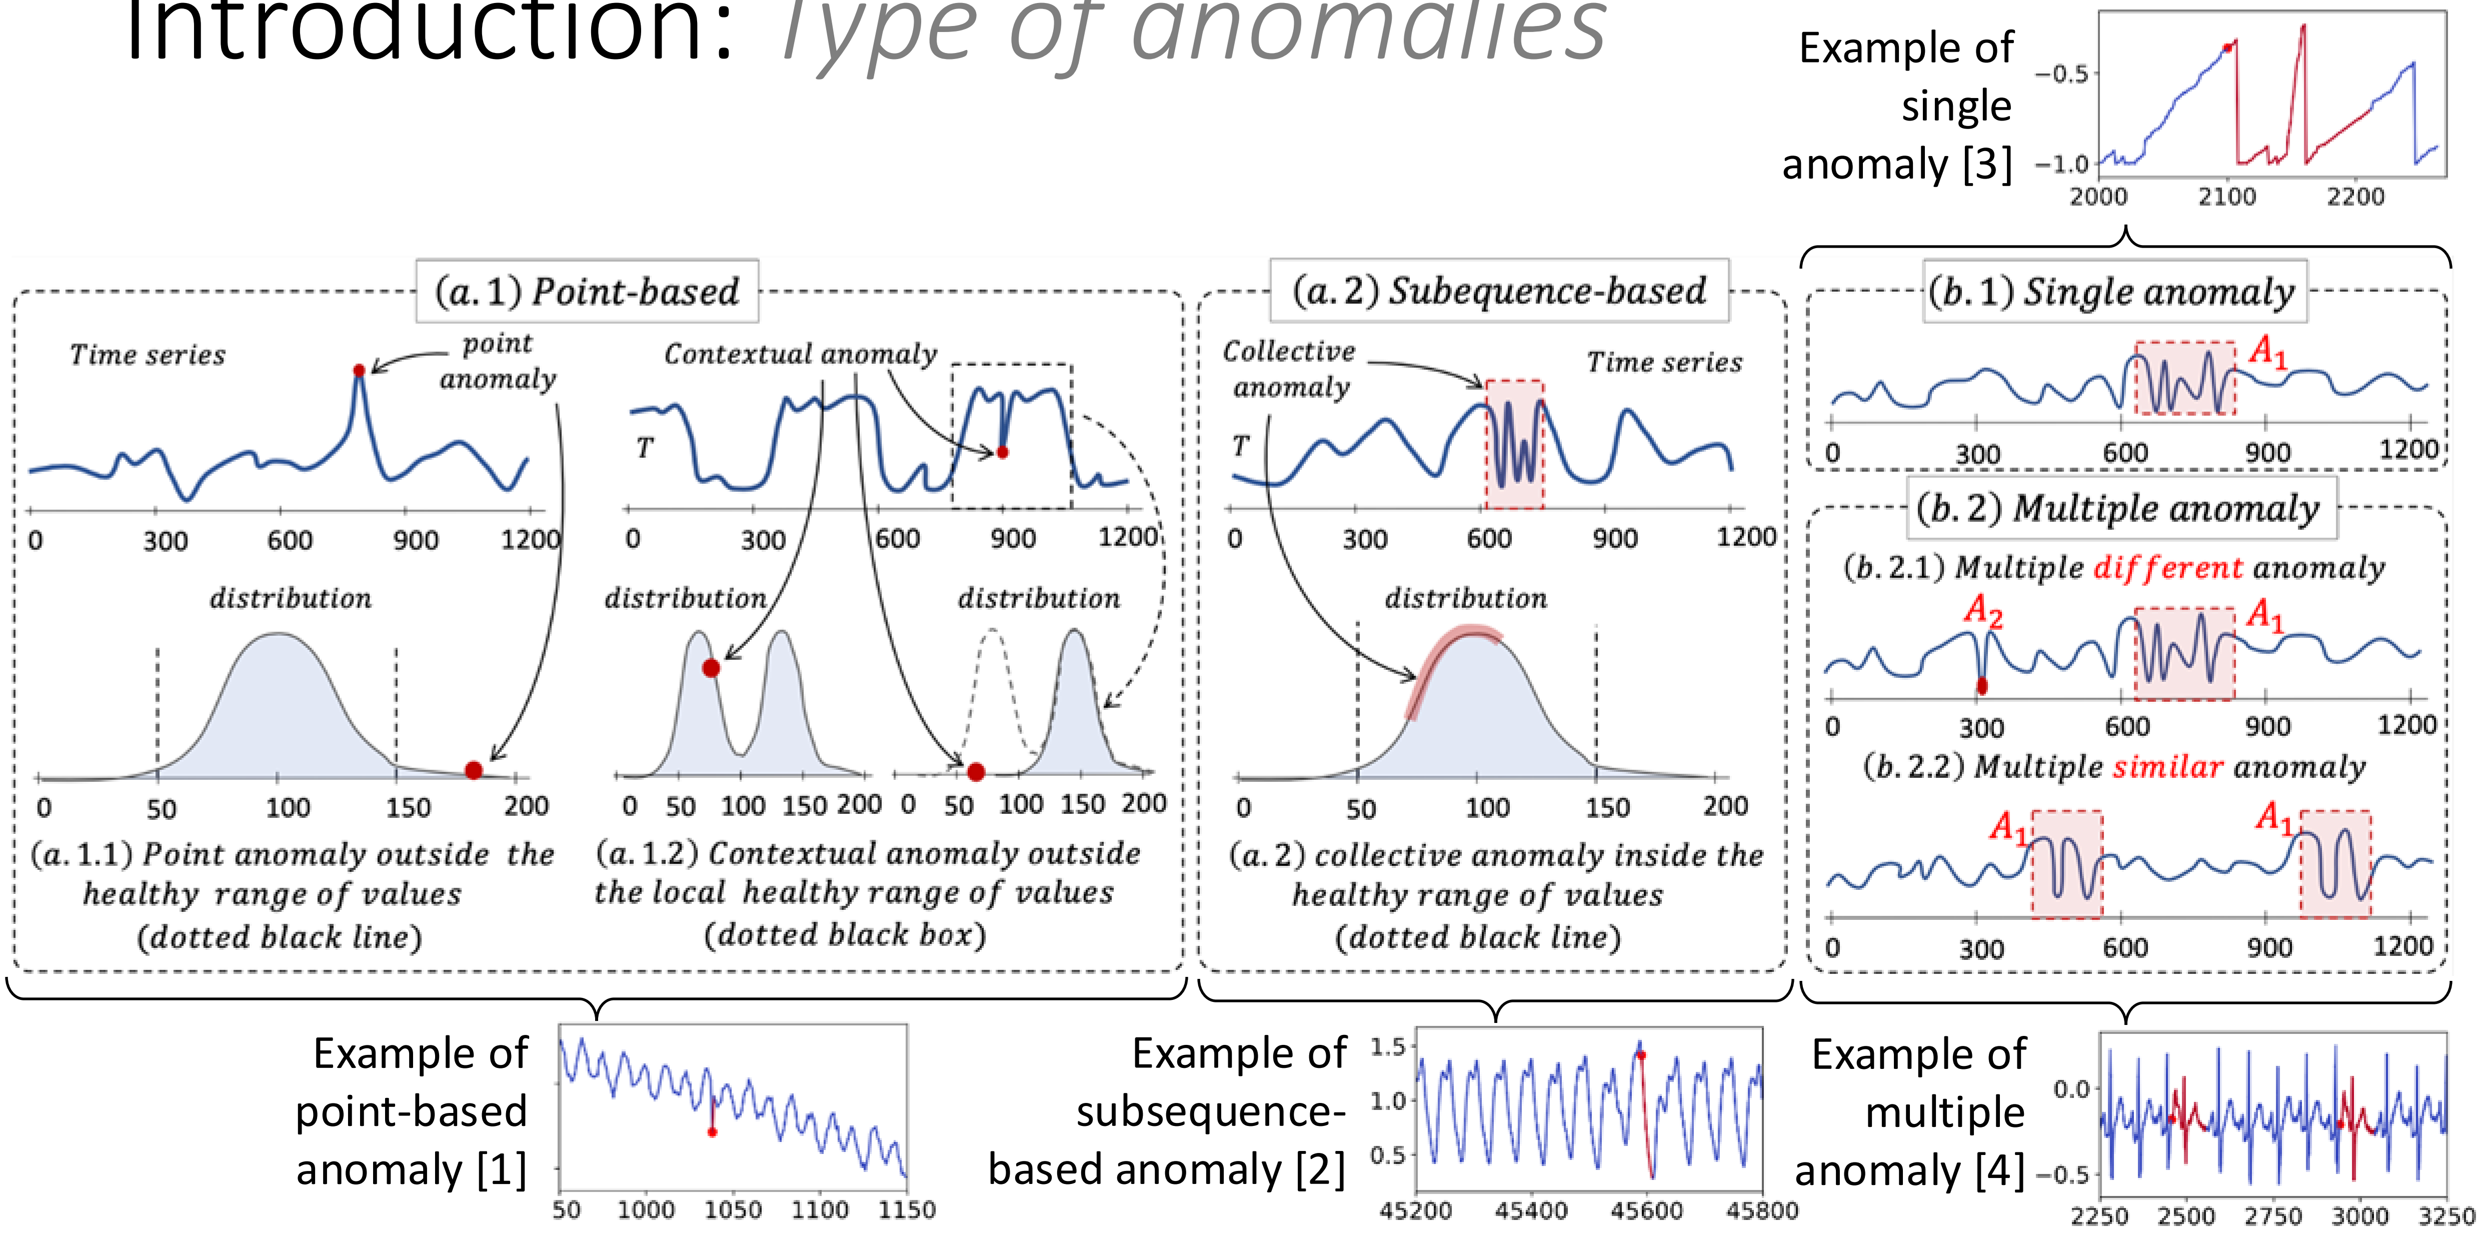

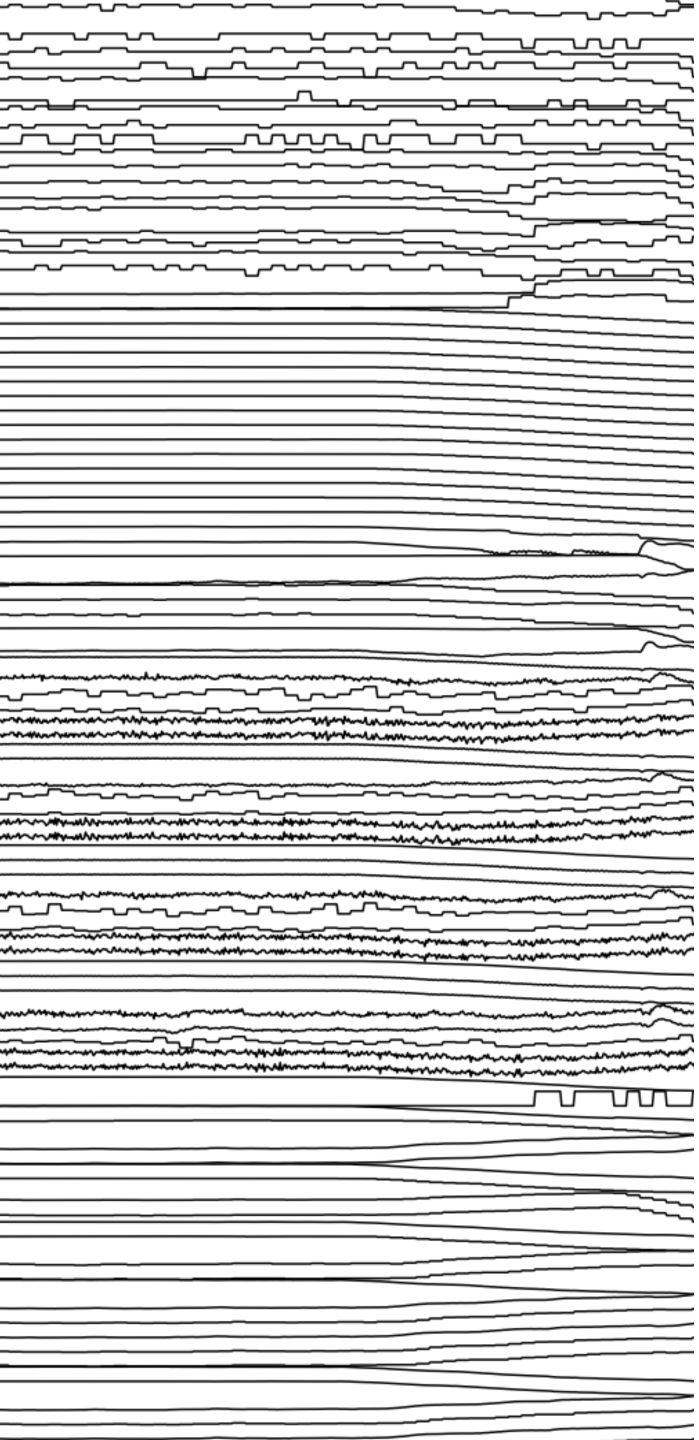

II. Time Series Anomaly Detection

How does it work?

Anomaly Detection methods: *A taxonomy*

Time Series

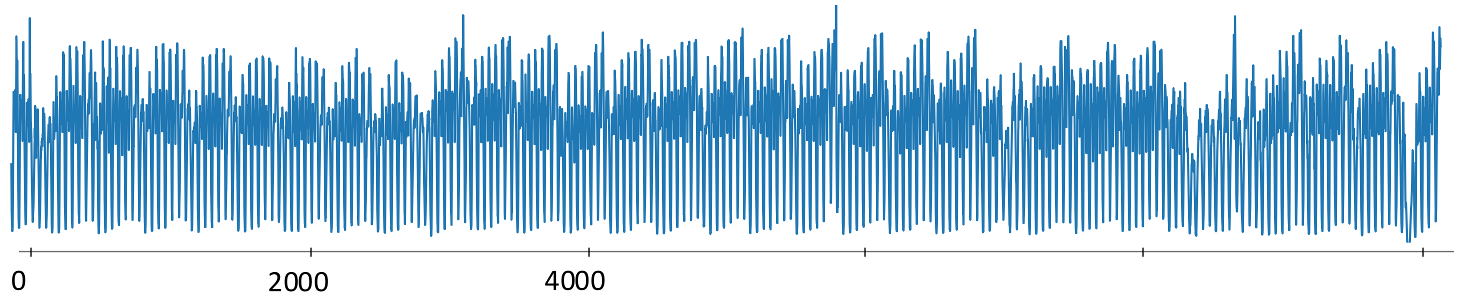

Anomaly Detection methods: *A taxonomy*

Time Series

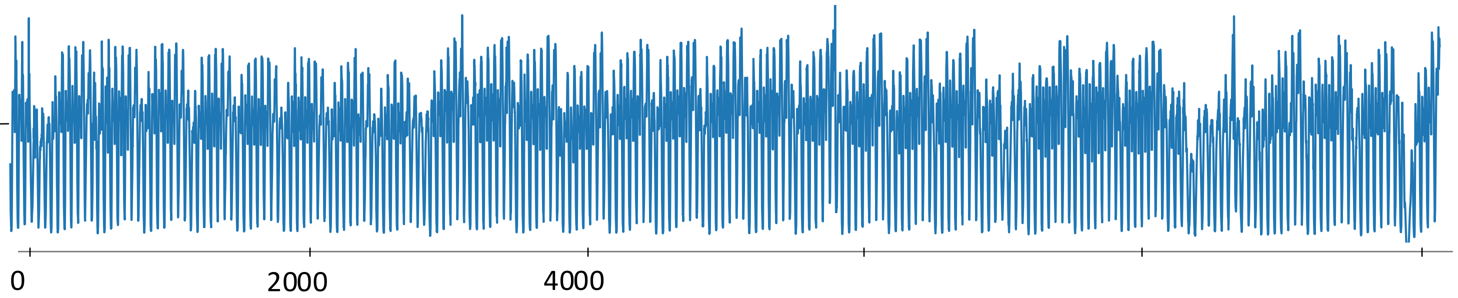

Pre-Processing

Anomaly Detection methods: *A taxonomy*

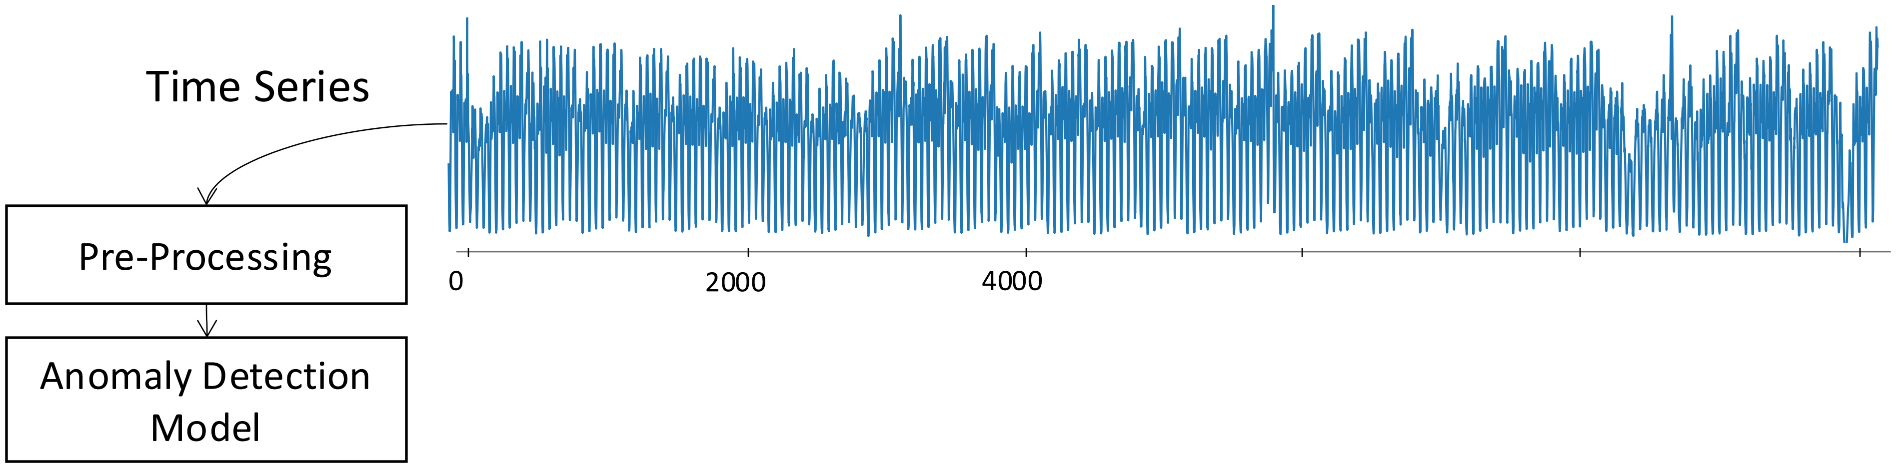

Anomaly Detection methods: *A taxonomy*

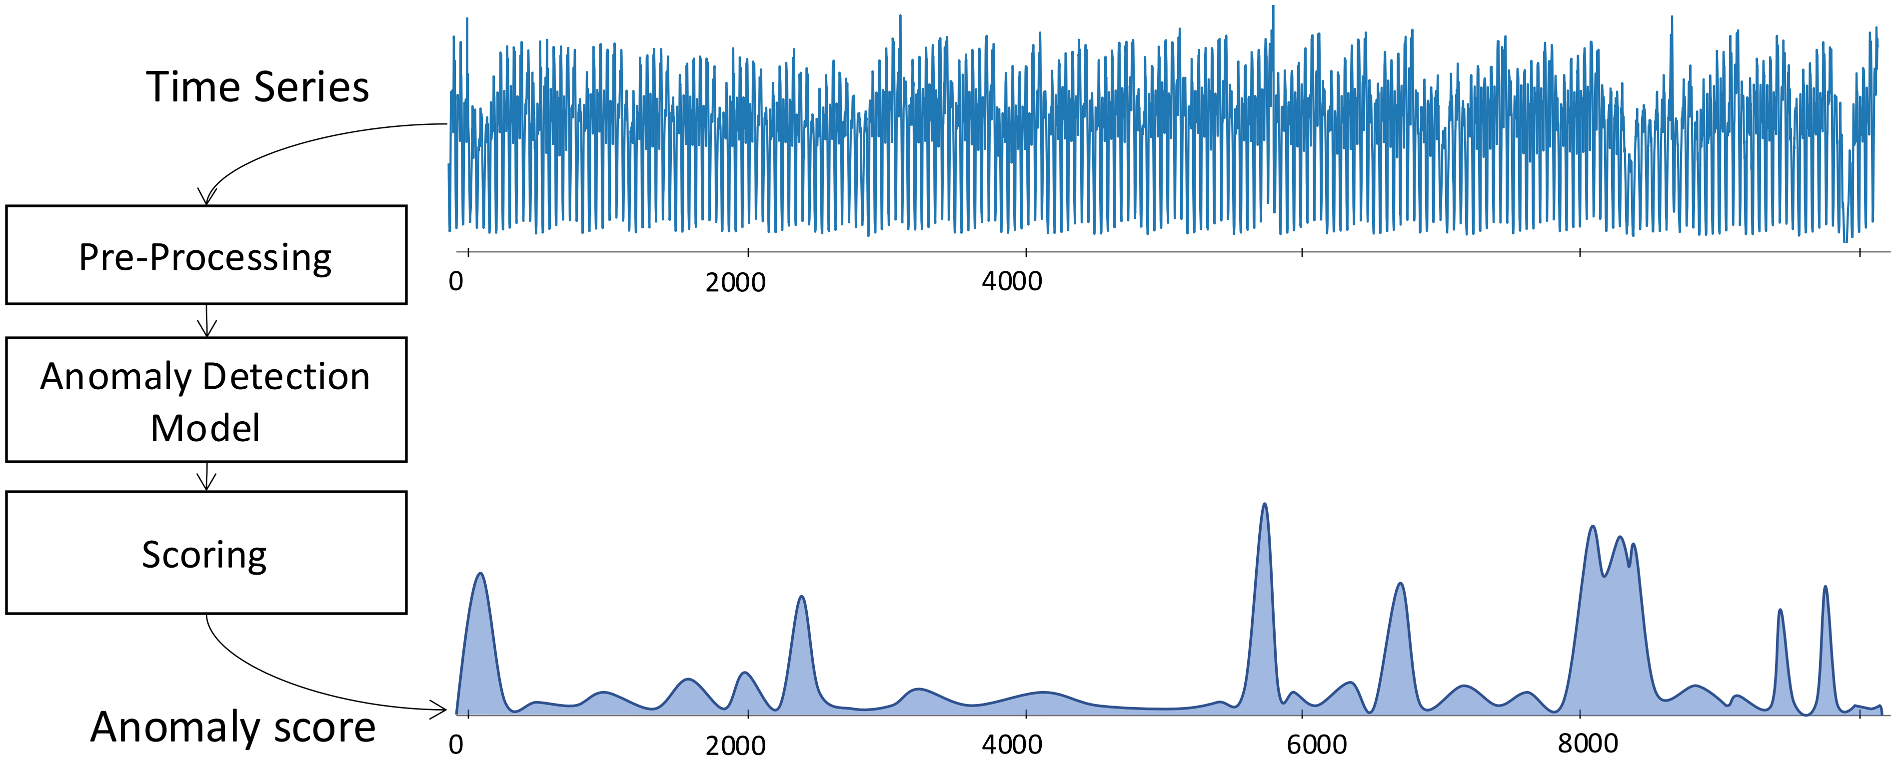

Anomaly Detection methods: *A taxonomy*

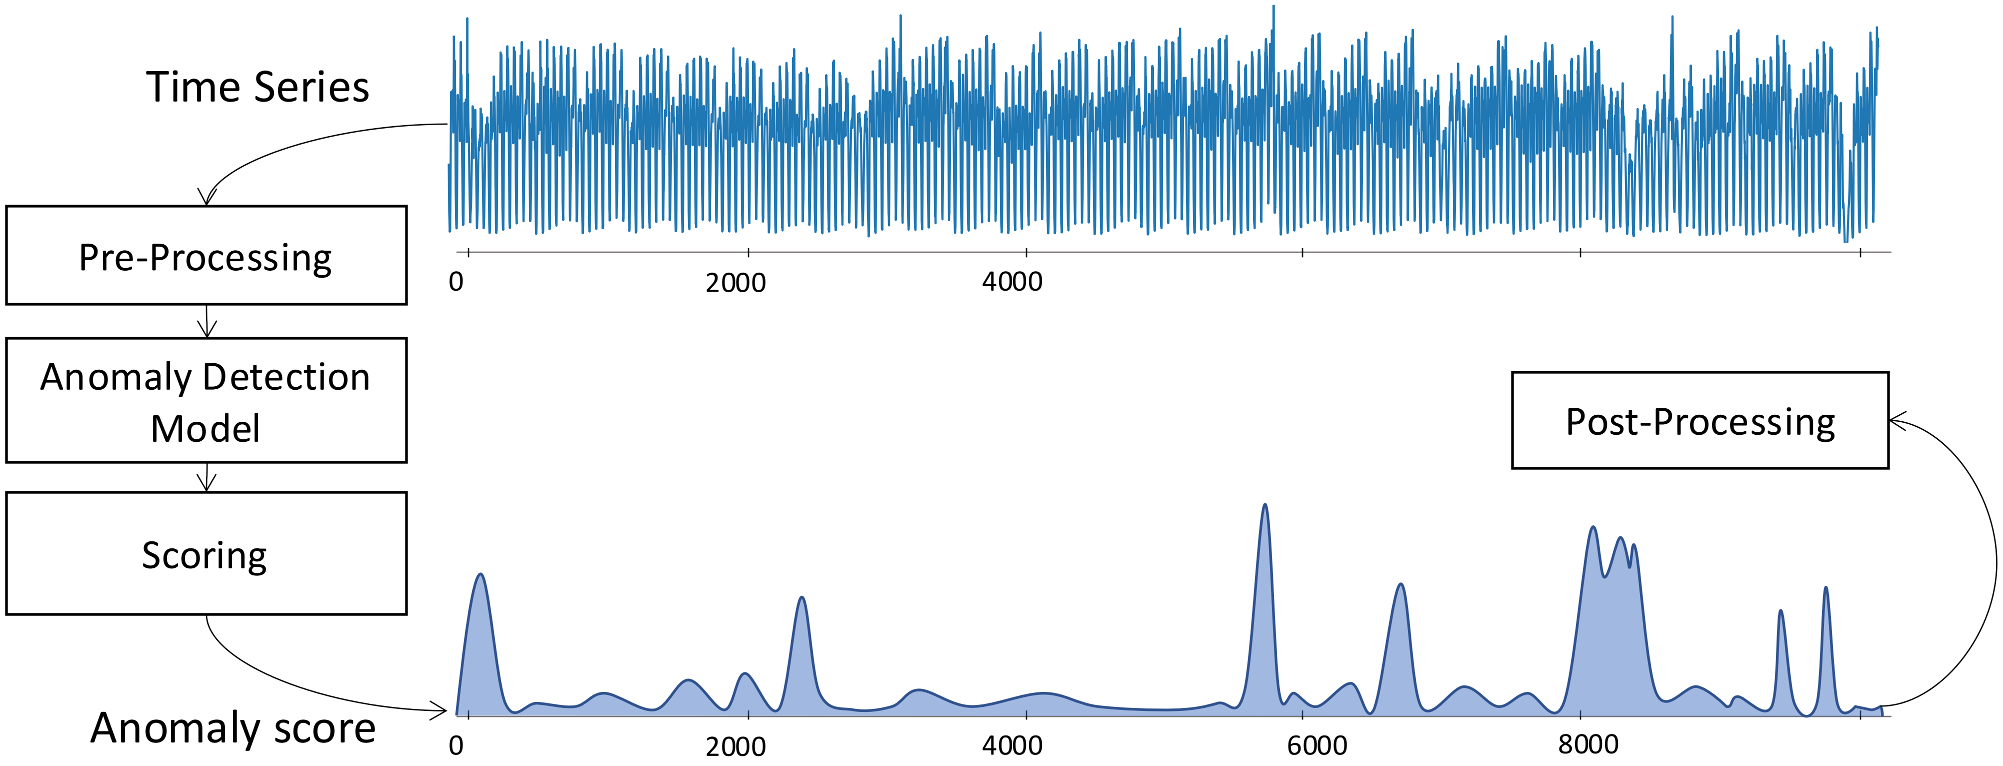

Anomaly Detection methods: *A taxonomy*

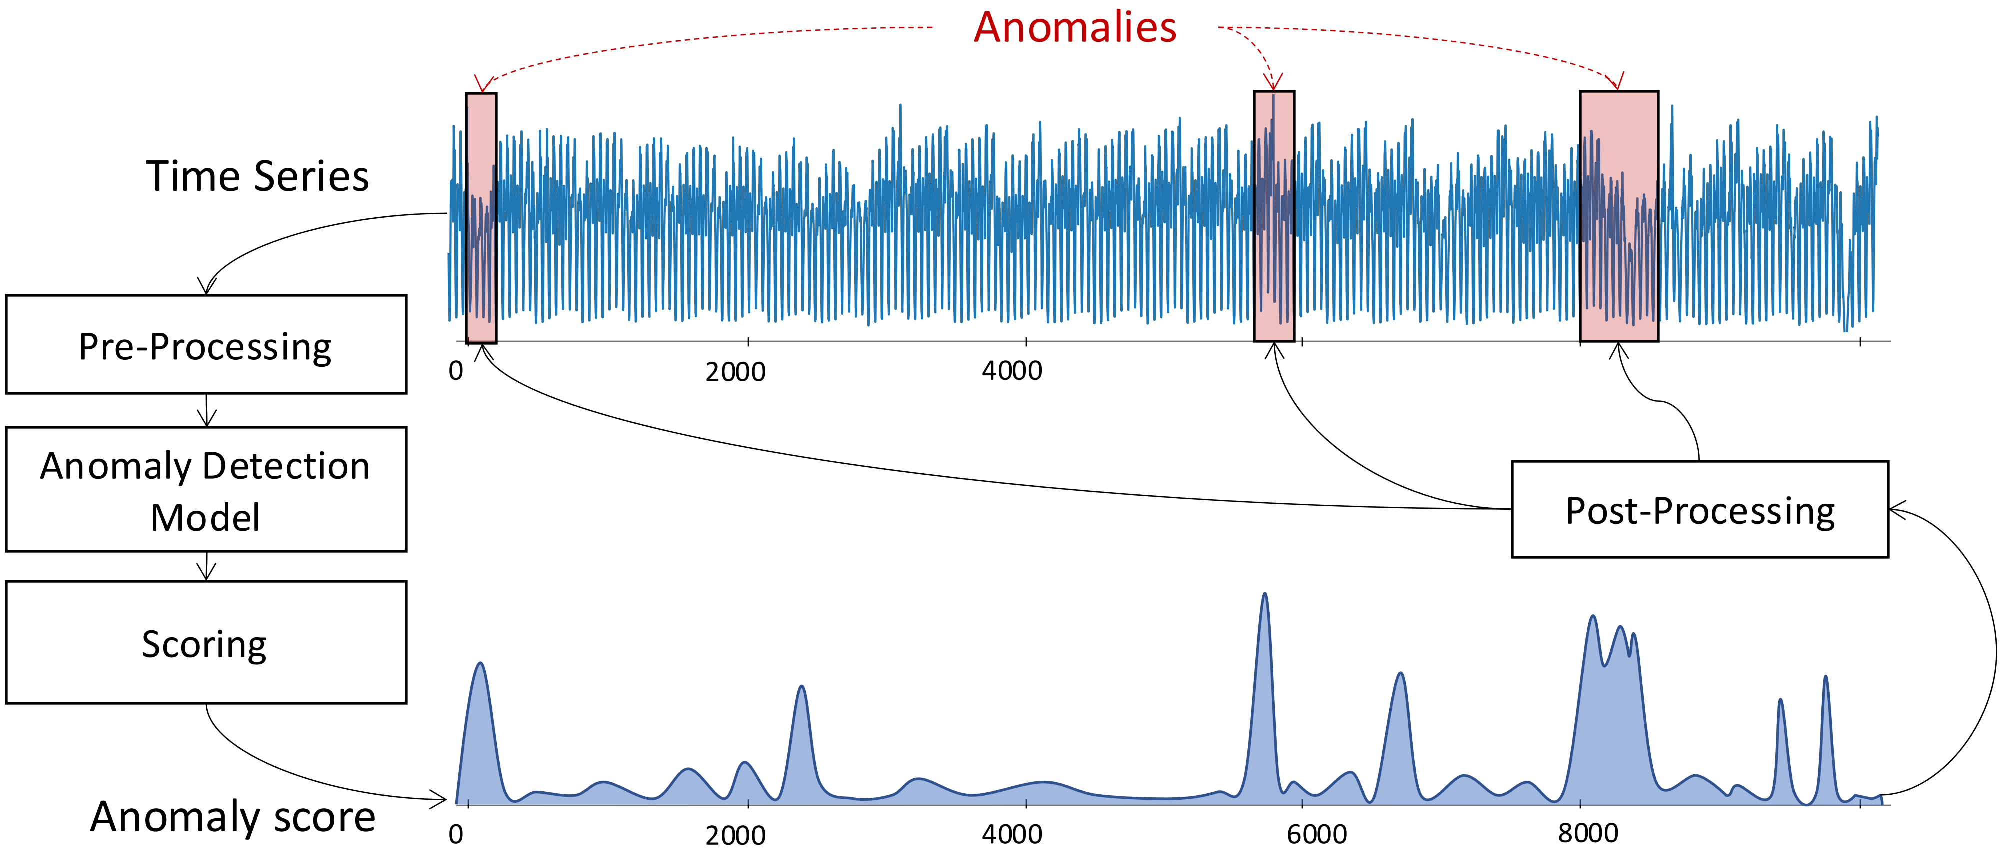

Anomaly Detection methods: *A taxonomy*

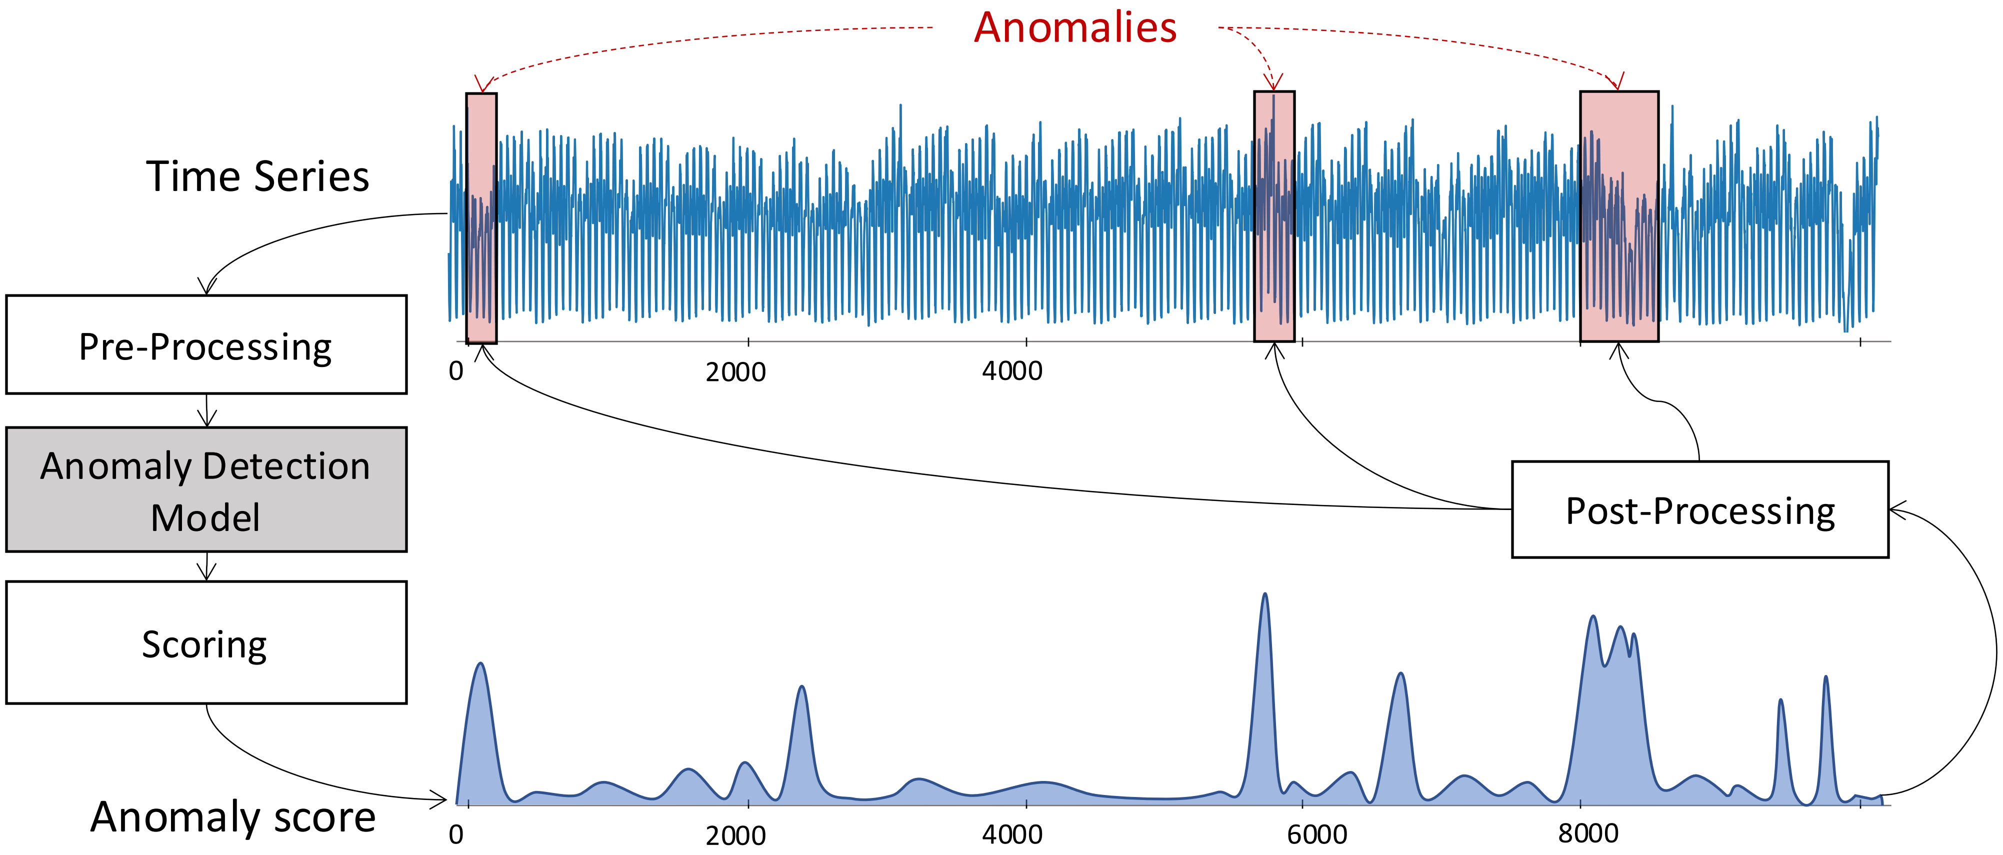

Anomaly Detection methods: *A taxonomy*

By domains [5] …

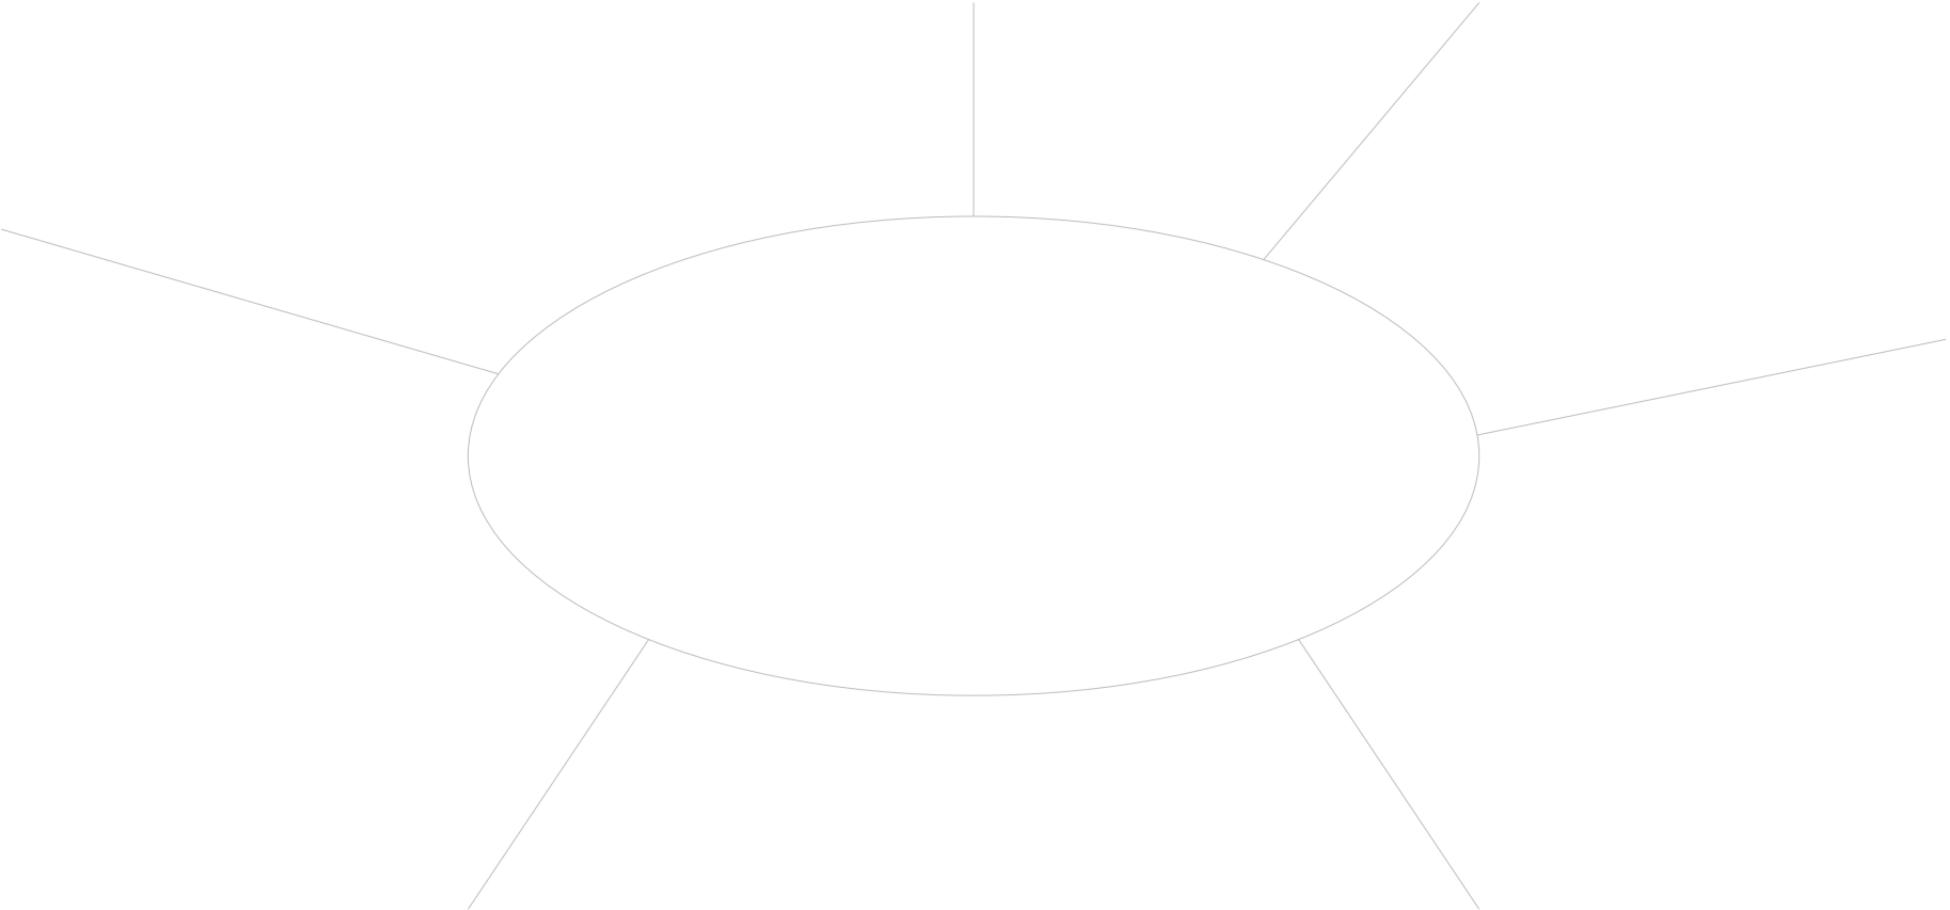

[5] Sebastian Schmidl, Phillip Wenig, and Thorsten Papenbrock. 2022. Anomaly detection in time series: a comprehensive evaluation. Proc. VLDB Endow. 15, 9 (May 2022), 1779–1797.

Anomaly Detection methods: *A taxonomy*

By domains [5] …

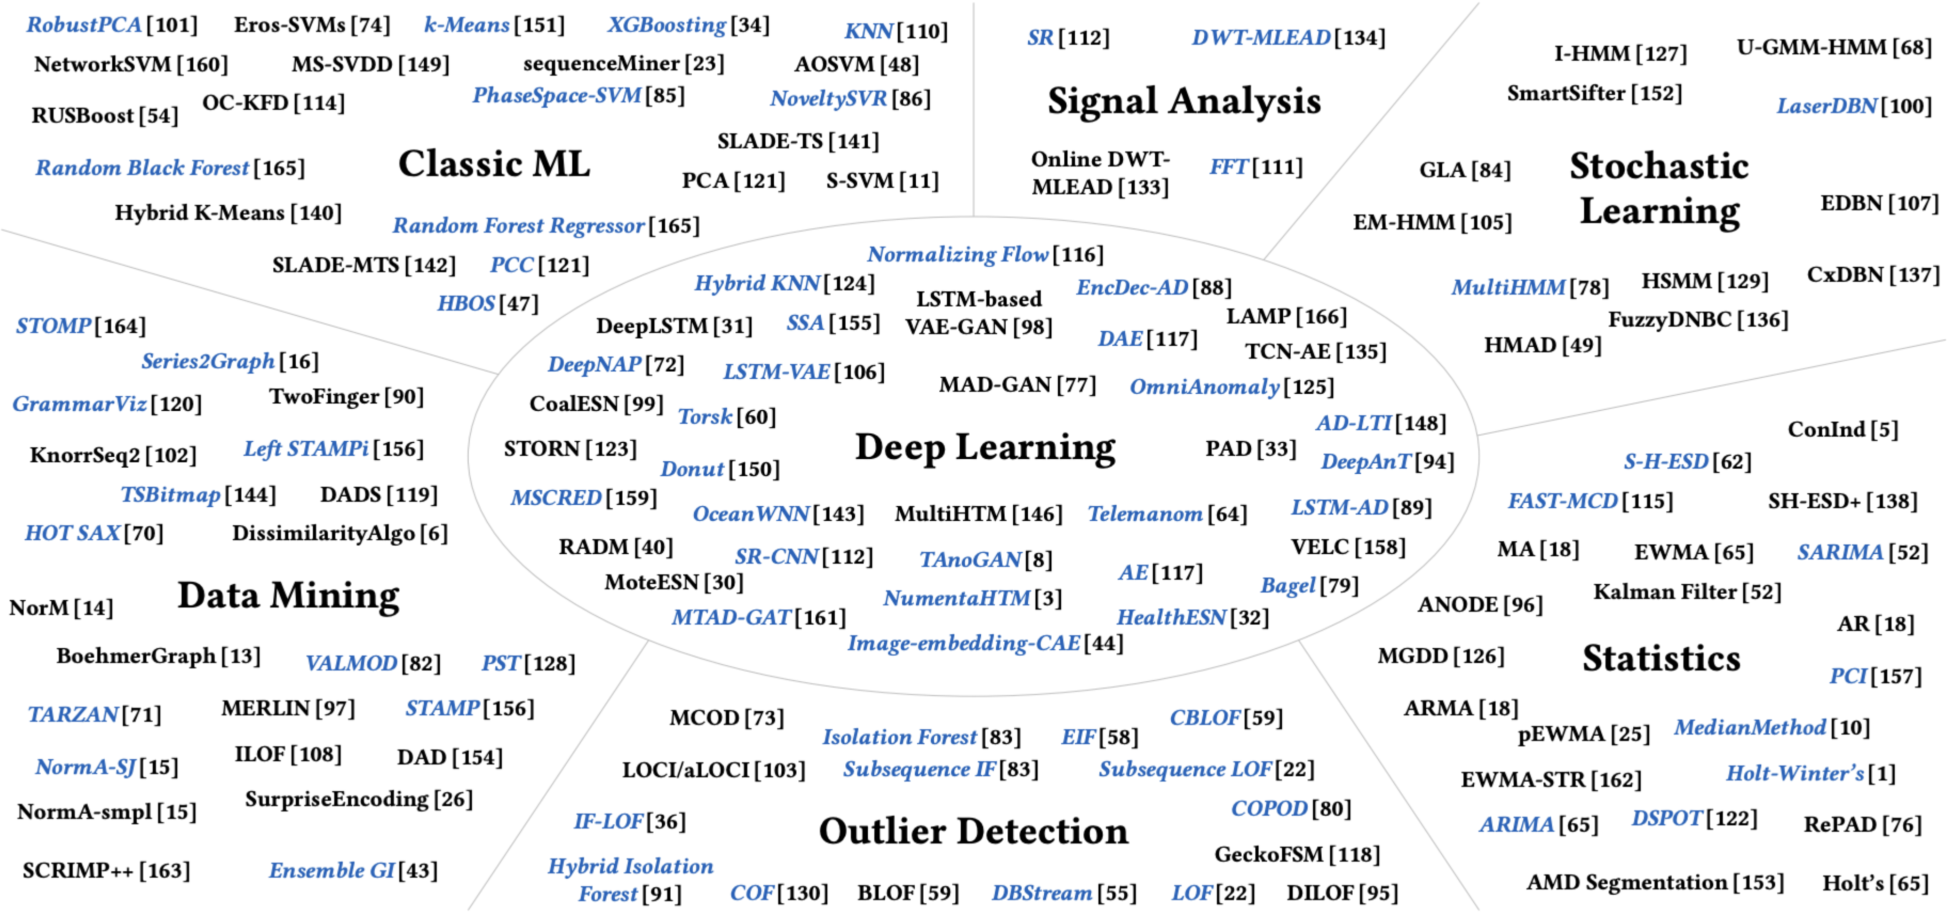

[5] Sebastian Schmidl, Phillip Wenig, and Thorsten Papenbrock. 2022. Anomaly detection in time series: a comprehensive evaluation. Proc. VLDB Endow. 15, 9 (May 2022), 1779–1797.

Anomaly Detection methods: *A taxonomy*

By inputs…

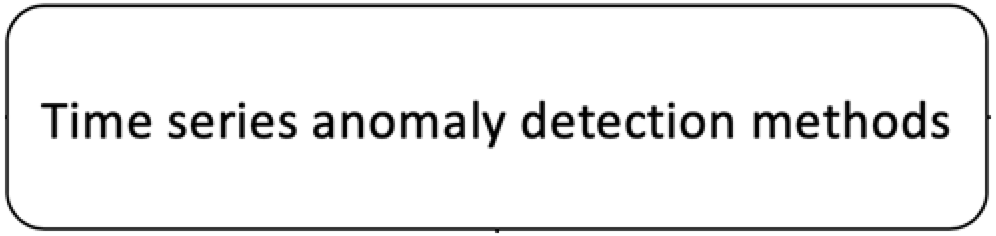

Anomaly Detection methods: *A taxonomy*

By inputs…

Time series anomaly detection methods

Supervised

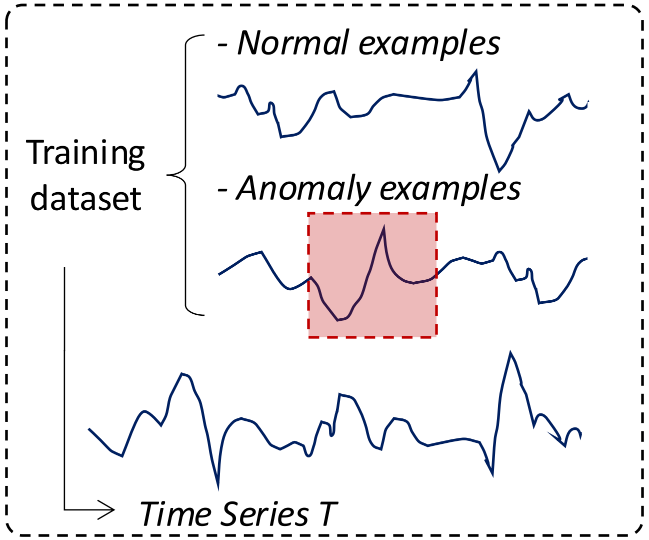

Anomaly Detection methods: *A taxonomy*

By inputs…

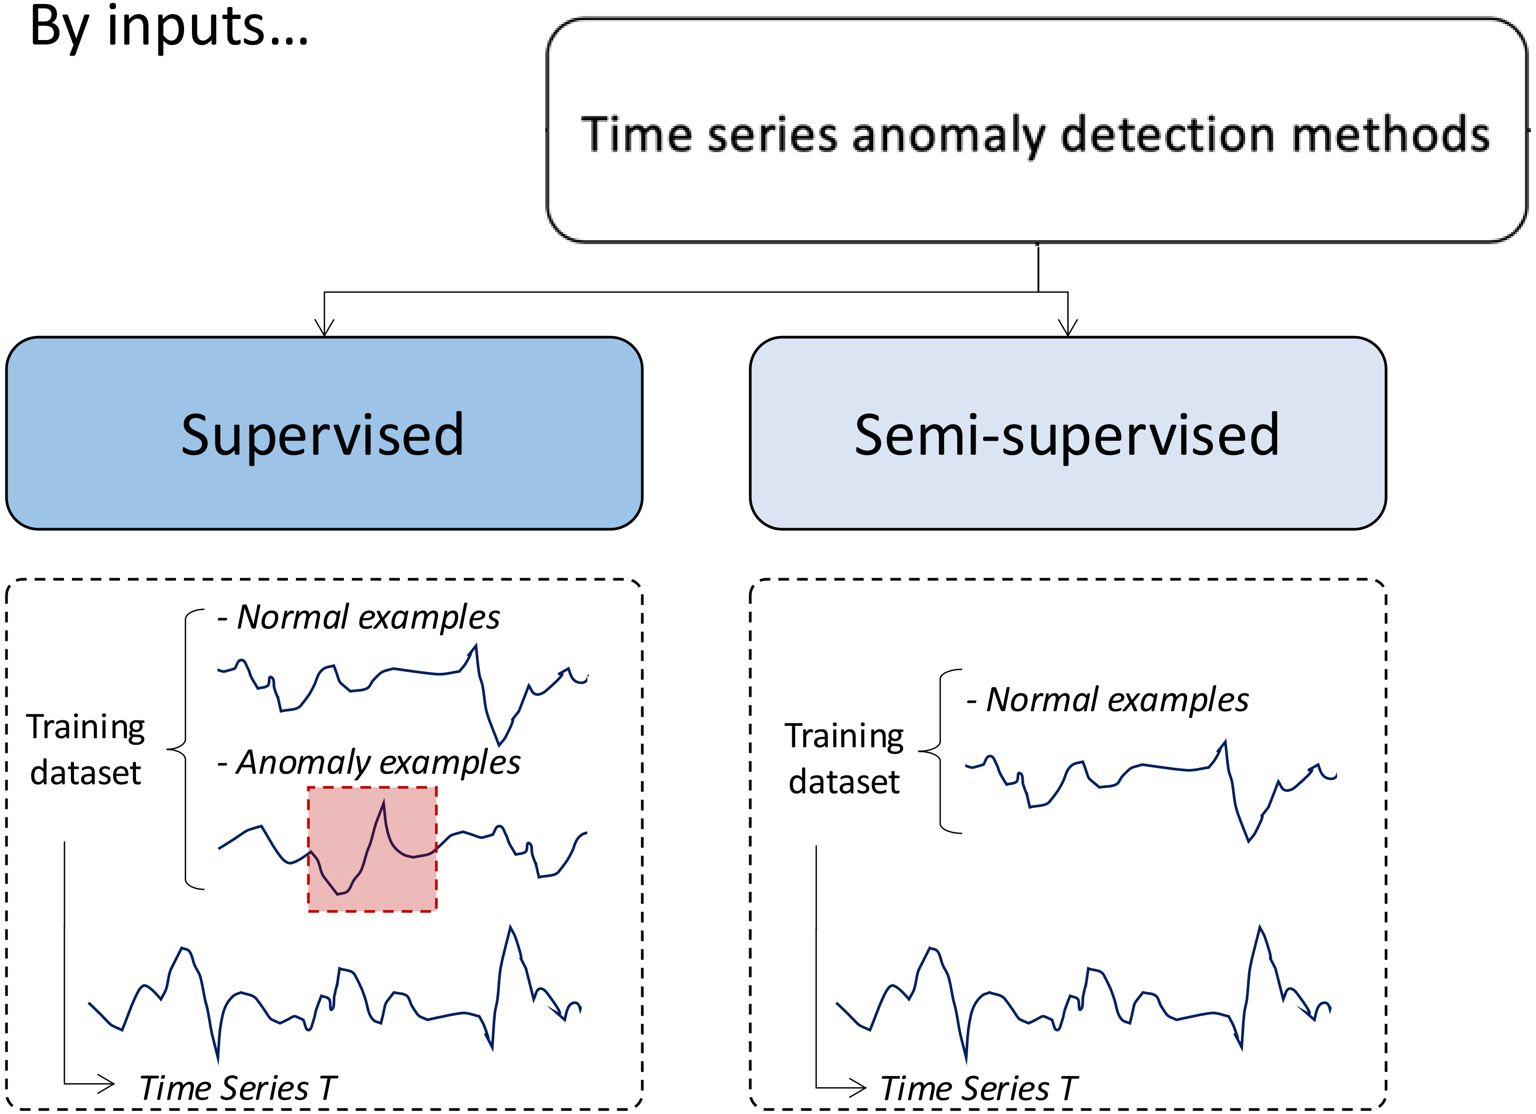

Anomaly Detection methods: *A taxonomy*

By inputs…

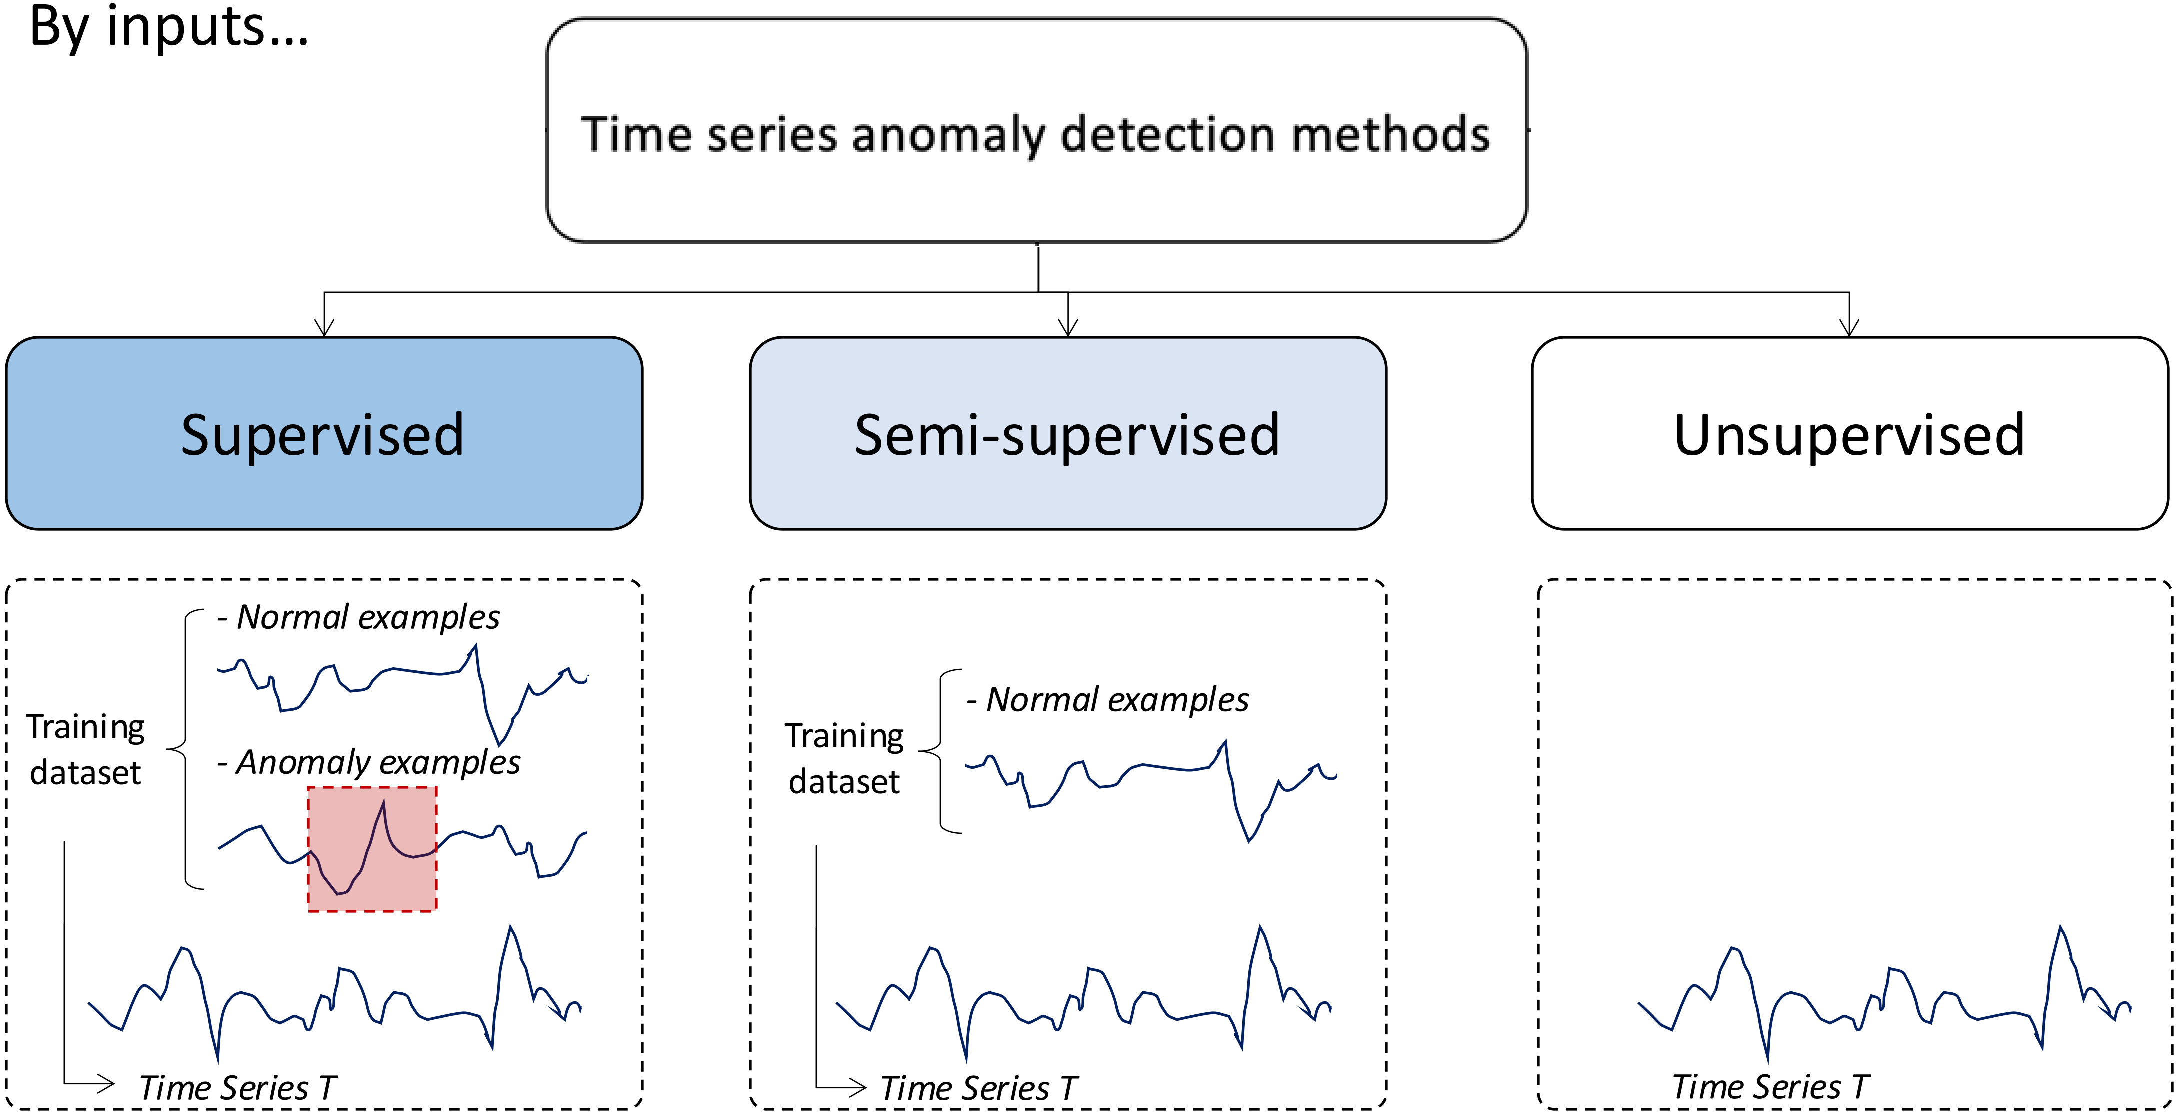

Anomaly Dete

By inputs…

Time

Supervised

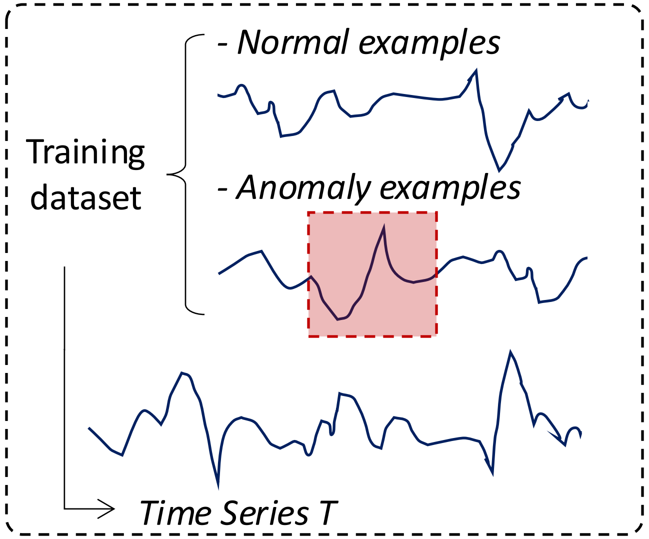

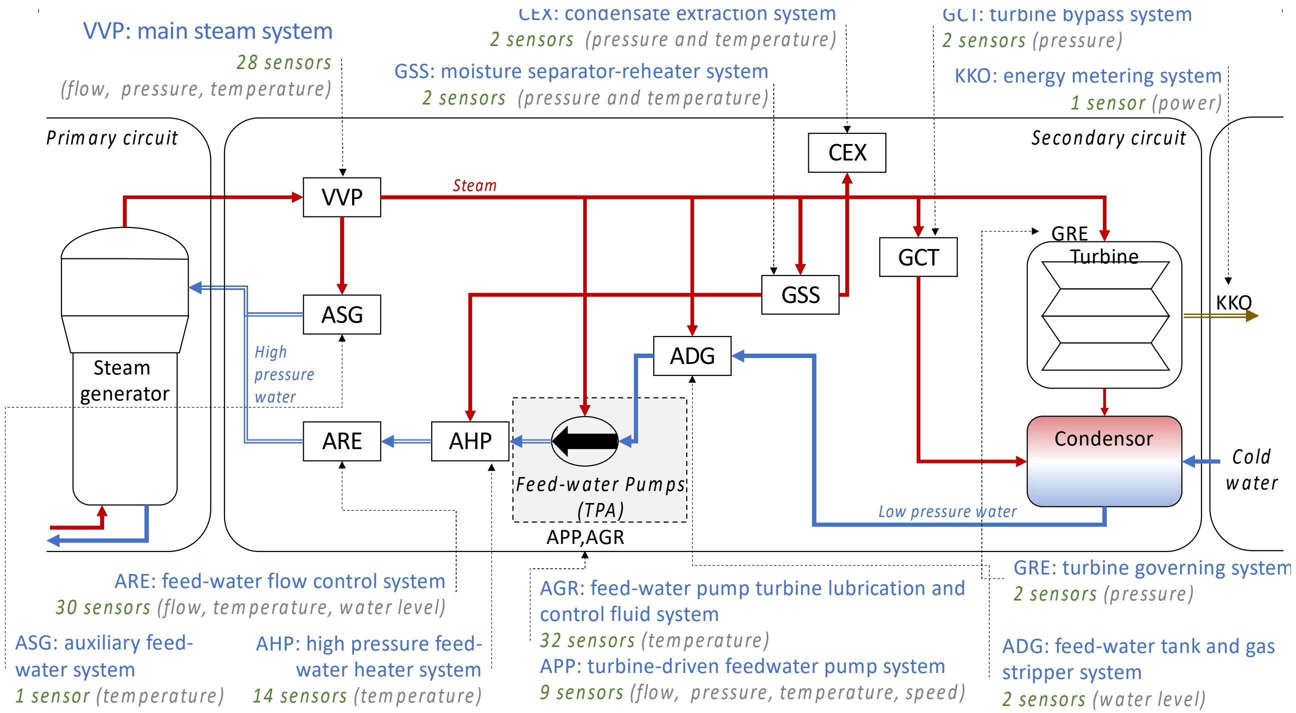

Anomaly Dete

By inputs…

Time

Supervised

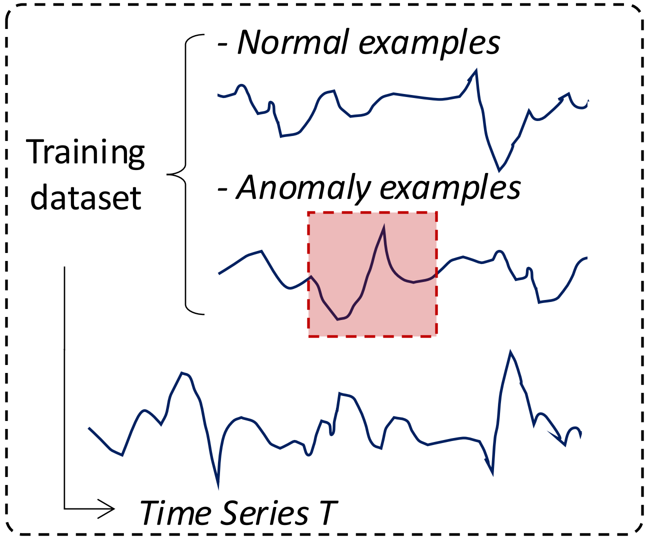

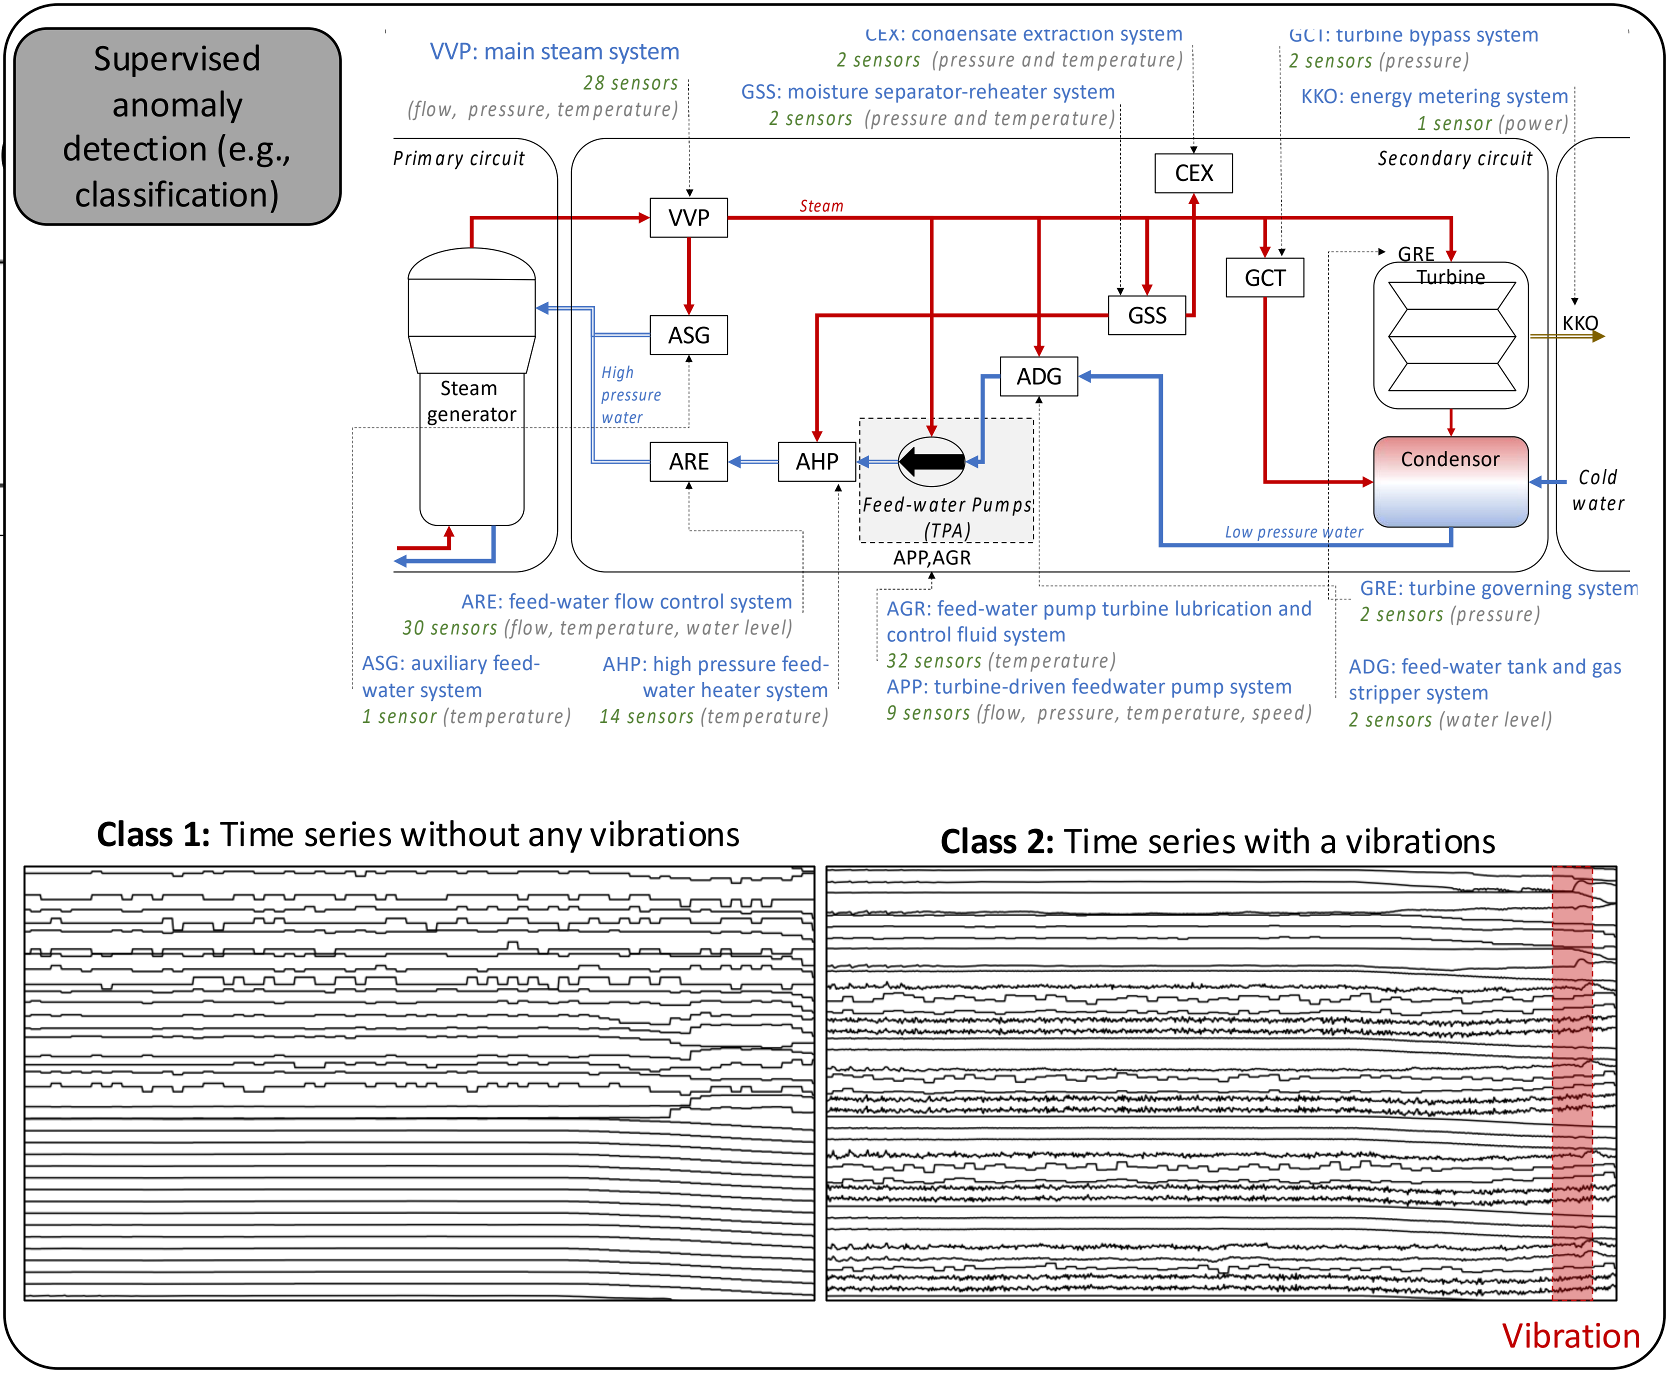

Anomaly Dete

By inputs…

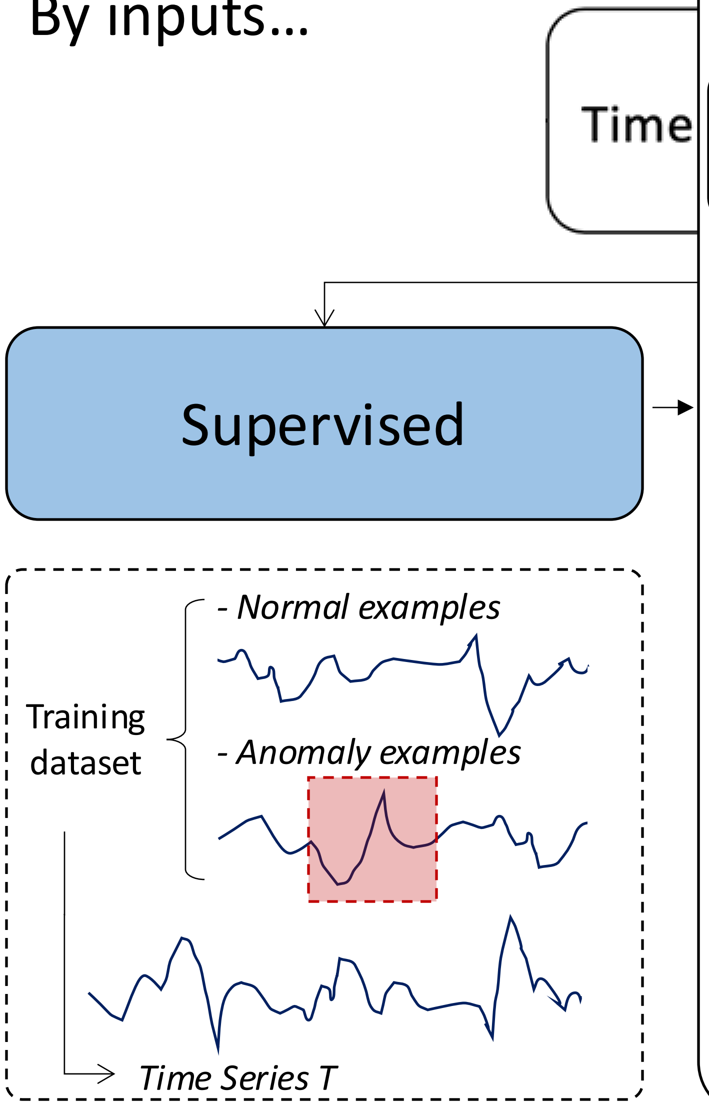

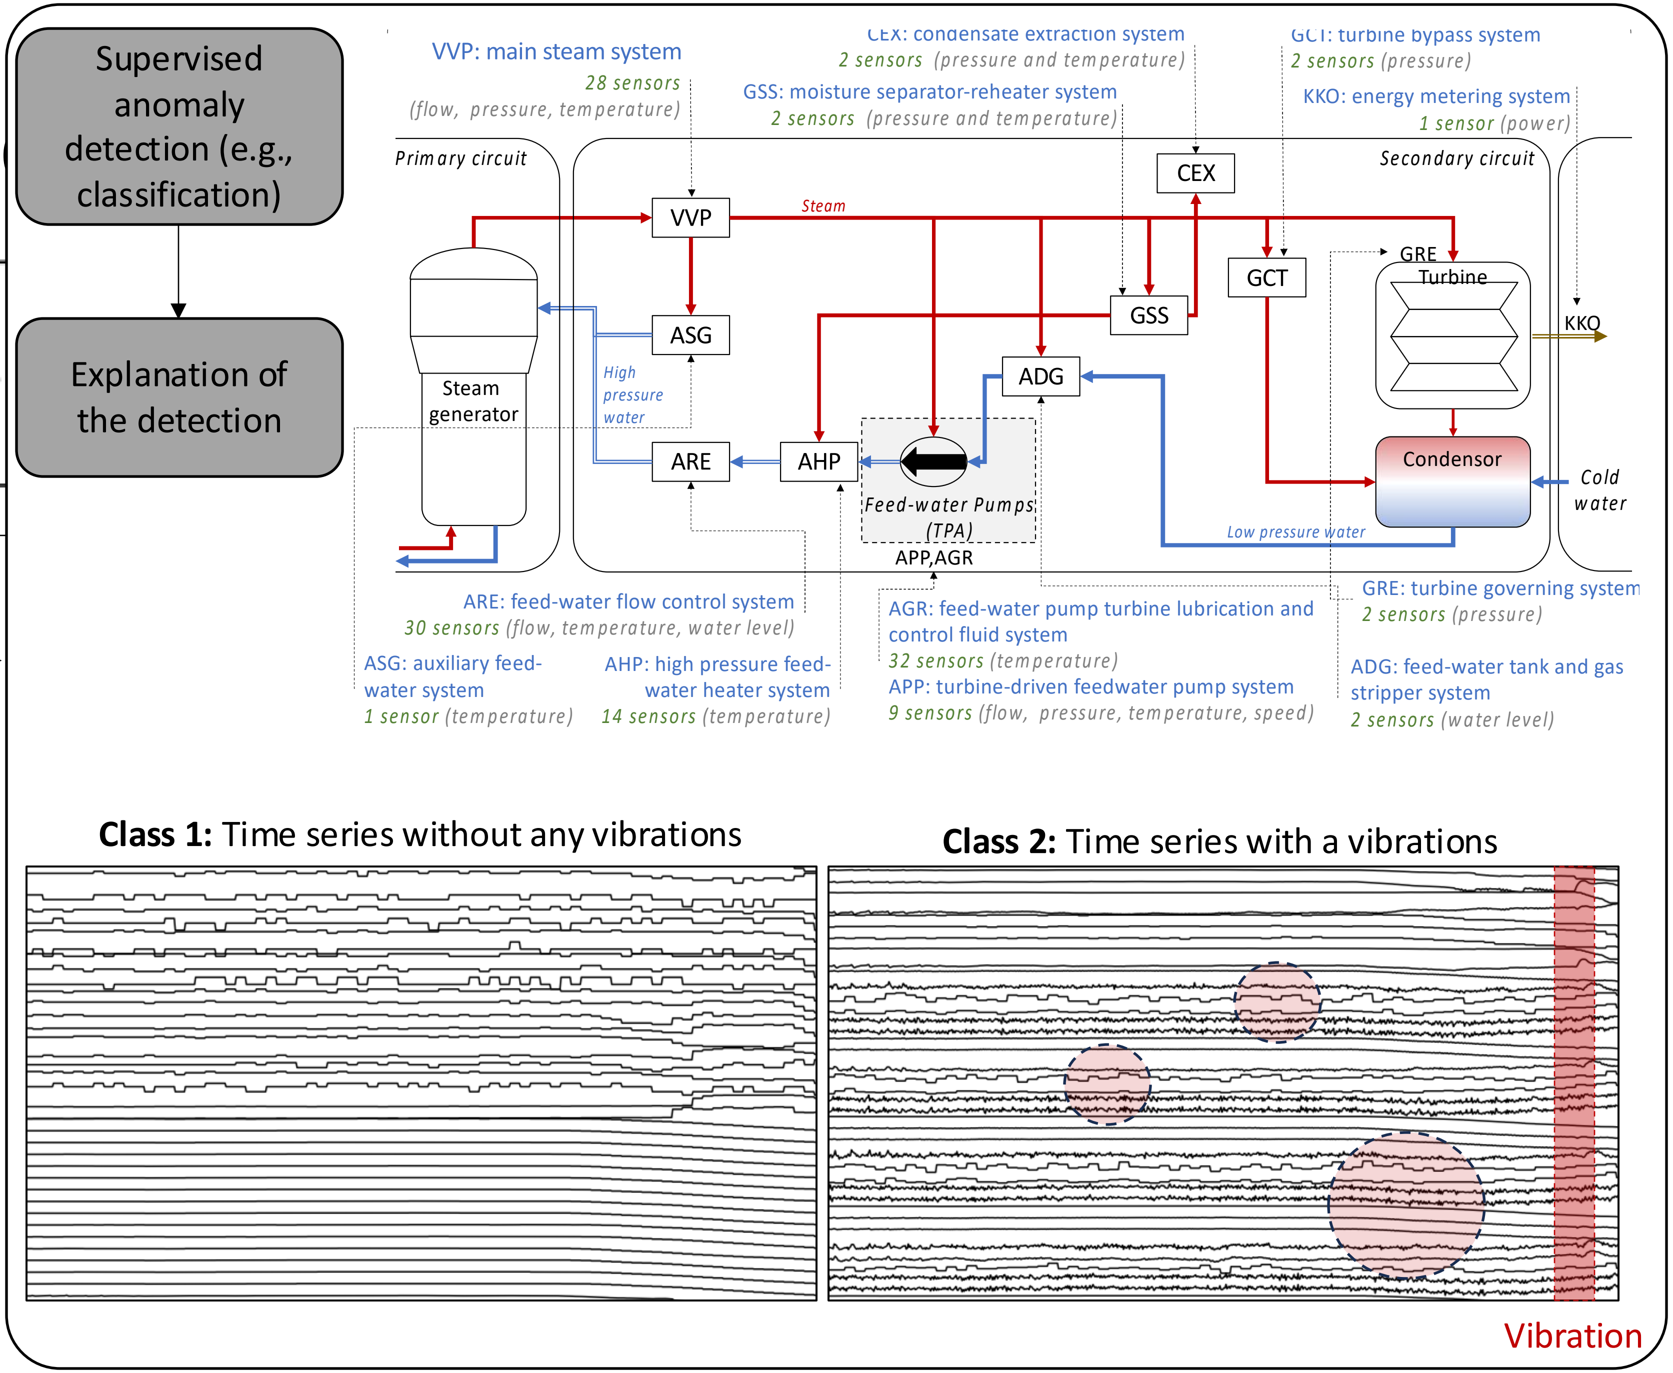

Anomaly Dete

By inputs…

Time

Supervised

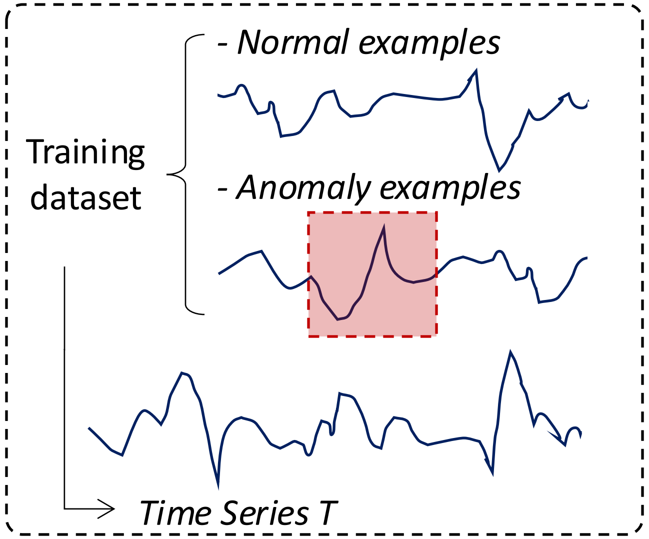

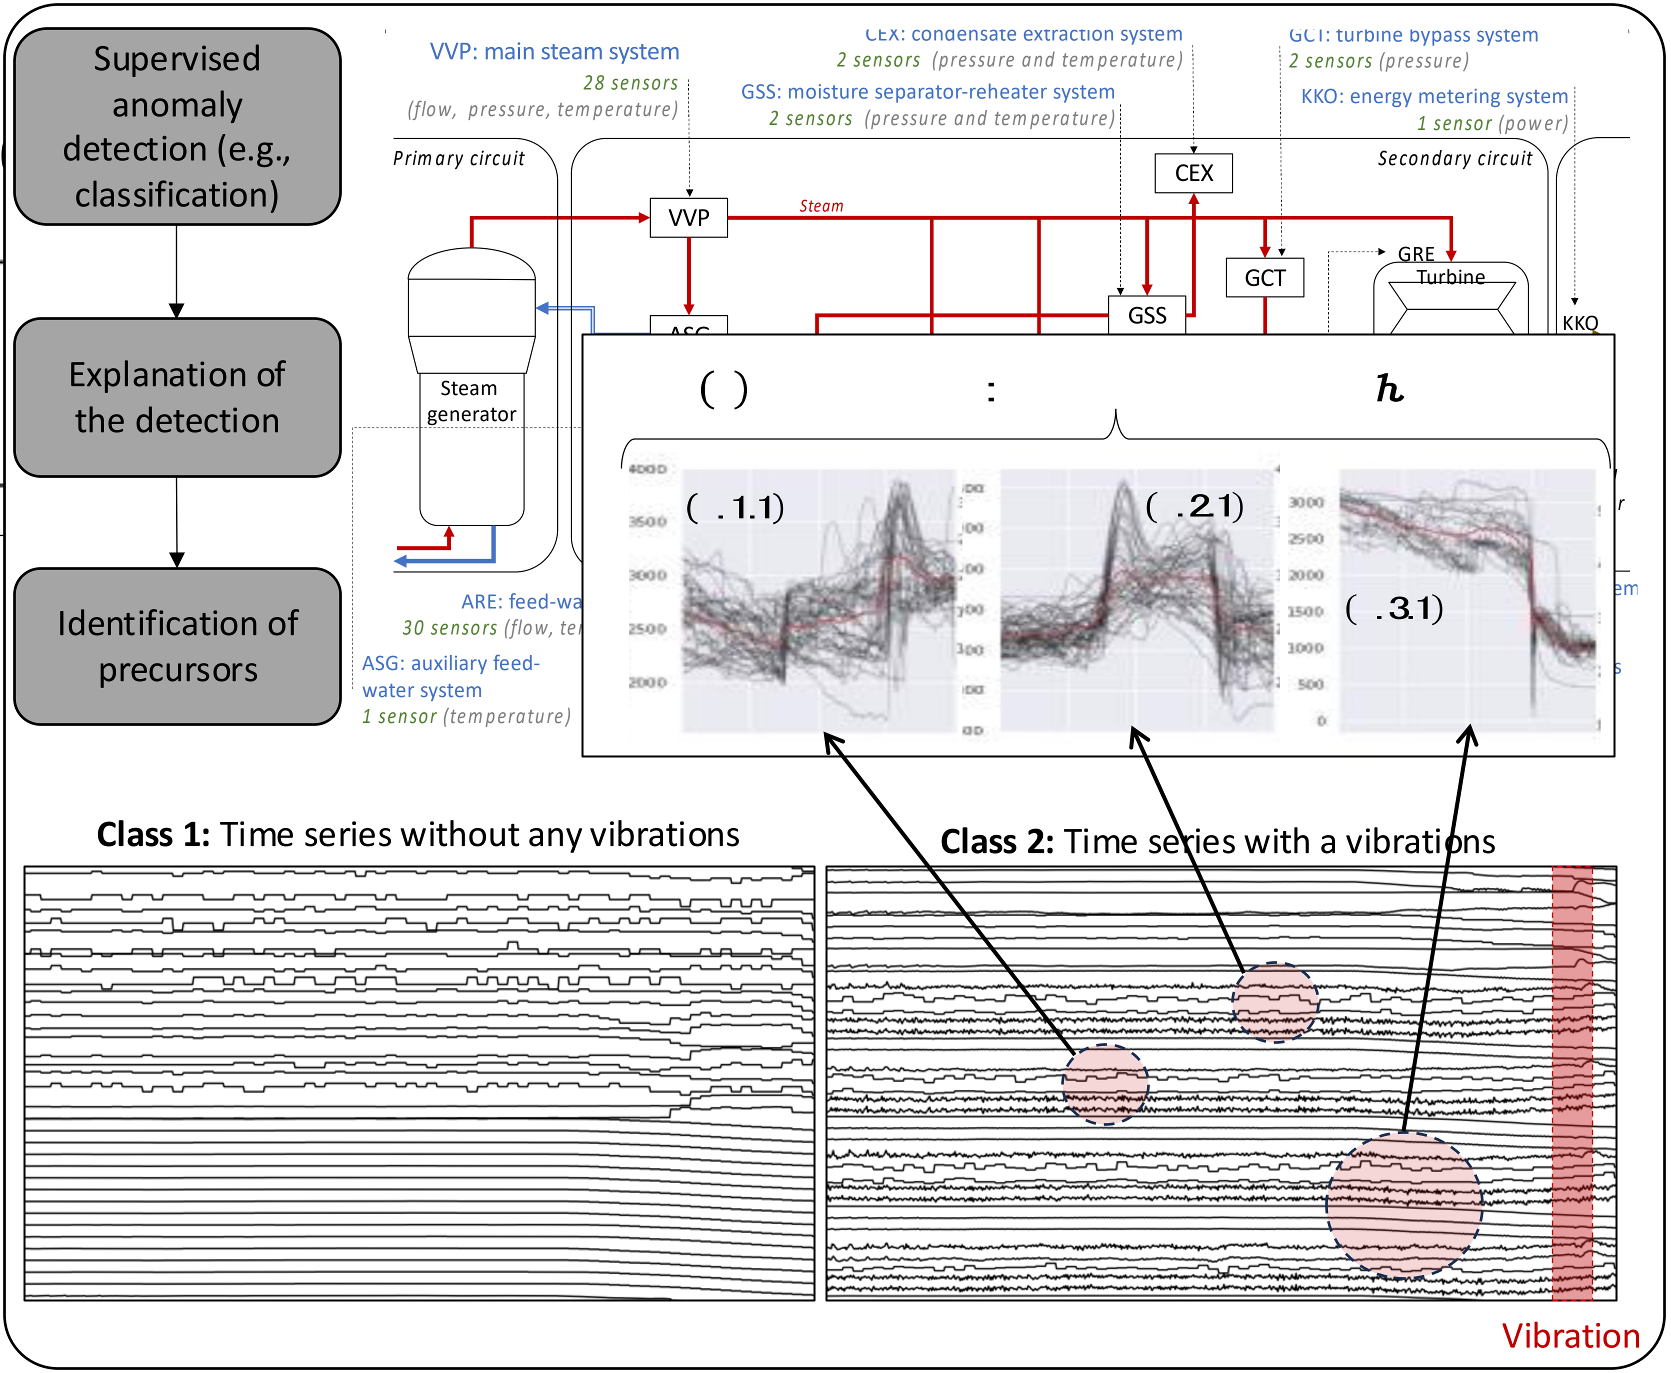

Anomaly Dete

By inputs…

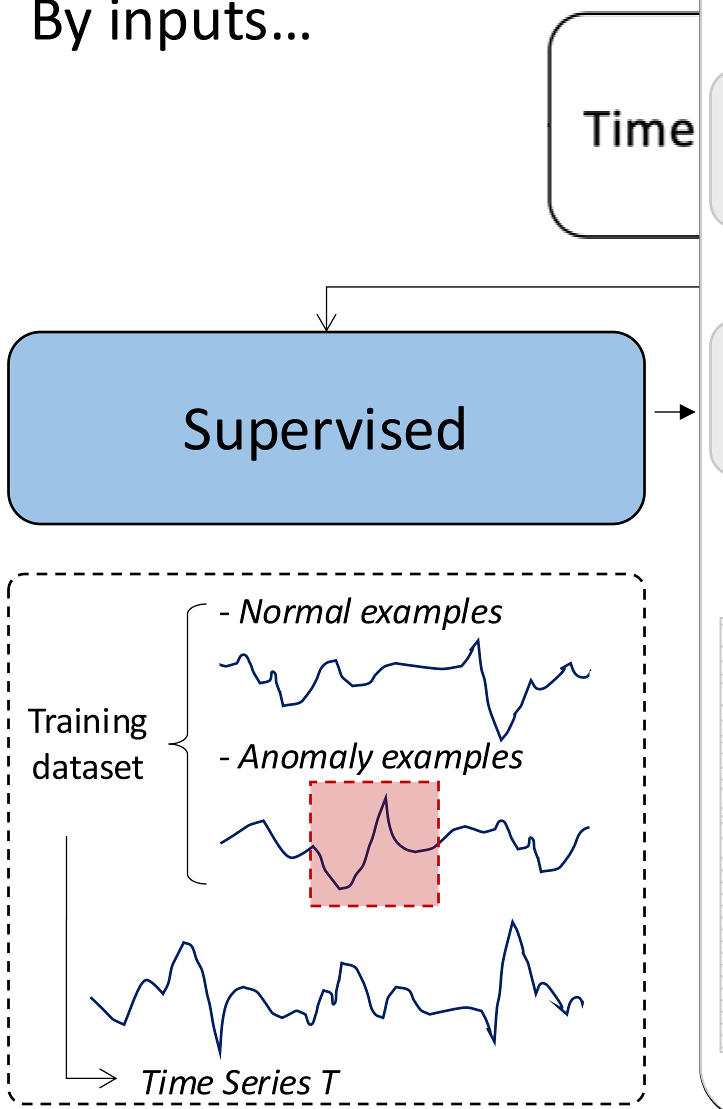

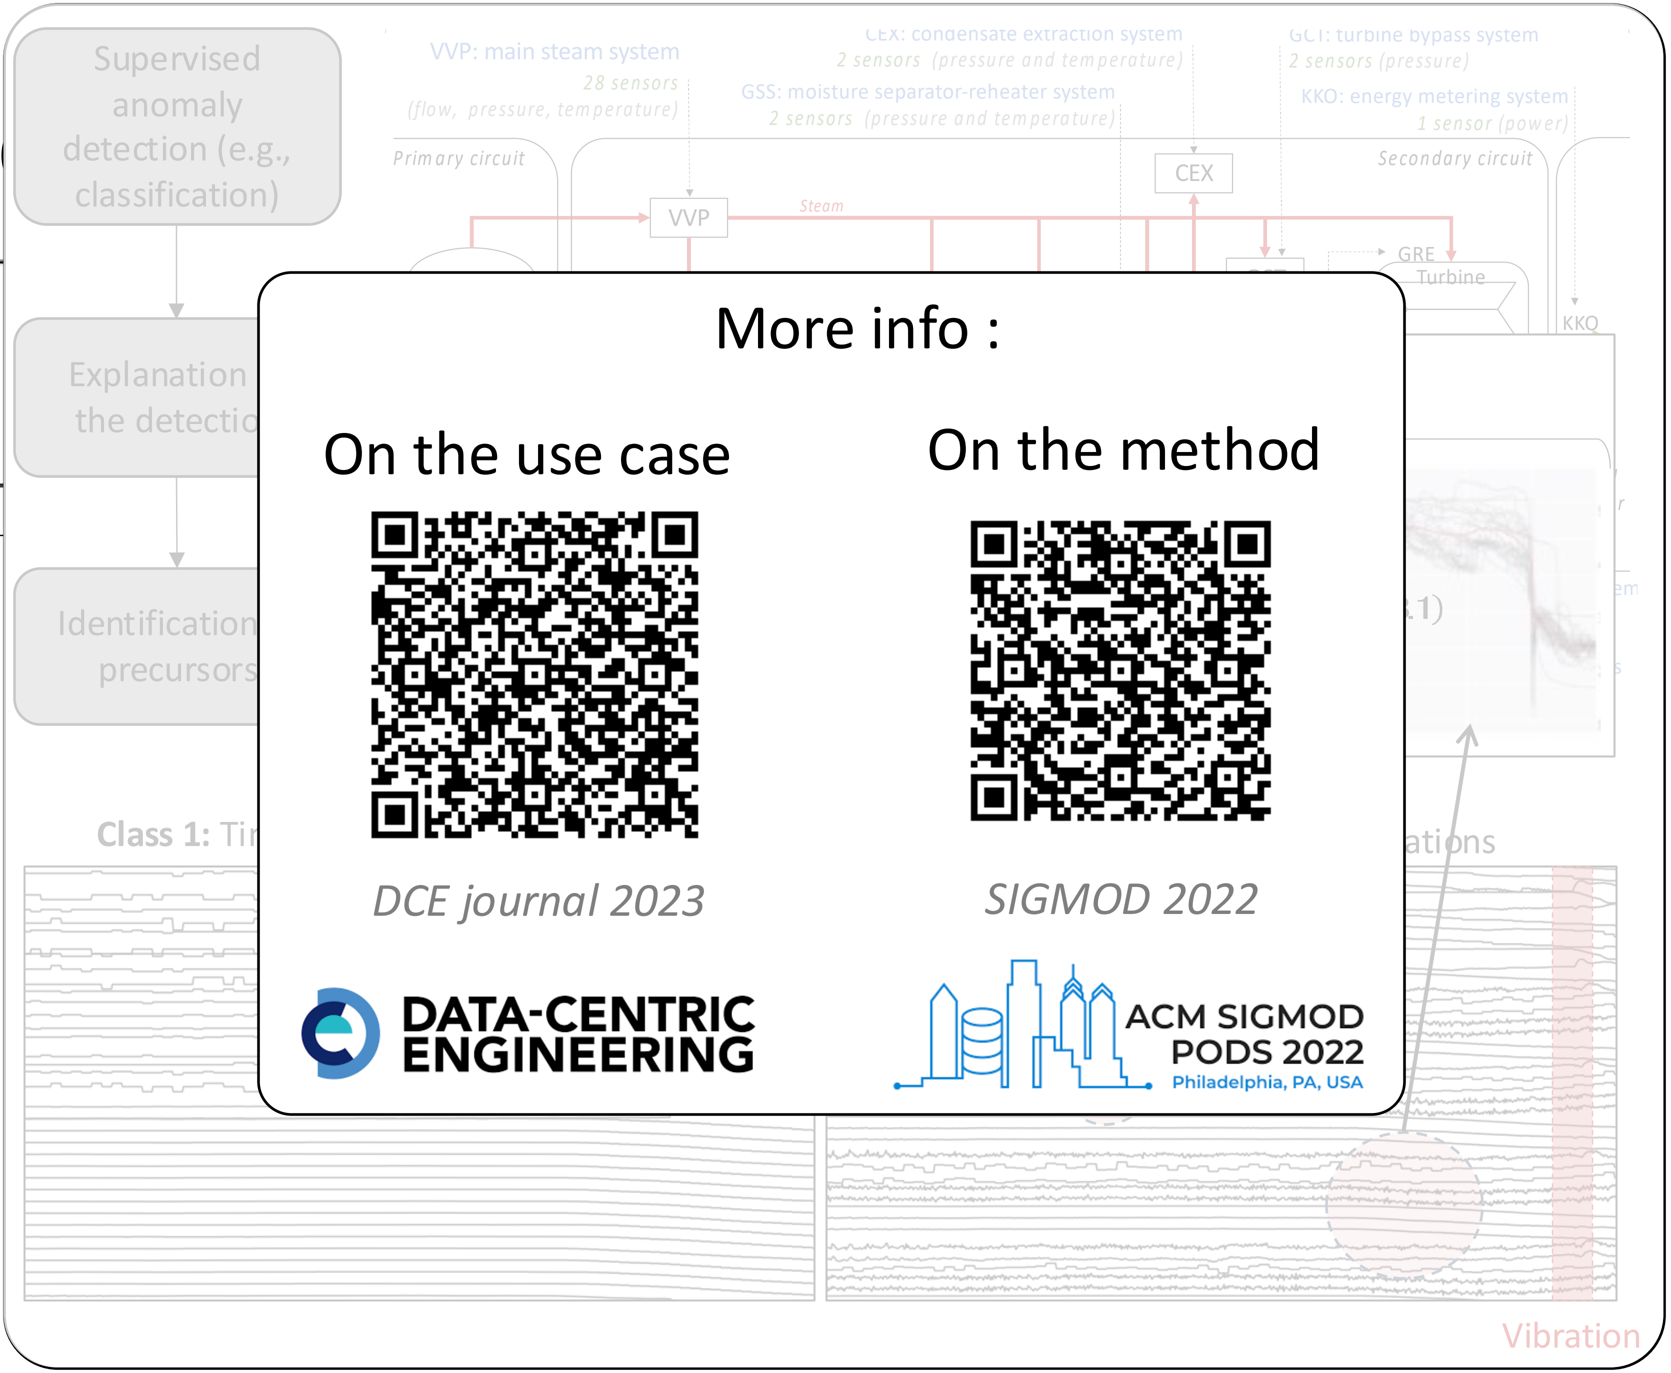

More info :

On the use case

DCE journal 2023

DATA-CENTRIC ENGINEERING

On the method

SIGMOD 2022

ACM SIGMOD PODS 2022
Philadelphia, PA, USA

Anomaly Detection methods: *A taxonomy*

By methods…

Time series anomaly detection methods

Anomaly Detection methods: *A taxonomy*

By methods…

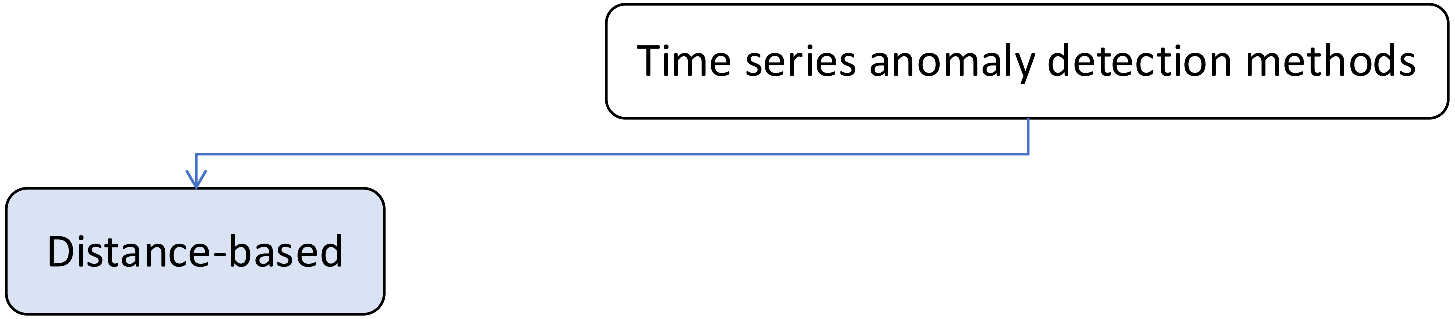

Anomaly Detection methods: *A taxonomy*

By methods…

Time series anomaly detection methods

Distance-based

Proximity-based

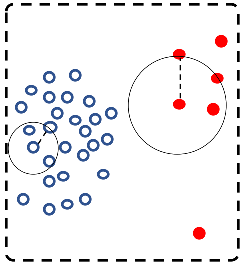

Anomaly Detection methods: *A taxonomy*

By methods…

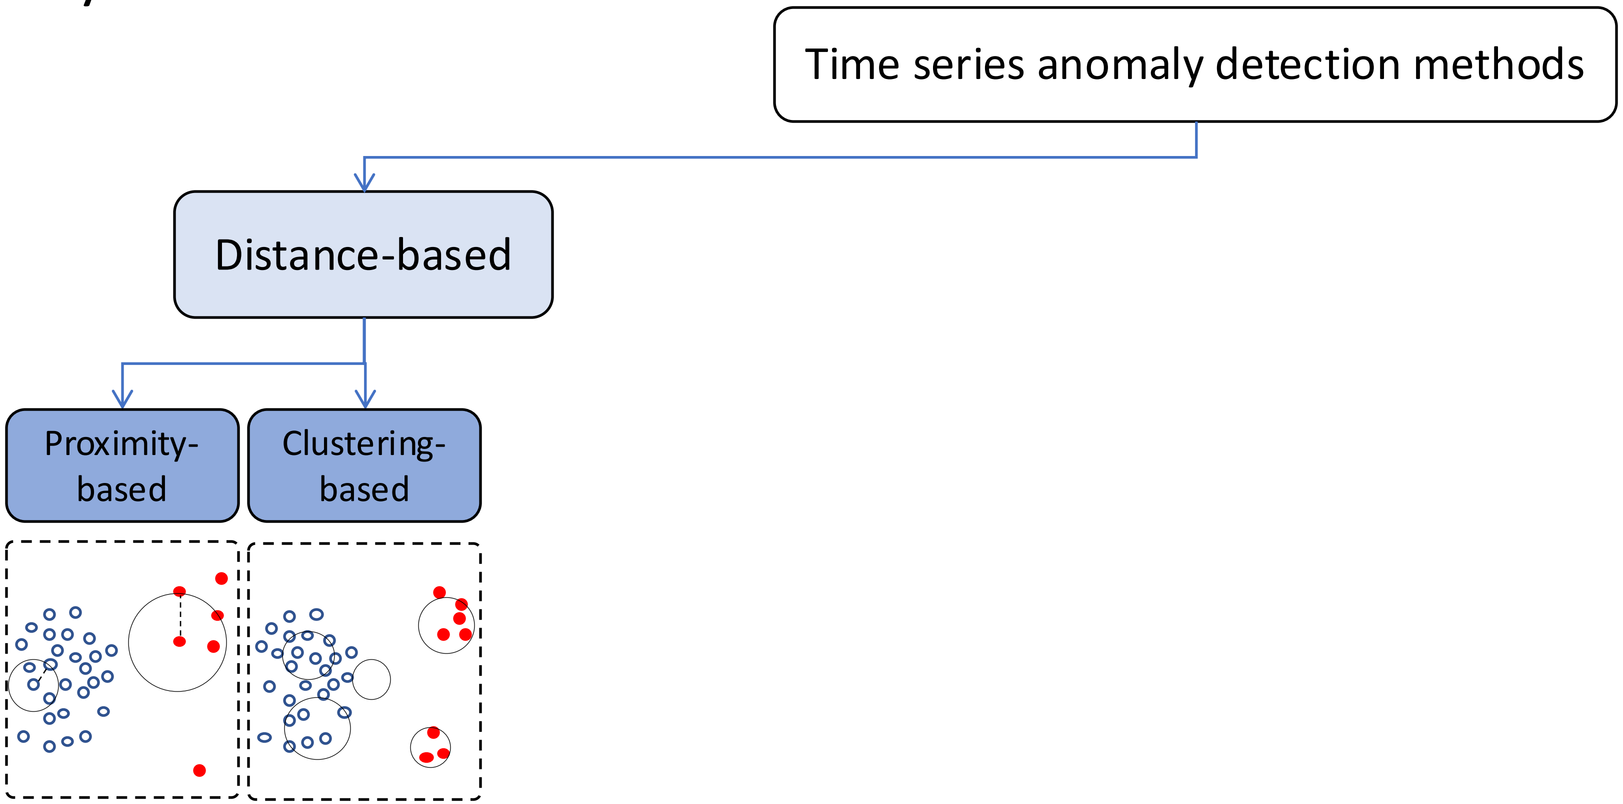

Anomaly Detection methods: *A taxonomy*

By methods…

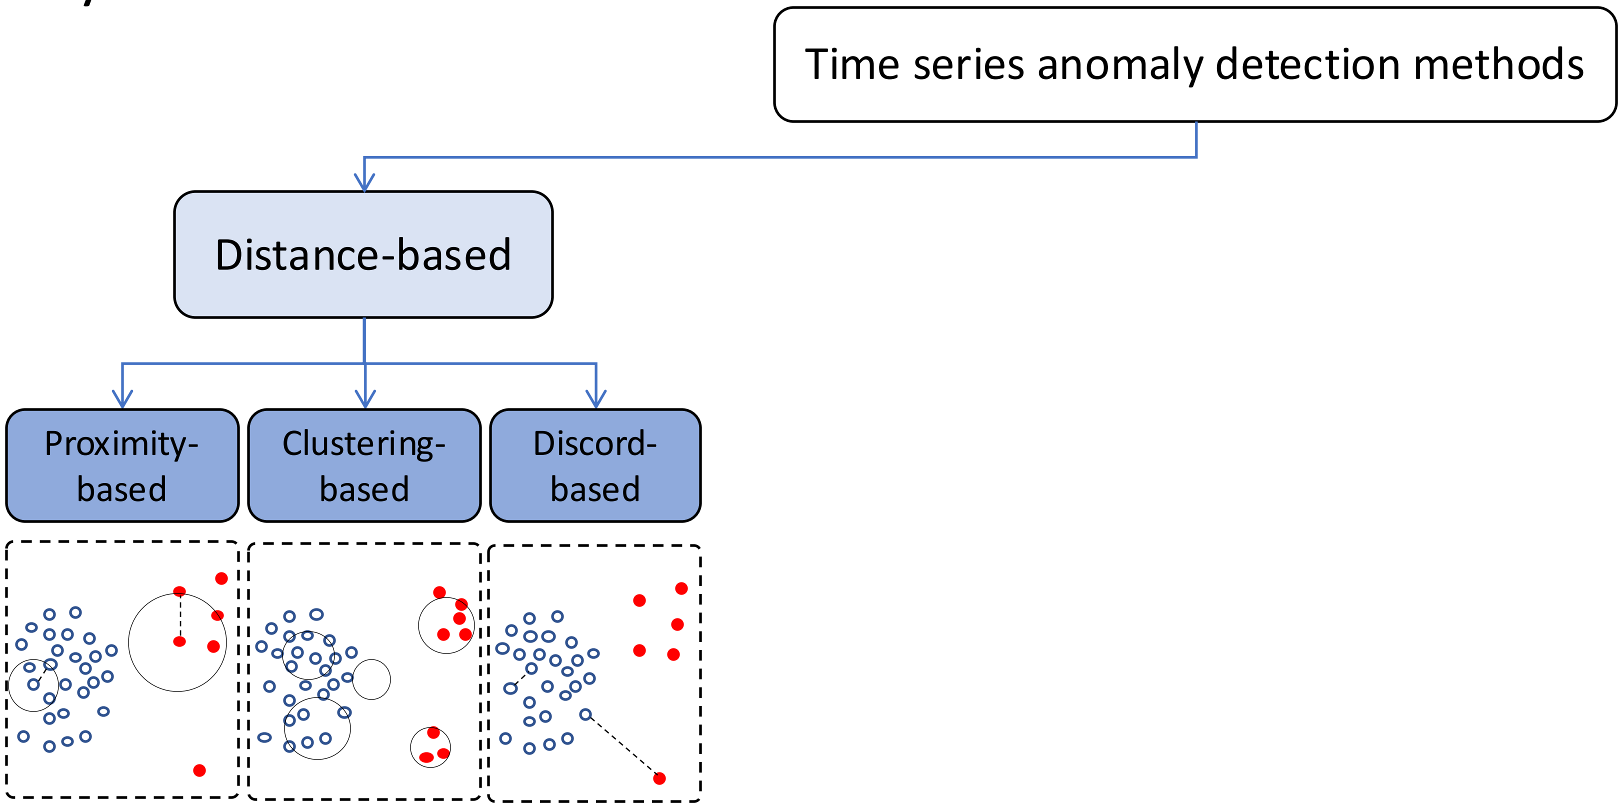

Anomaly Detection methods: *A taxonomy*

By methods…

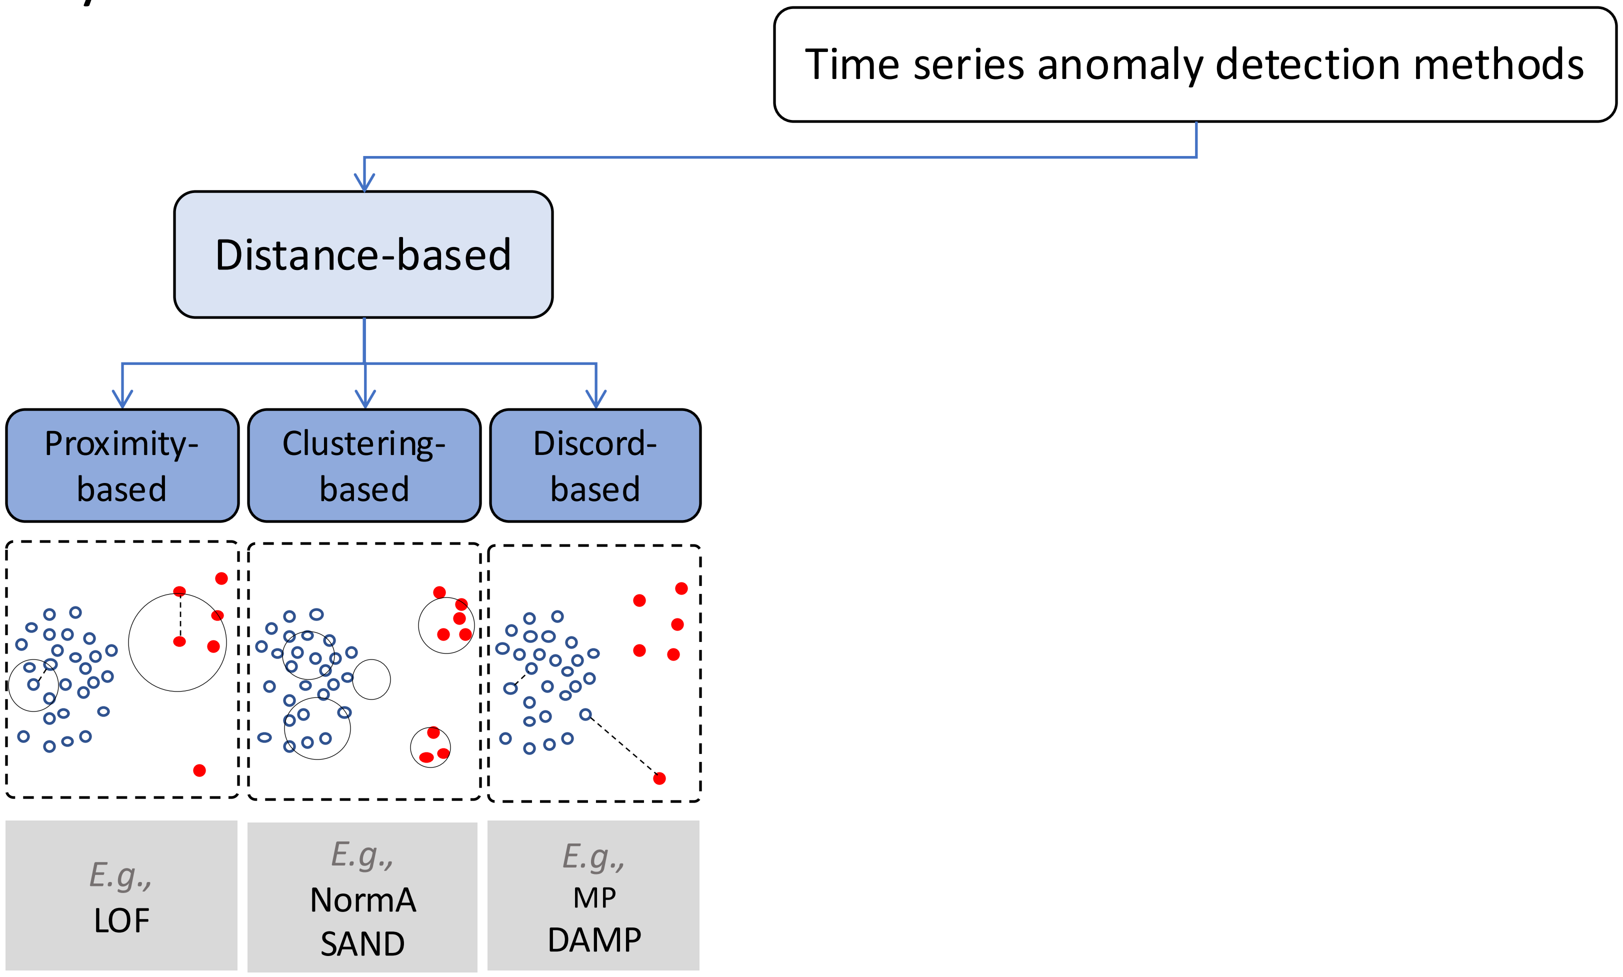

Anomaly Detection methods: *A taxonomy*

By methods…

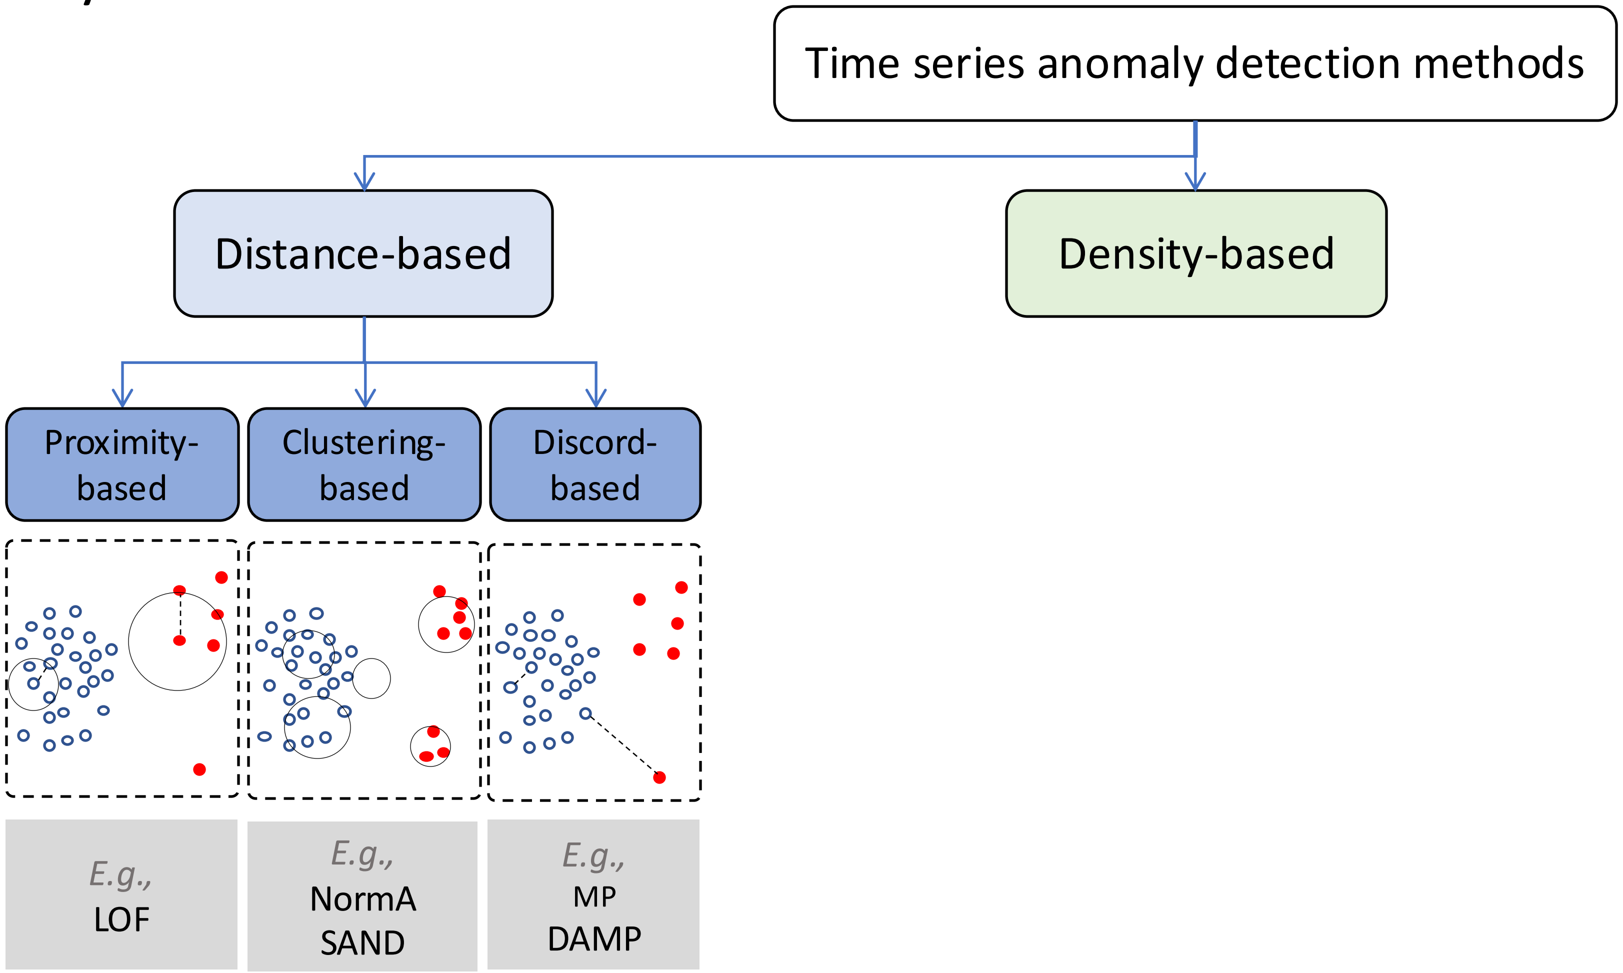

Anomaly Detection methods: *A taxonomy*

By methods…

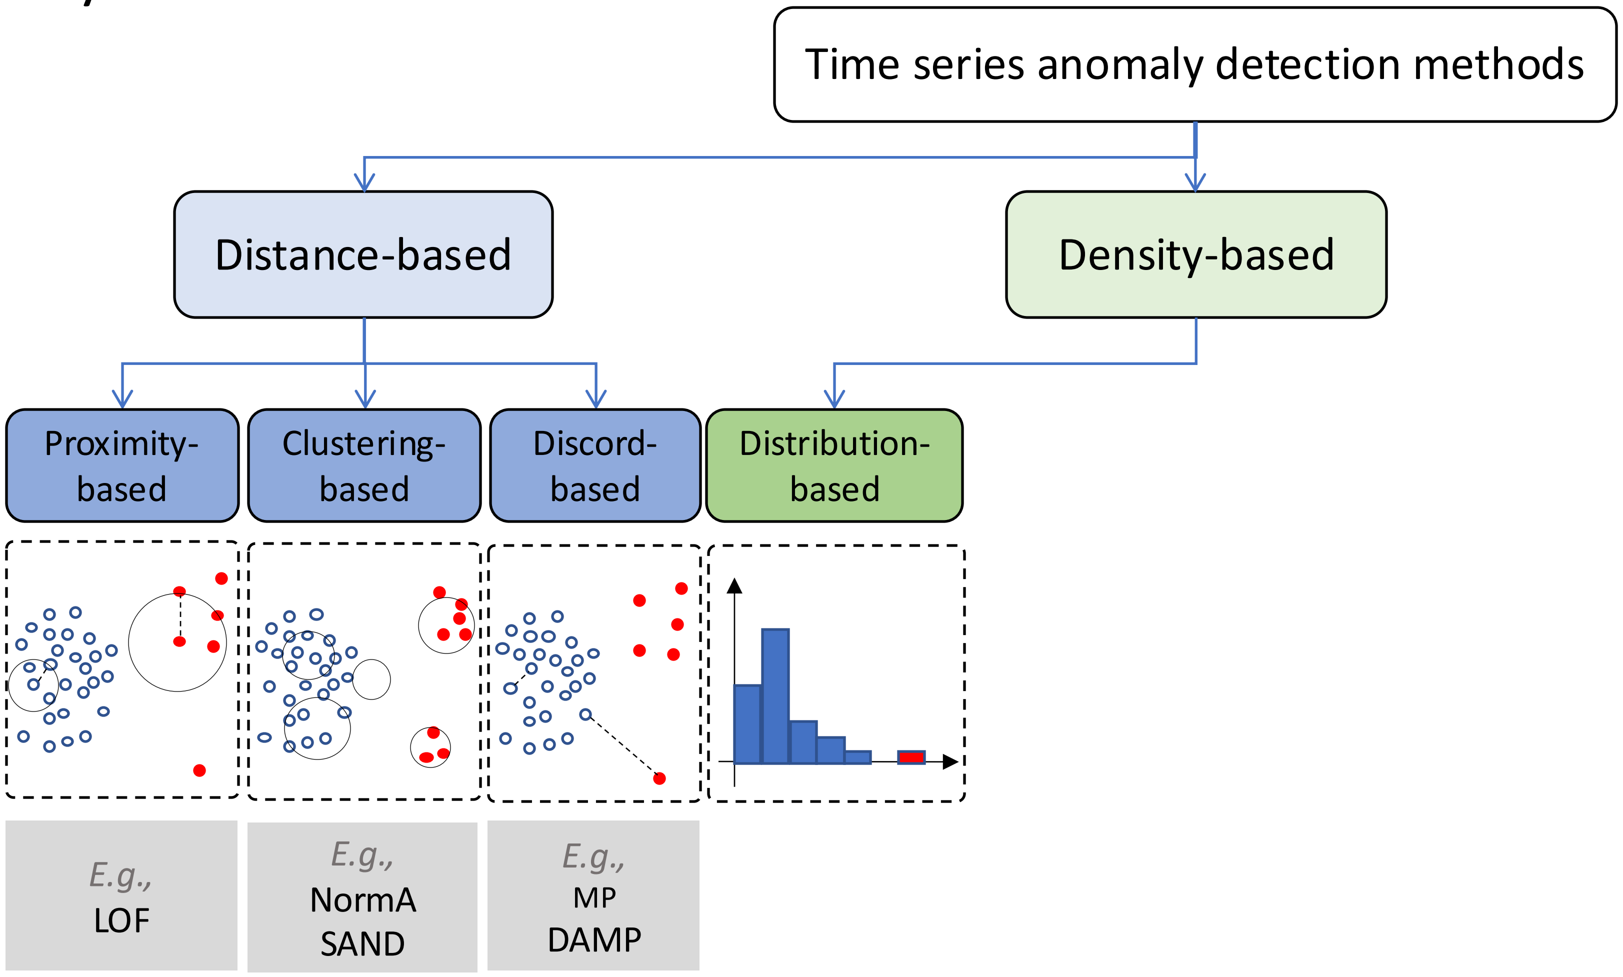

Anomaly Detection methods: *A taxonomy*

By methods…

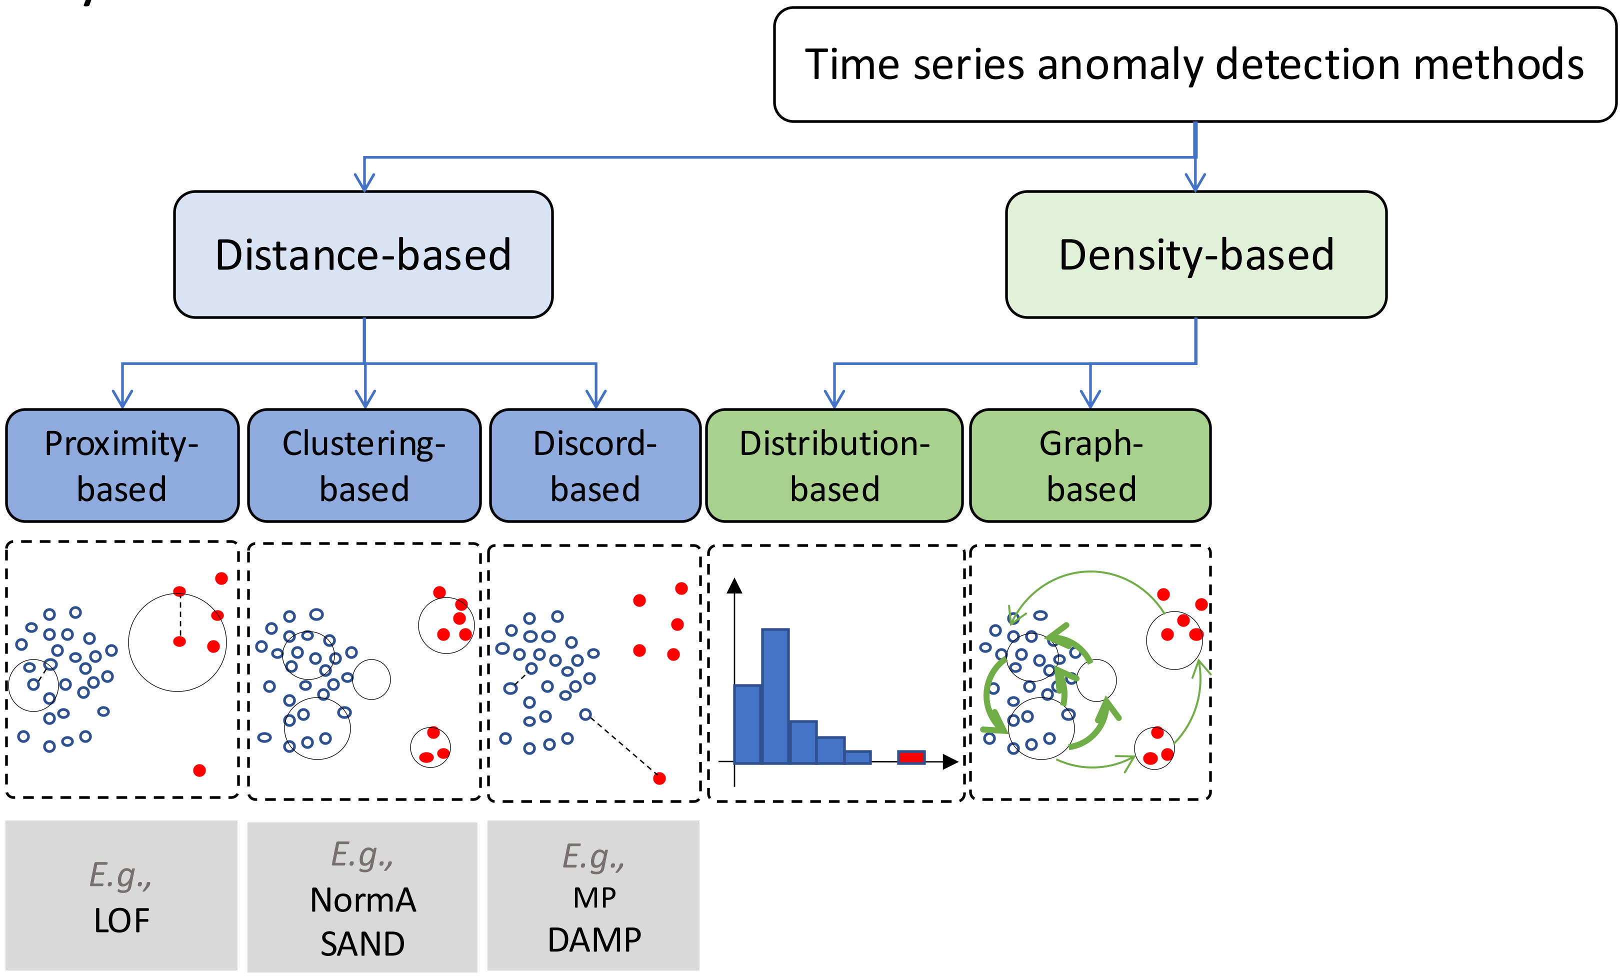

Anomaly Detection methods: *A taxonomy*

By methods…

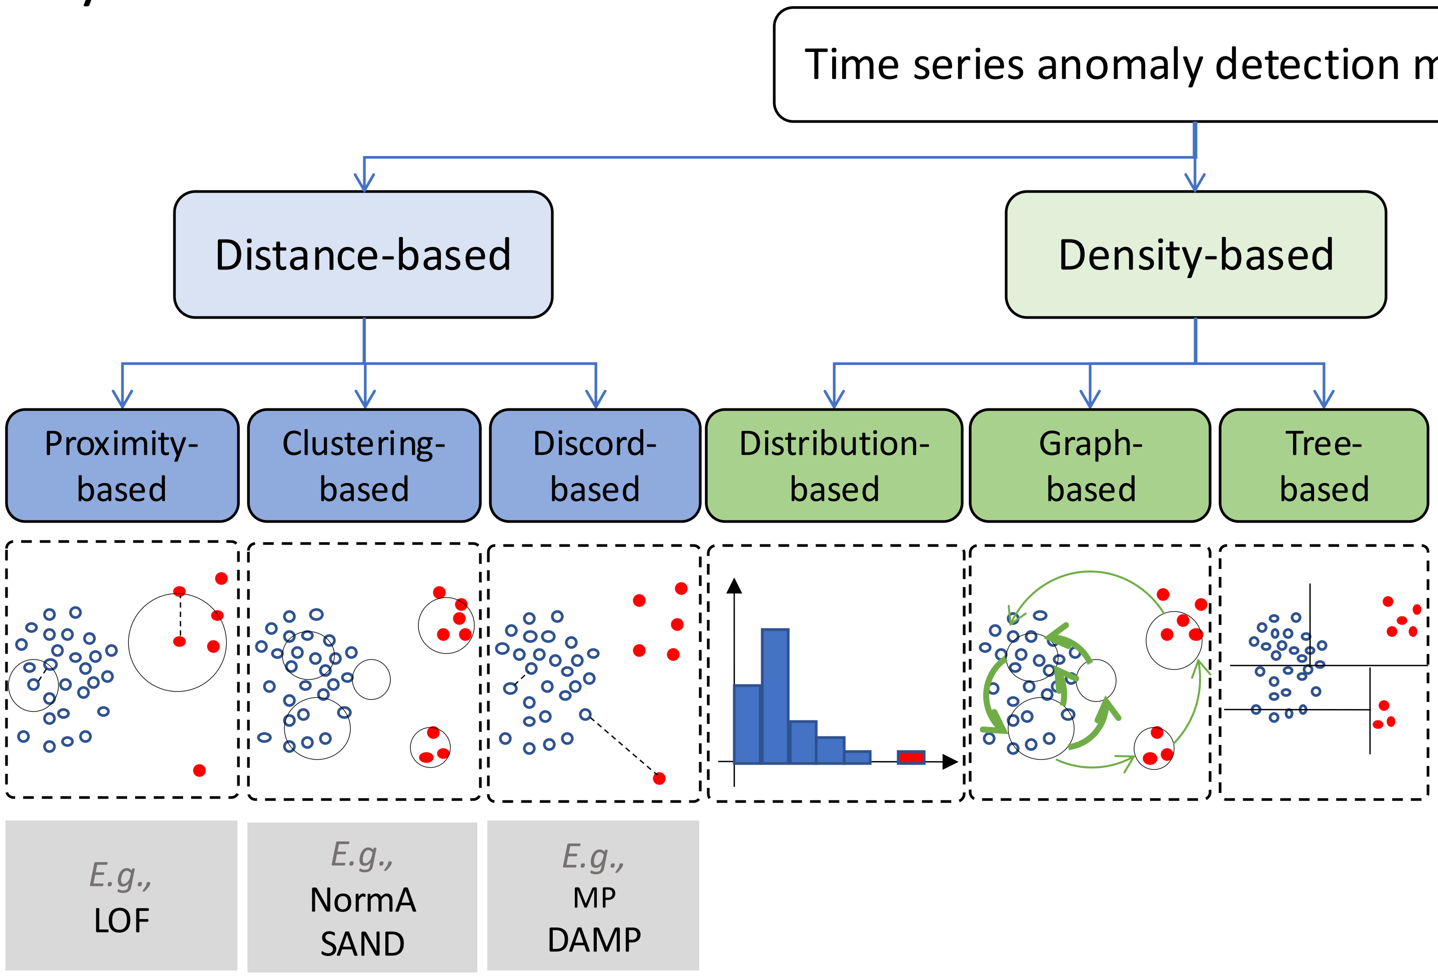

Anomaly Detection methods: *A taxonomy*

By methods…

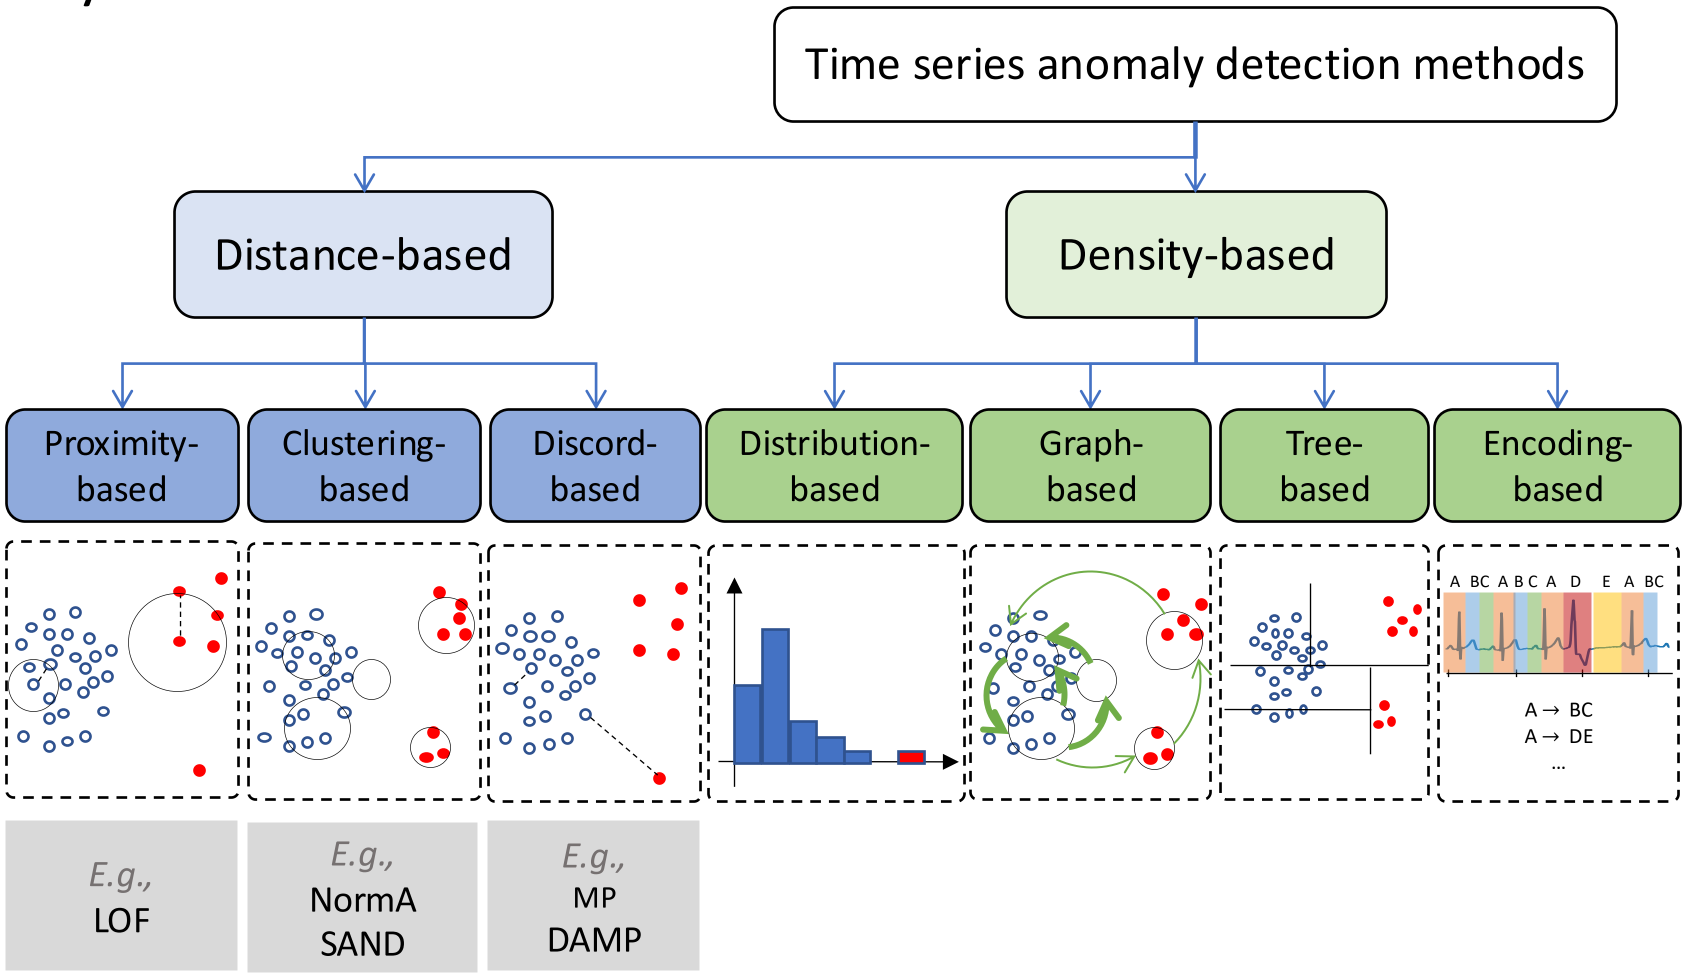

Anomaly Detection methods: *A taxonomy*

By methods…

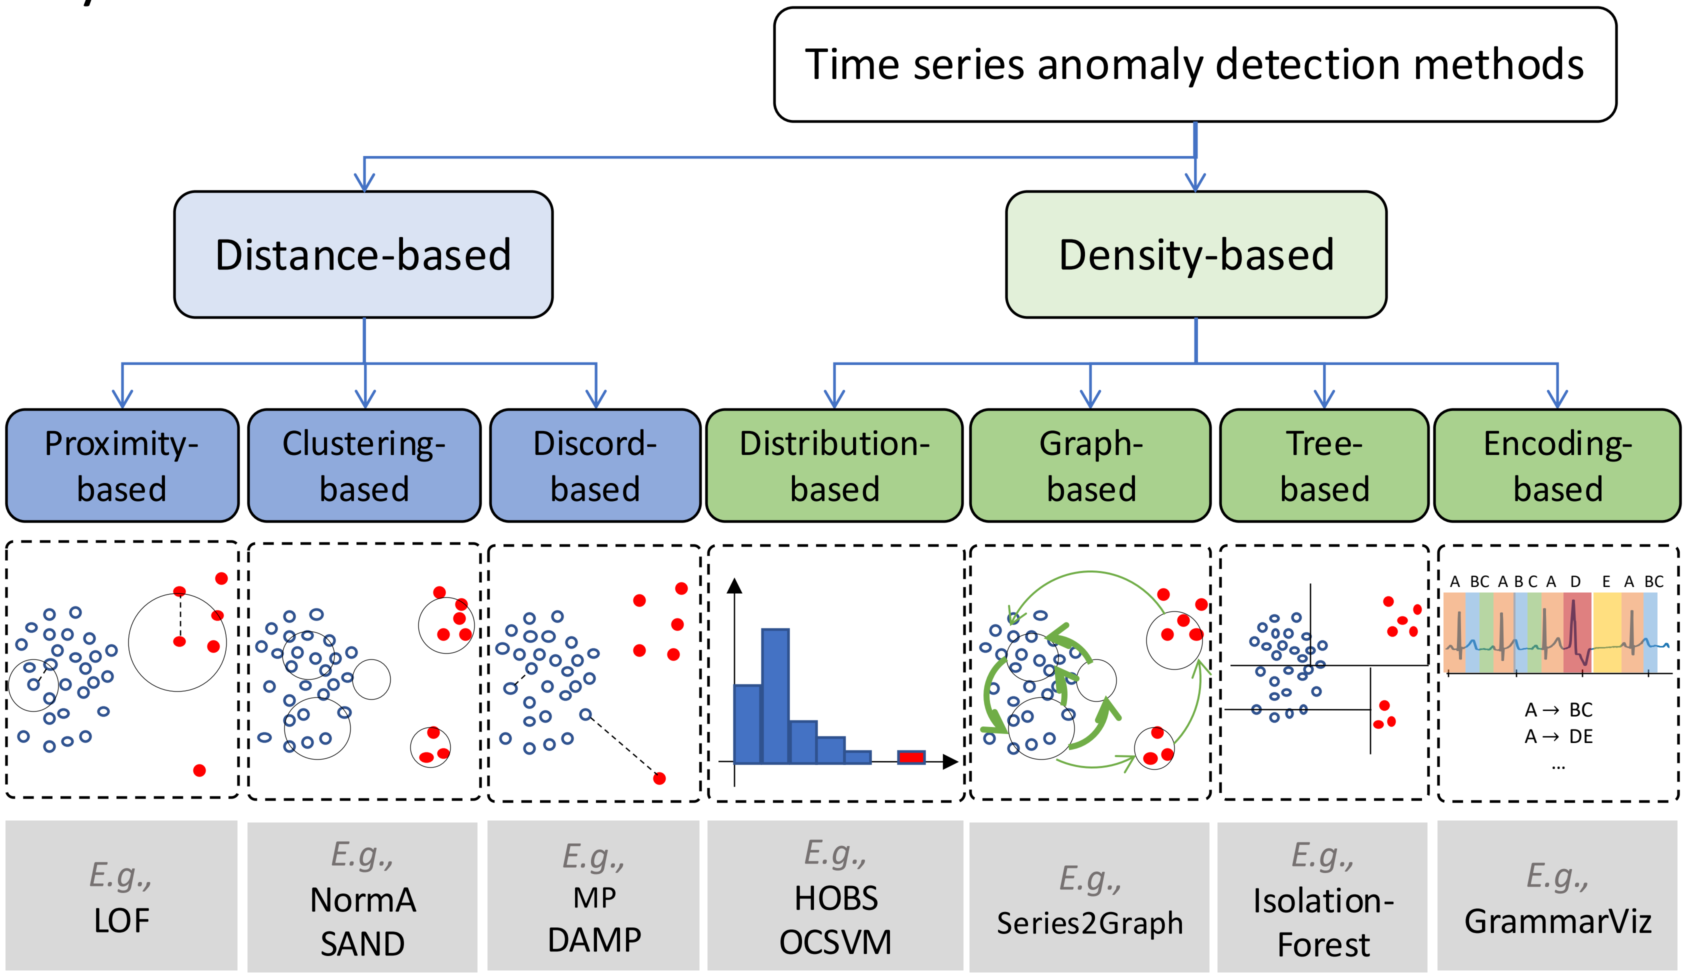

Anomaly Detection methods: *A taxonomy*

By methods…

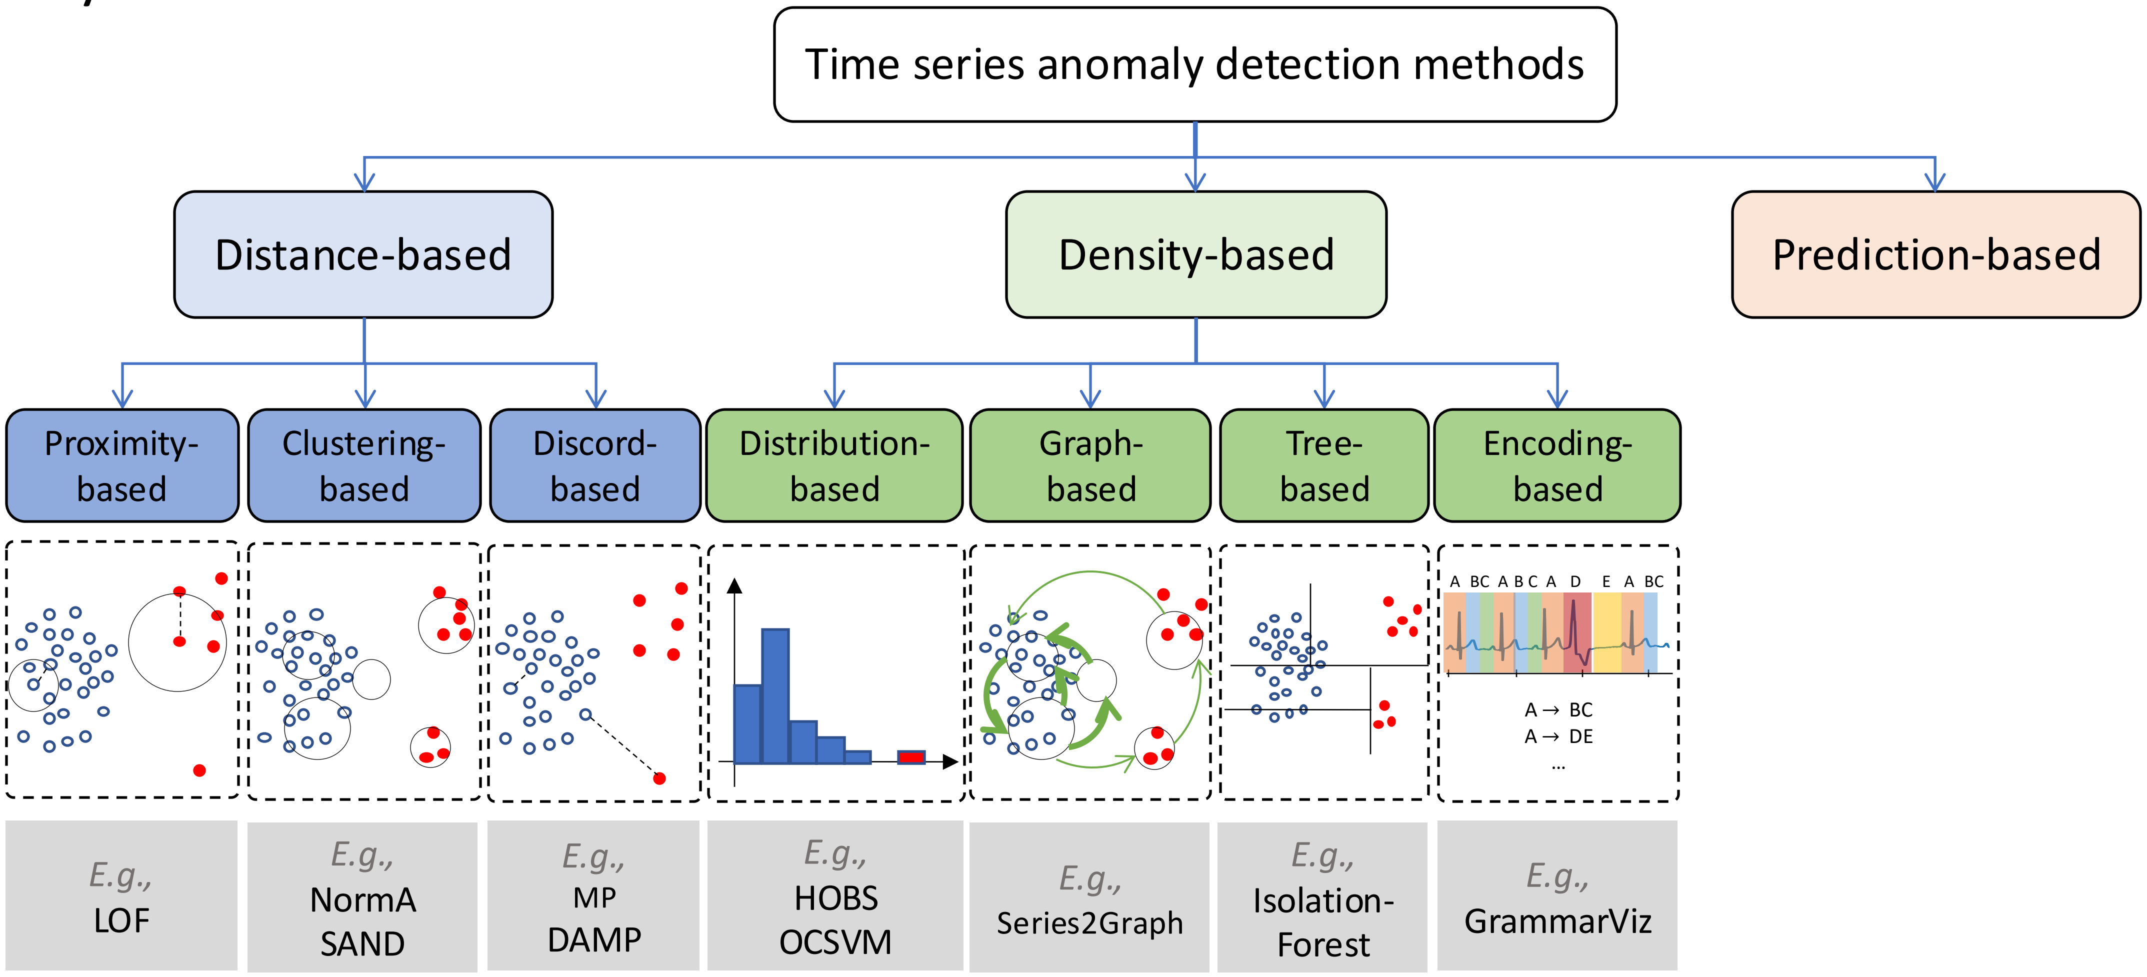

Anomaly Detection methods: *A taxonomy*

By methods…

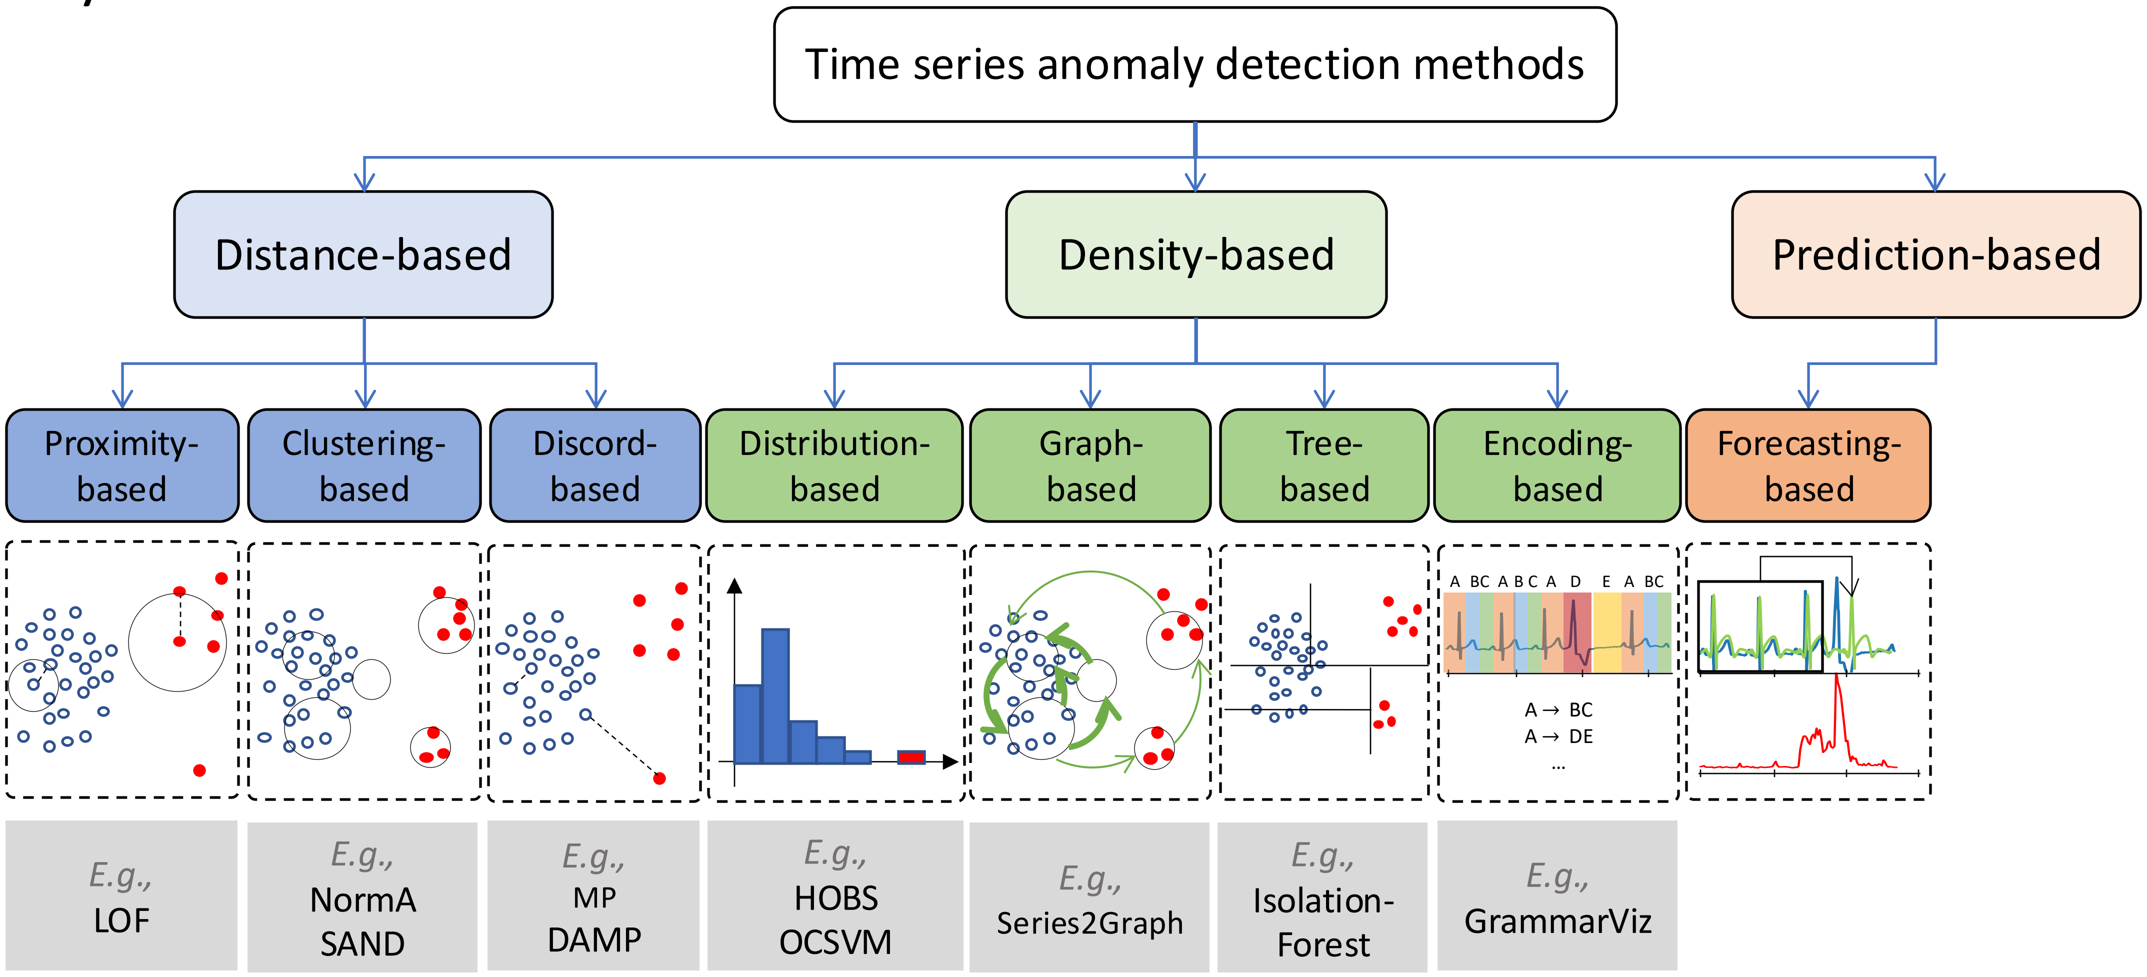

Anomaly Detection methods: *A taxonomy*

By methods…

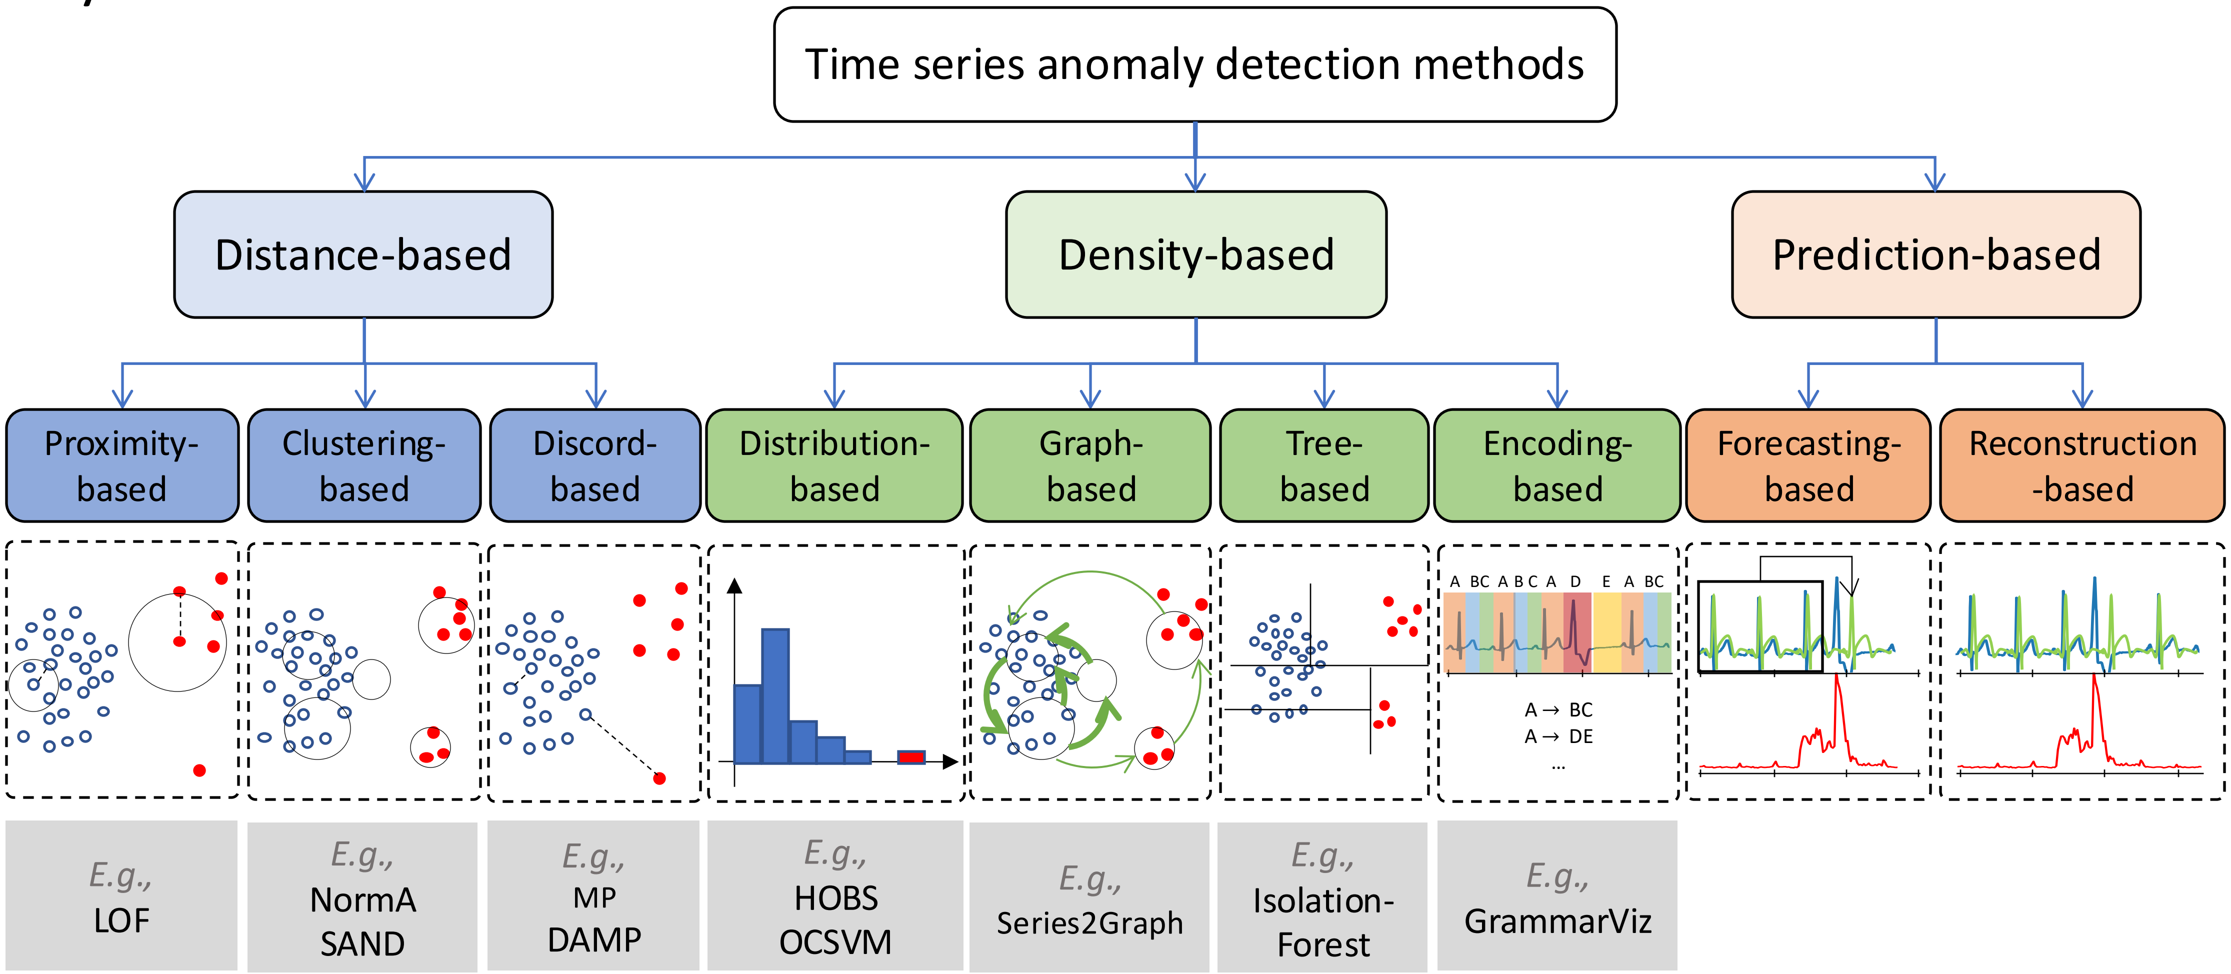

Anomaly Detection methods: *A taxonomy*

By methods…

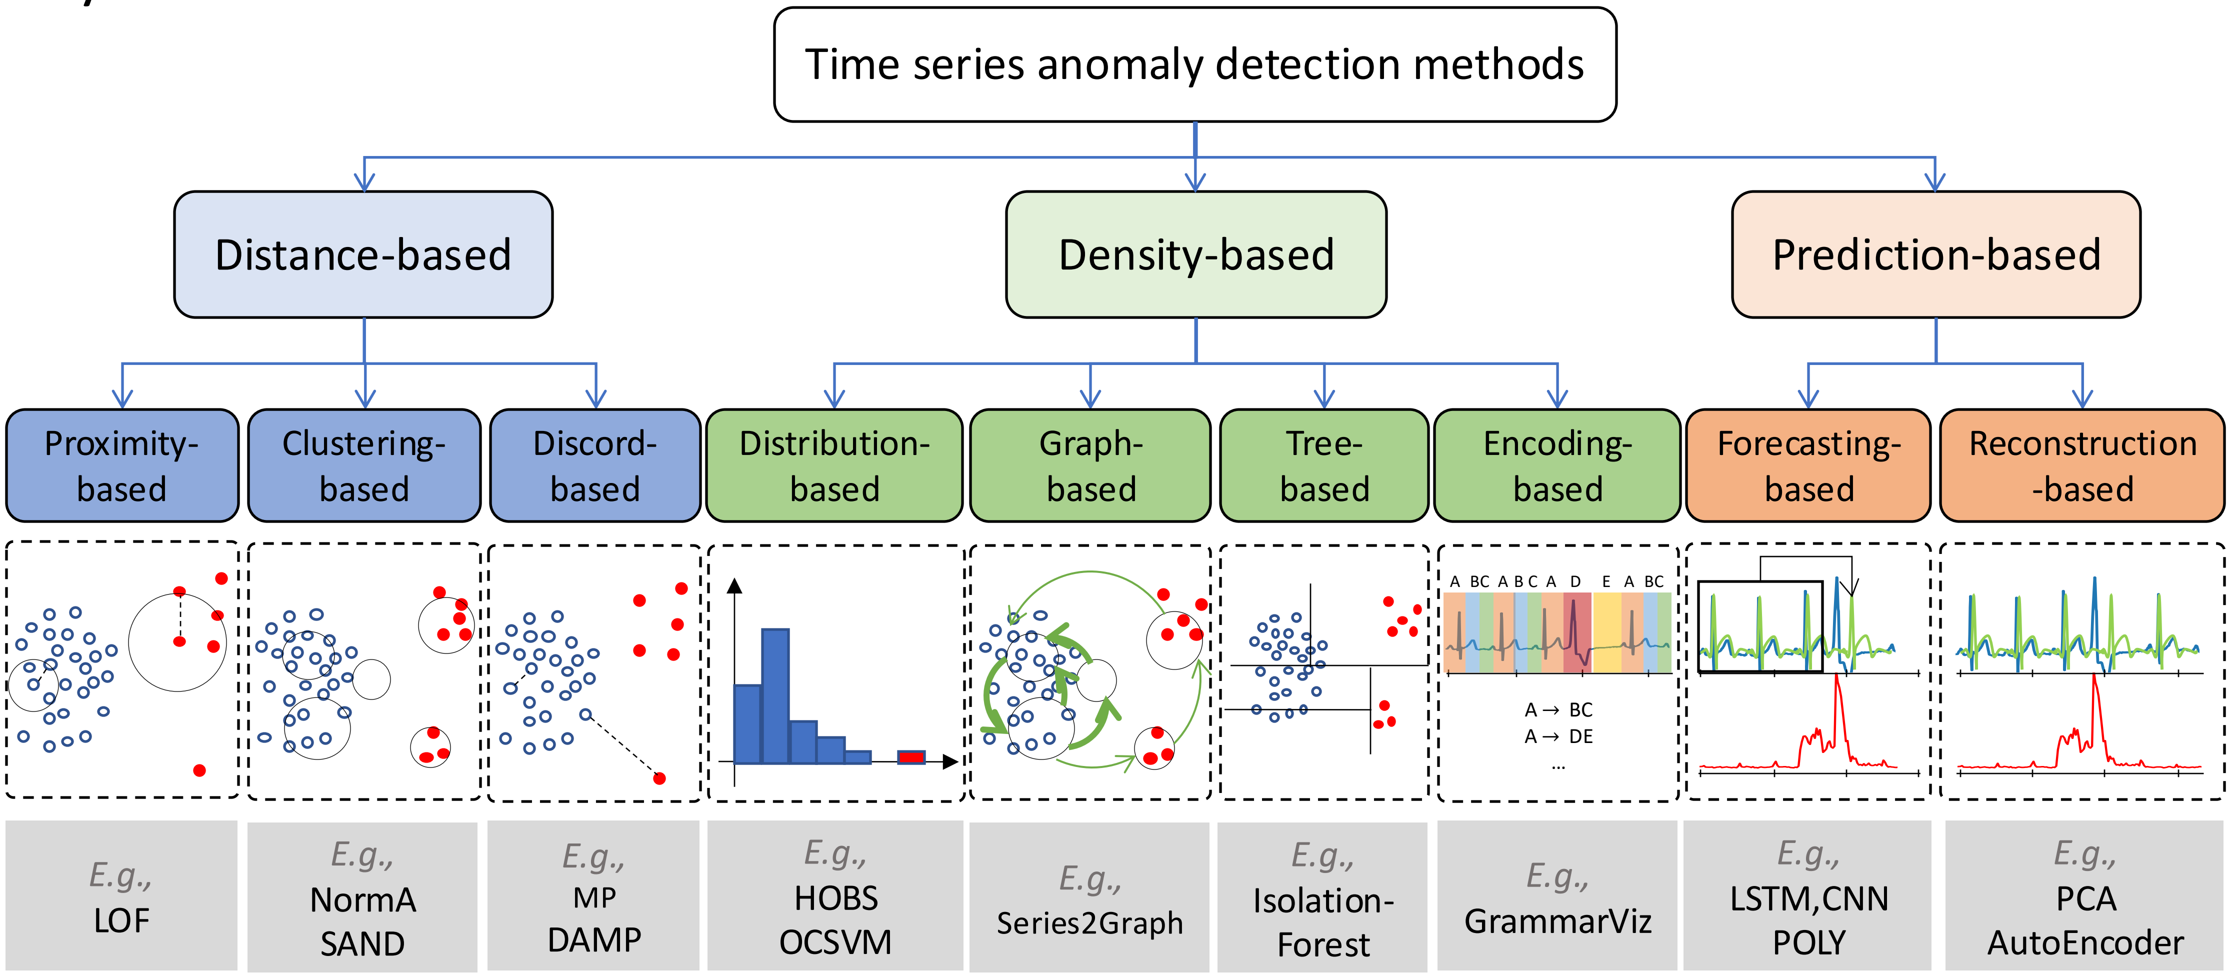

Anomaly Detection methods: *A taxonomy*

By time…

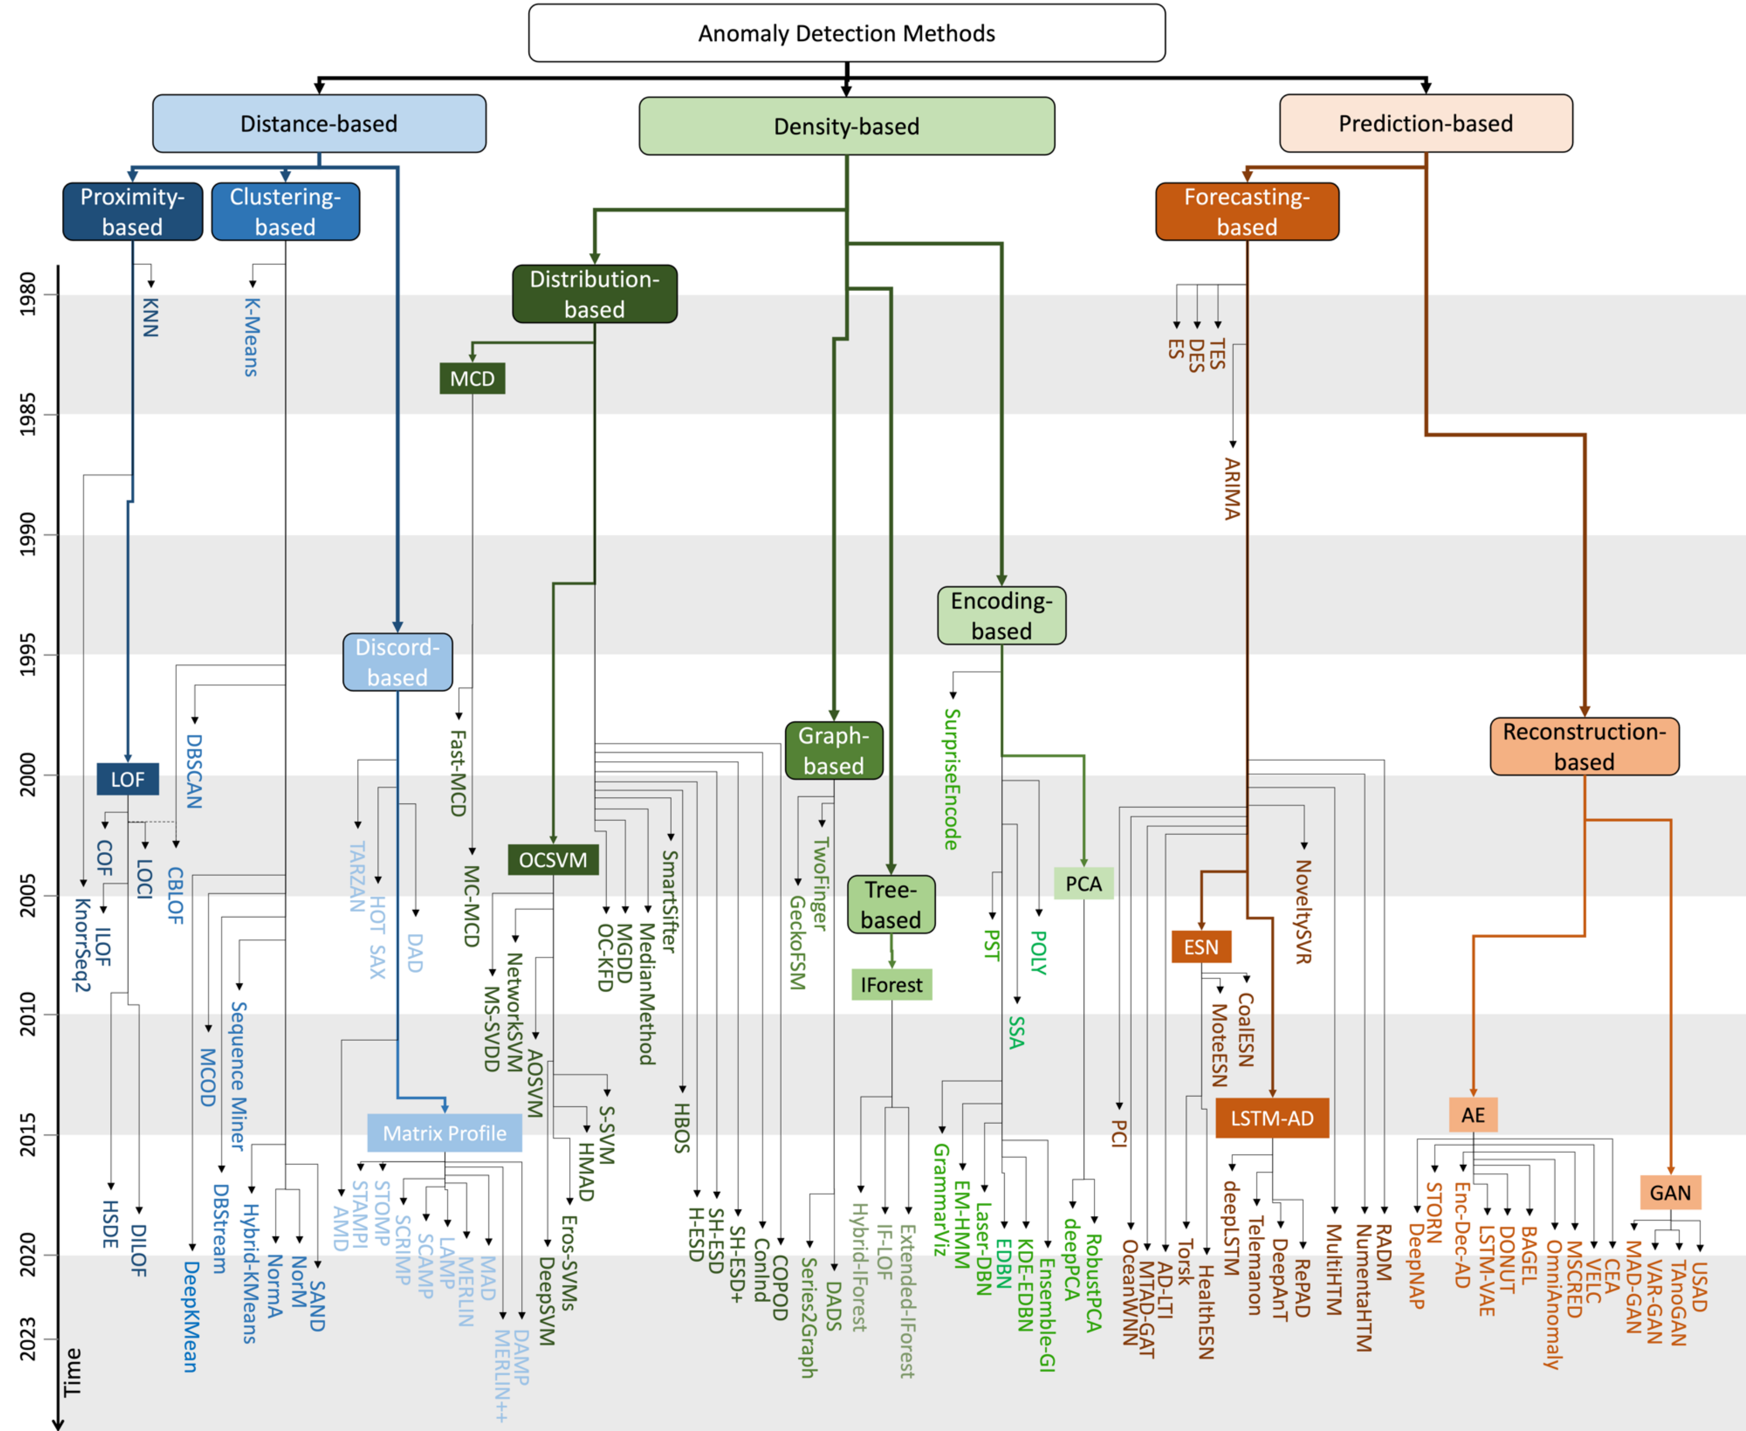

Ano
Dete
met
A ta

By time

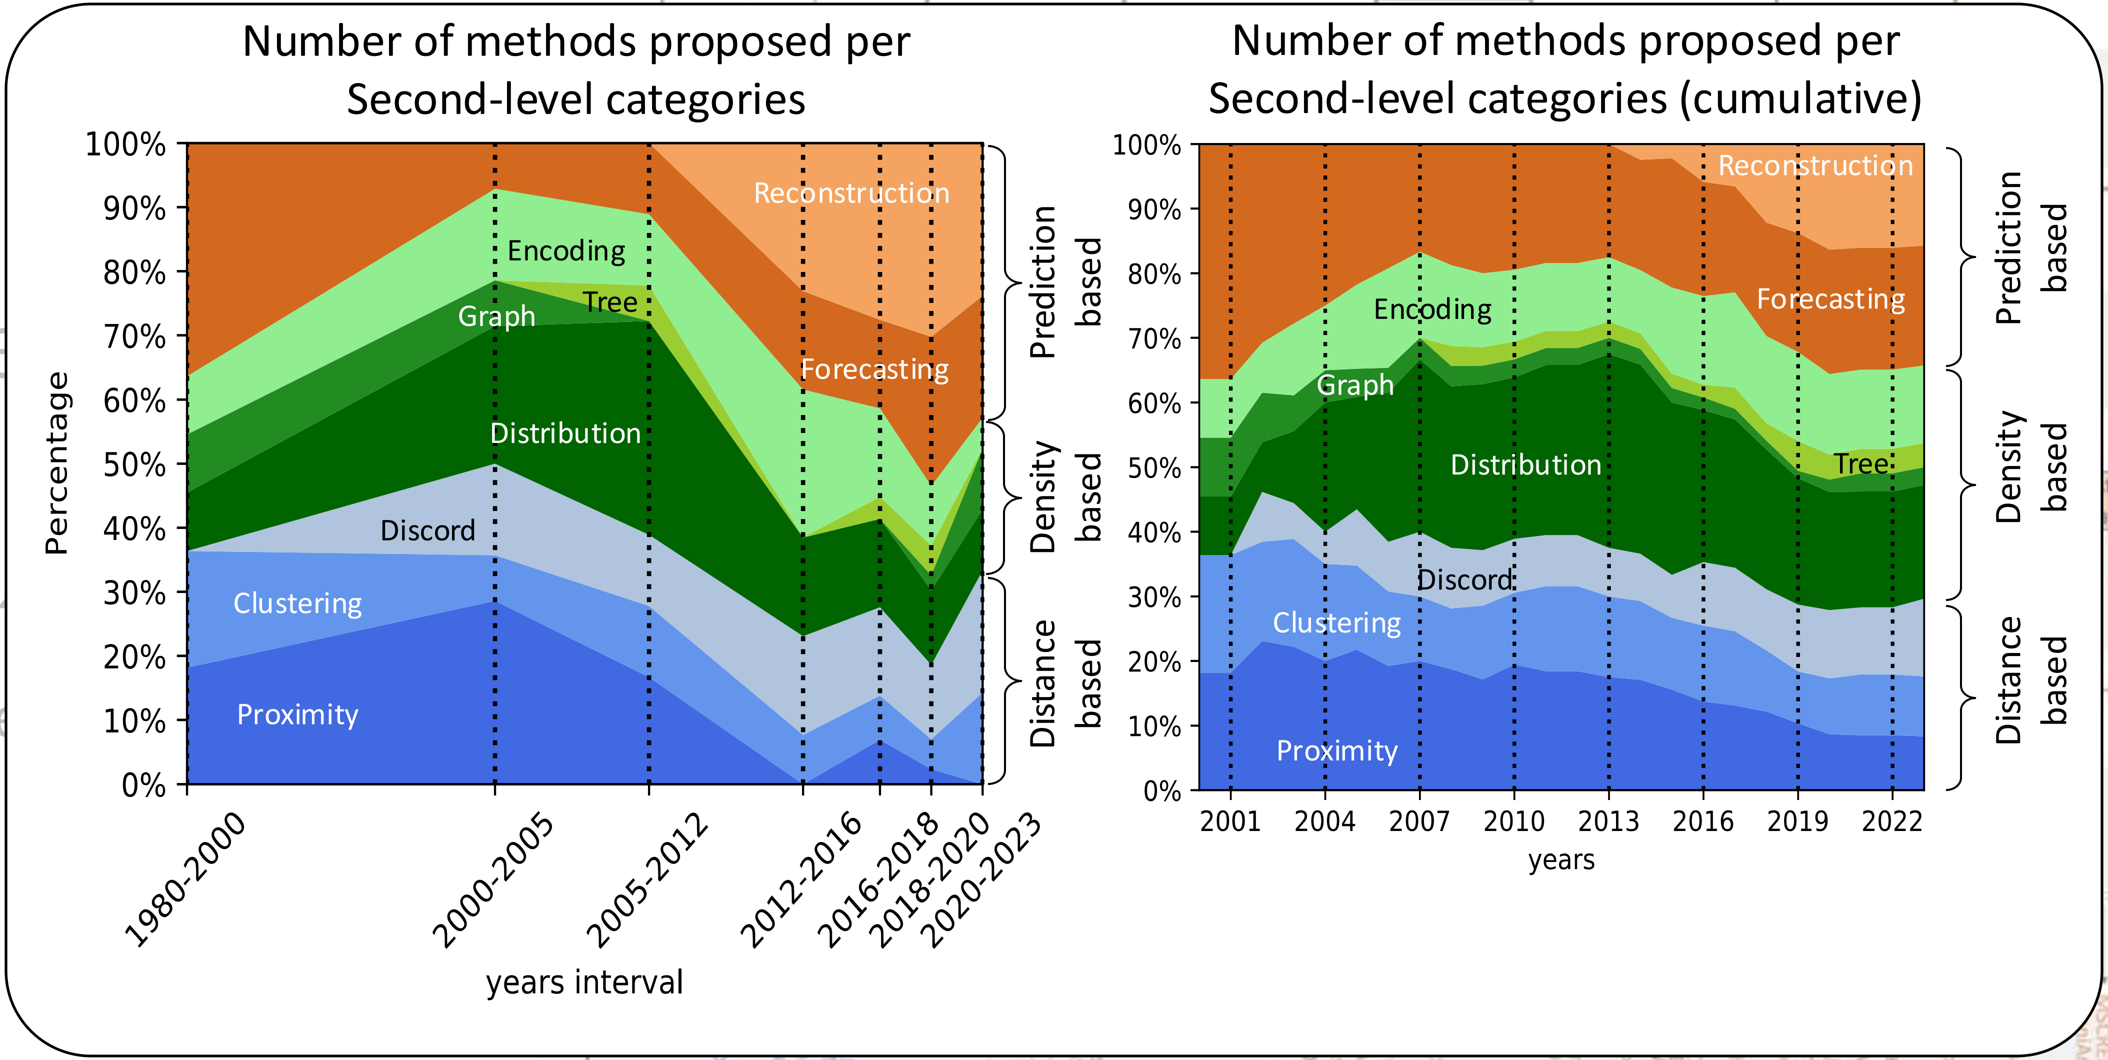

Ano
Dete
met
A ta

By time

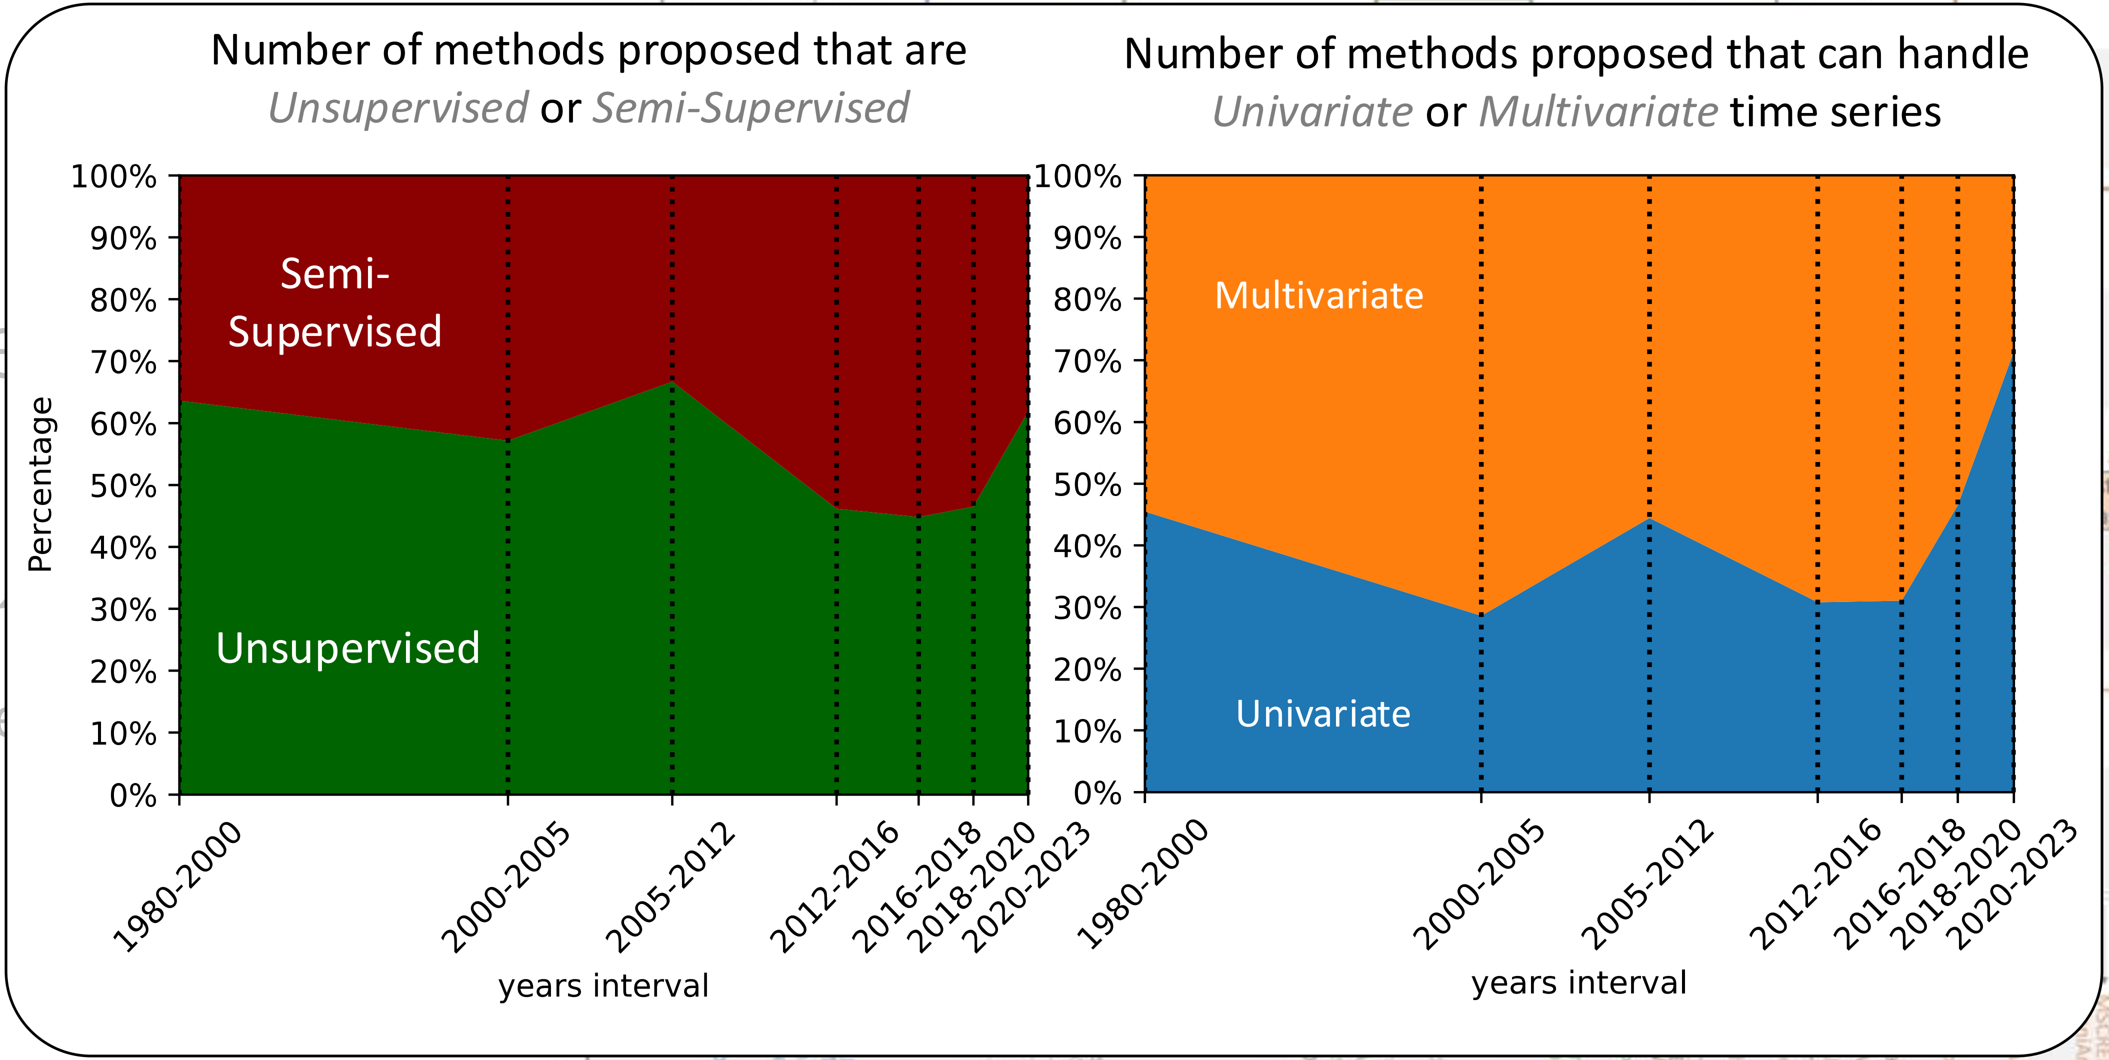

Ano
Dete
met
A ta

By time

Time Series Anomaly Detection

Paul Boniol, Qinghua Liu, John Paparrizos, and Themis Palpanas.

Video (EDBT 2023 Tutorial)

https://www.youtube.com/watch?v=96869qimXAA&t=1s

Slides (VLDB 2024 Tutorial)

https://drive.google.com/file/d/1Vyz6H0E16IpbVZXgtiZVnU9Ie8zAJaog/view

SIGMOD Blogpost

https://wp.sigmod.org/?p=3739

Anomaly Detection methods: *A taxonomy*

By time…

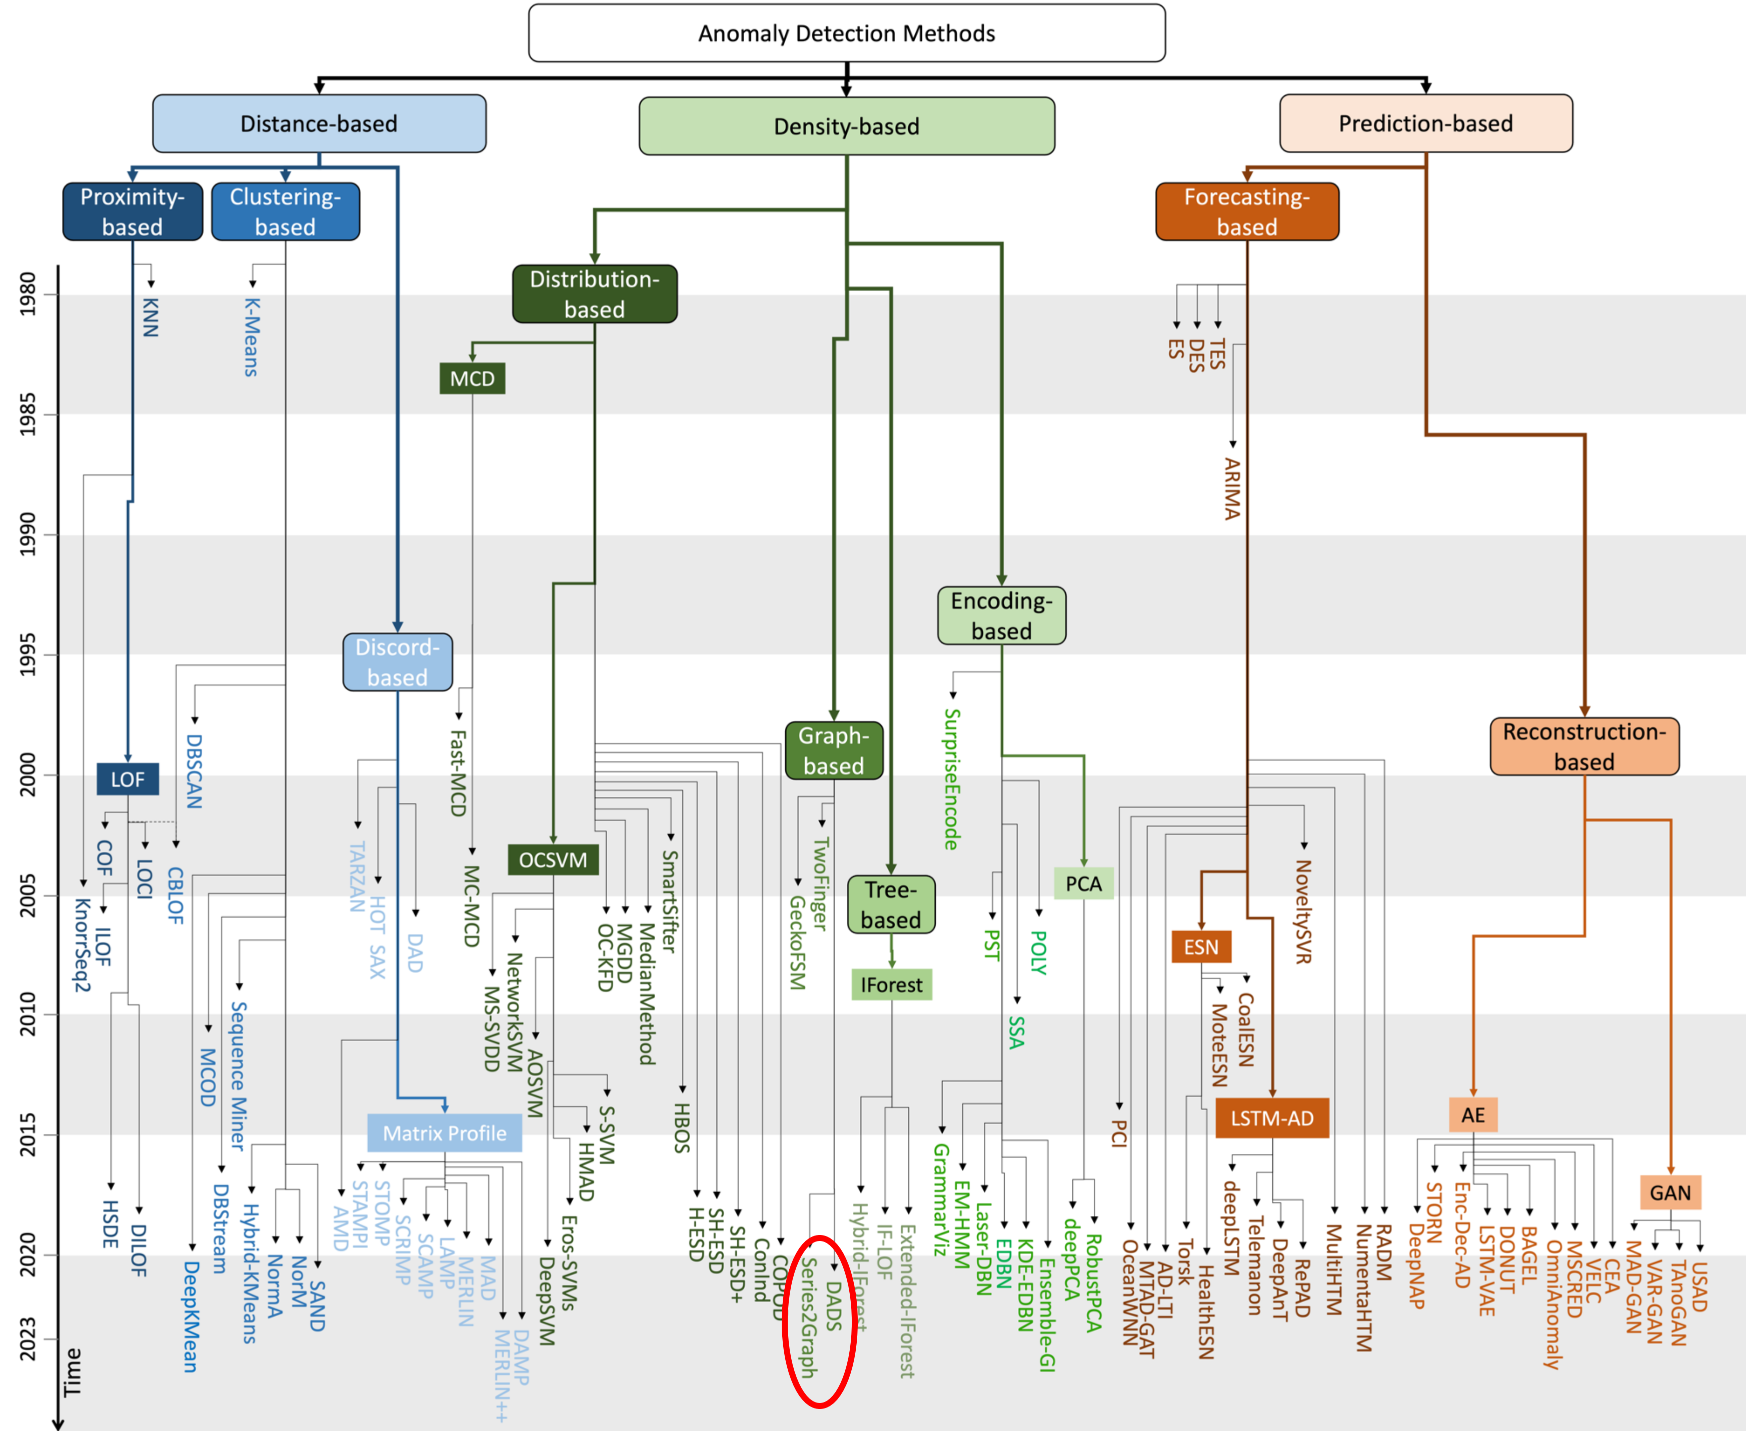

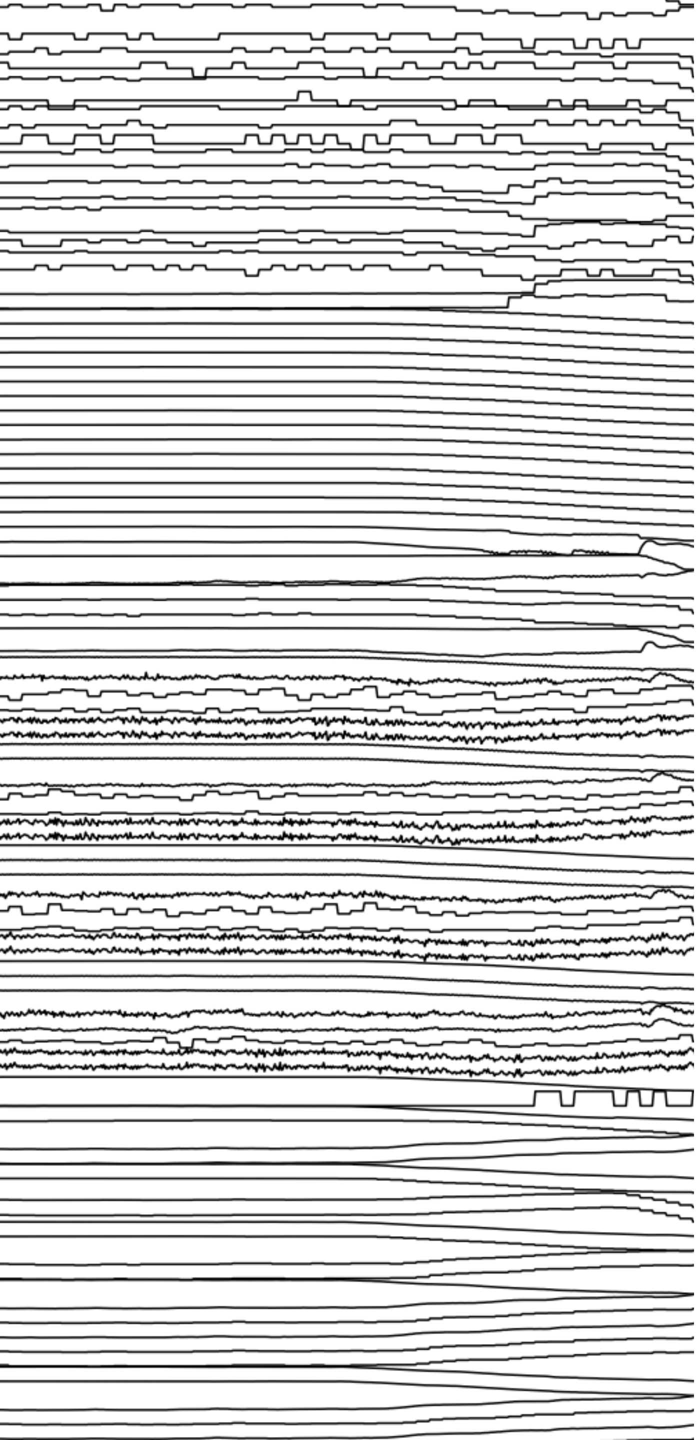

III. Series2Graph

A graph-based approach

Series2Graph: *From time series to a graph*

Converting the time series to a graph:

- Existing solutions create a node per point (e.g., Visibility Graph [6,7])
- Do not scale for large time series

Visibility Graph

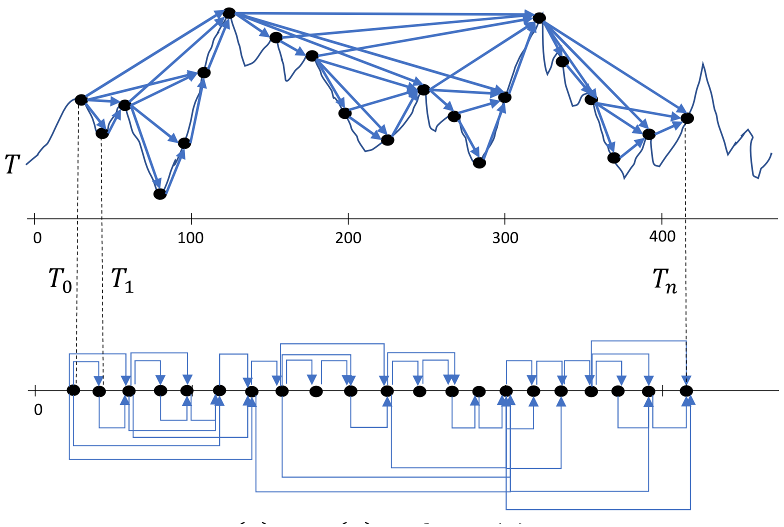

Complex network for time series [8]

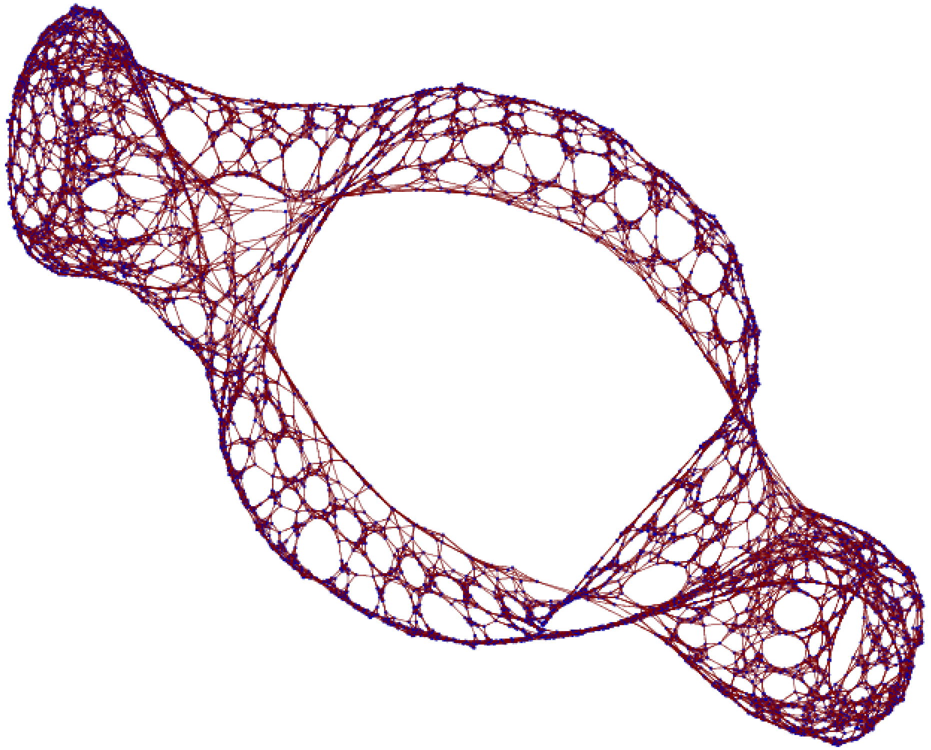

M. Small et al. Transforming Time seires into Complex Networks, Complex Sciences (2009)

Series2Graph: *From time series to a graph*

Graph G_{ℓ_G} [9]:

Given a data series T, and an input length ℓ_G, we build a graph $G_{\ell_G}(\mathcal{N}, \mathcal{E})$ for which:

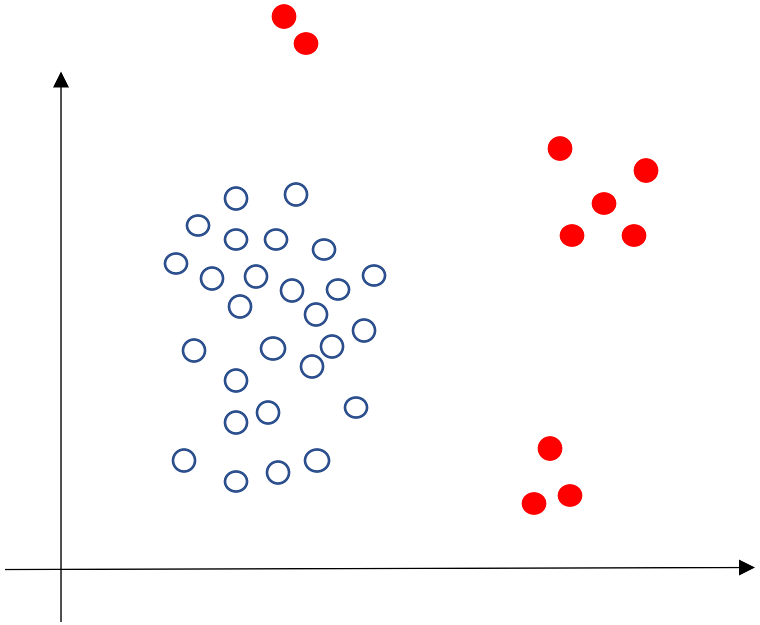

Series2Graph: *From time series to a graph*

Graph G_{ℓ_G} [9]:

Given a data series T, and an input length ℓ_G, we build a graph $G_{\ell_G}(\mathcal{N}, \mathcal{E})$ for which:

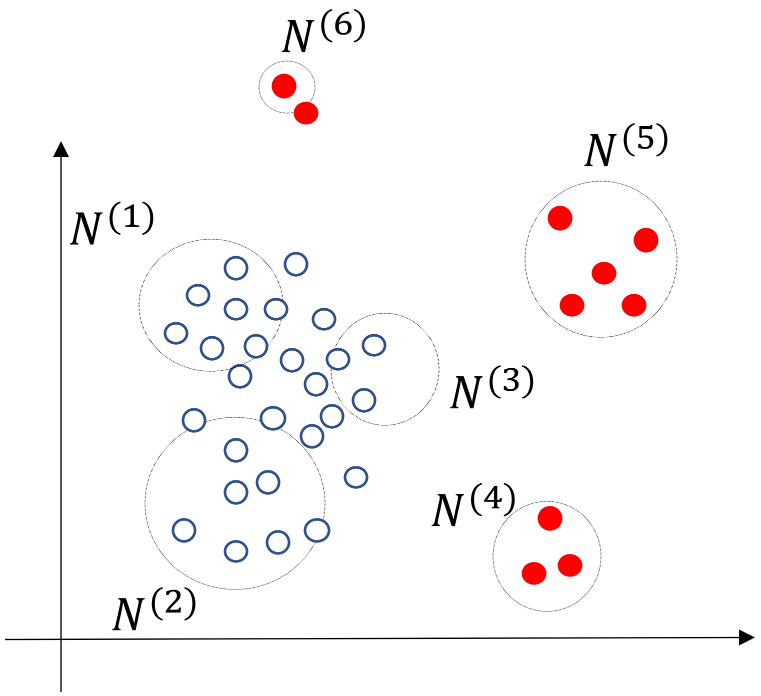

Each node is an ensemble of similar subsequences.

Series2Graph: *From time series to a graph*

Graph G_{ℓ_G} [9]:
Given a data series T, and an input length ℓ_G, we build a graph $G_{\ell_G}(\mathcal{N},\mathcal{E})$ for which:

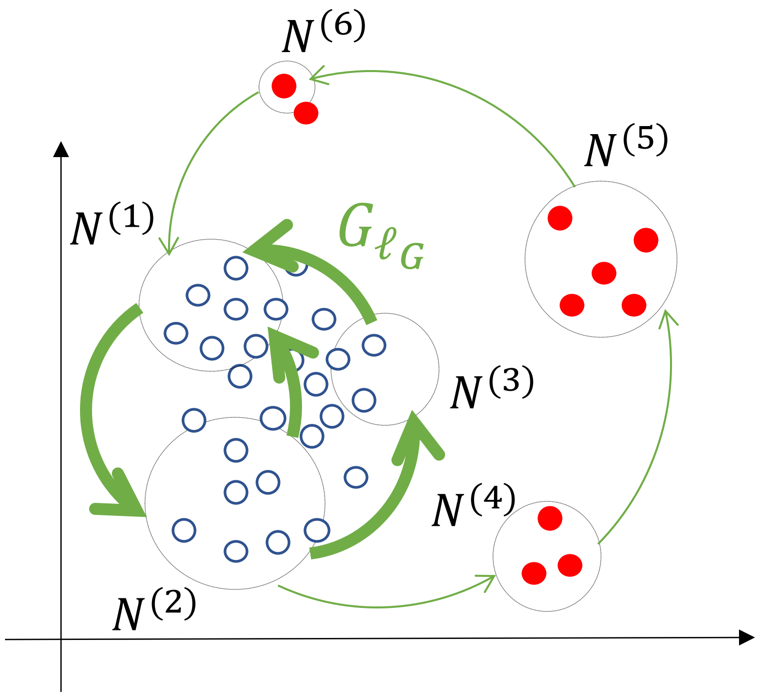

Each node is an ensemble of similar subsequences.

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

A subsequence $T_{i,\ell}$ (with $\ell > \ell_G$) is a path in G_{ℓ_G}.

Series2Graph: *From time series to a graph*

Graph G_{ℓ_G} [9]:

Given a data series T, and an input length ℓ_G, we build a graph $G_{\ell_G}(\mathcal{N}, \mathcal{E})$ for which:

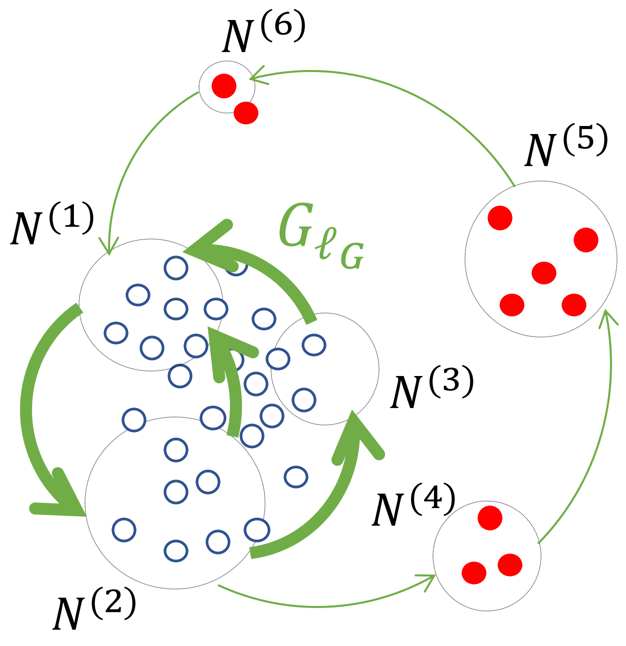

For a given subsequence $T_{i,\ell}$ and its corresponding path

$P_{th} = < N^{(i)}, N^{(i+1)}, \ldots, N^{(i+\ell)} >$,

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: *From time series to a graph*

Graph G_{ℓ_G} [9]:

Given a data series T, and an input length ℓ_G, we build a graph $G_{\ell_G}(\mathcal{N}, \mathcal{E})$ for which:

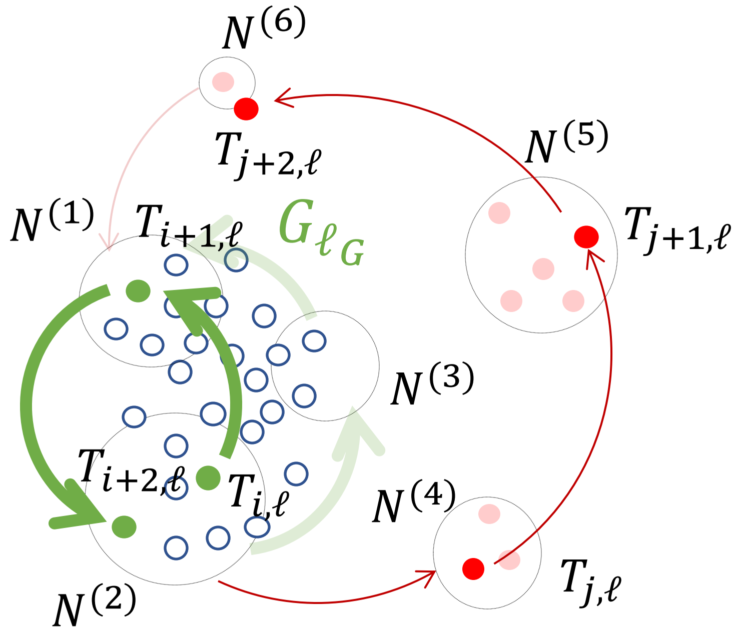

$$Norm\left(P_{th}(T_{j,\ell+2})\right) \ll Norm\left(P_{th}(T_{i,\ell+2})\right)$$

Series2Graph: *Computation Steps*

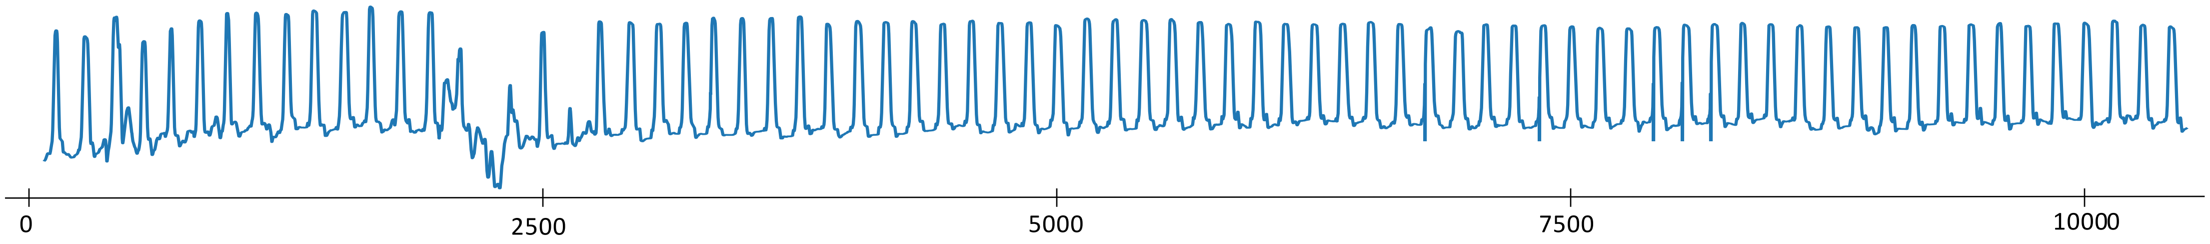

Series2Graph: *Computation Steps*

1 3 components of the *Principal Component Analysis* applied on all subsequences of T

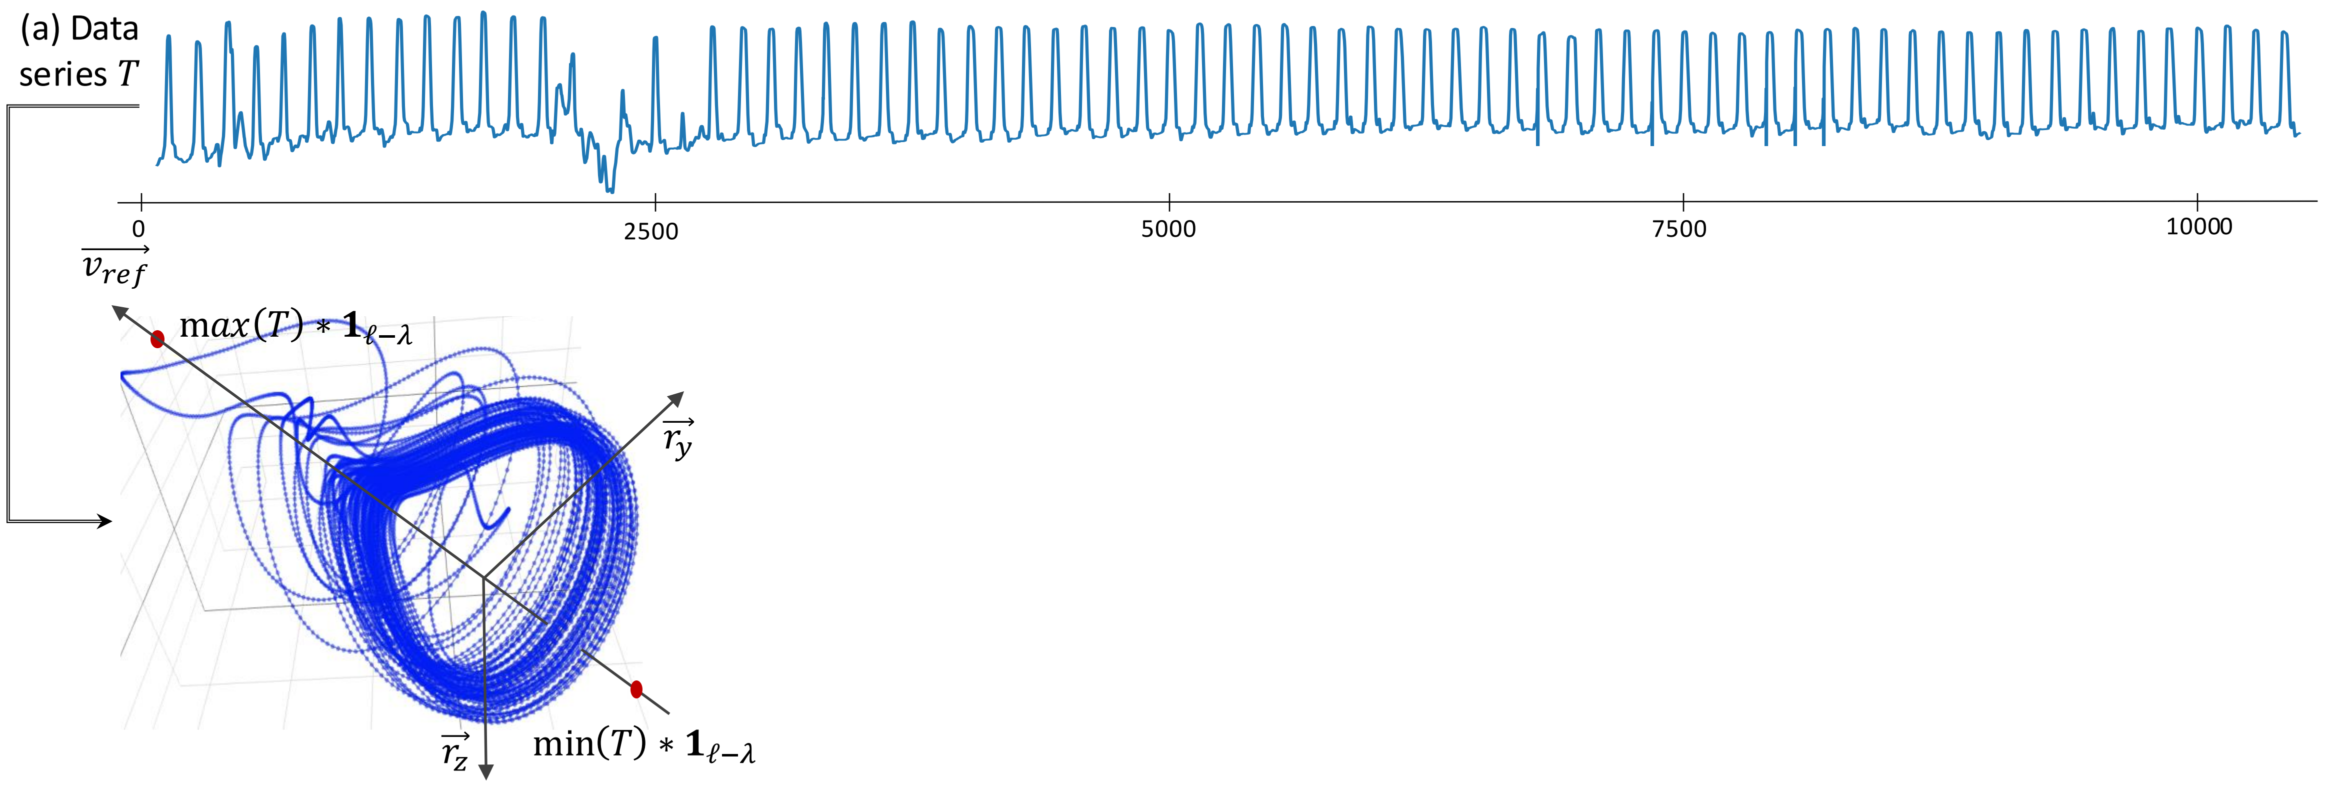

(b) First step: subsequence embedding

Series2Graph: *Computation Steps*

1 3 components of the *Principal Component Analysis* applied on all subsequences of T

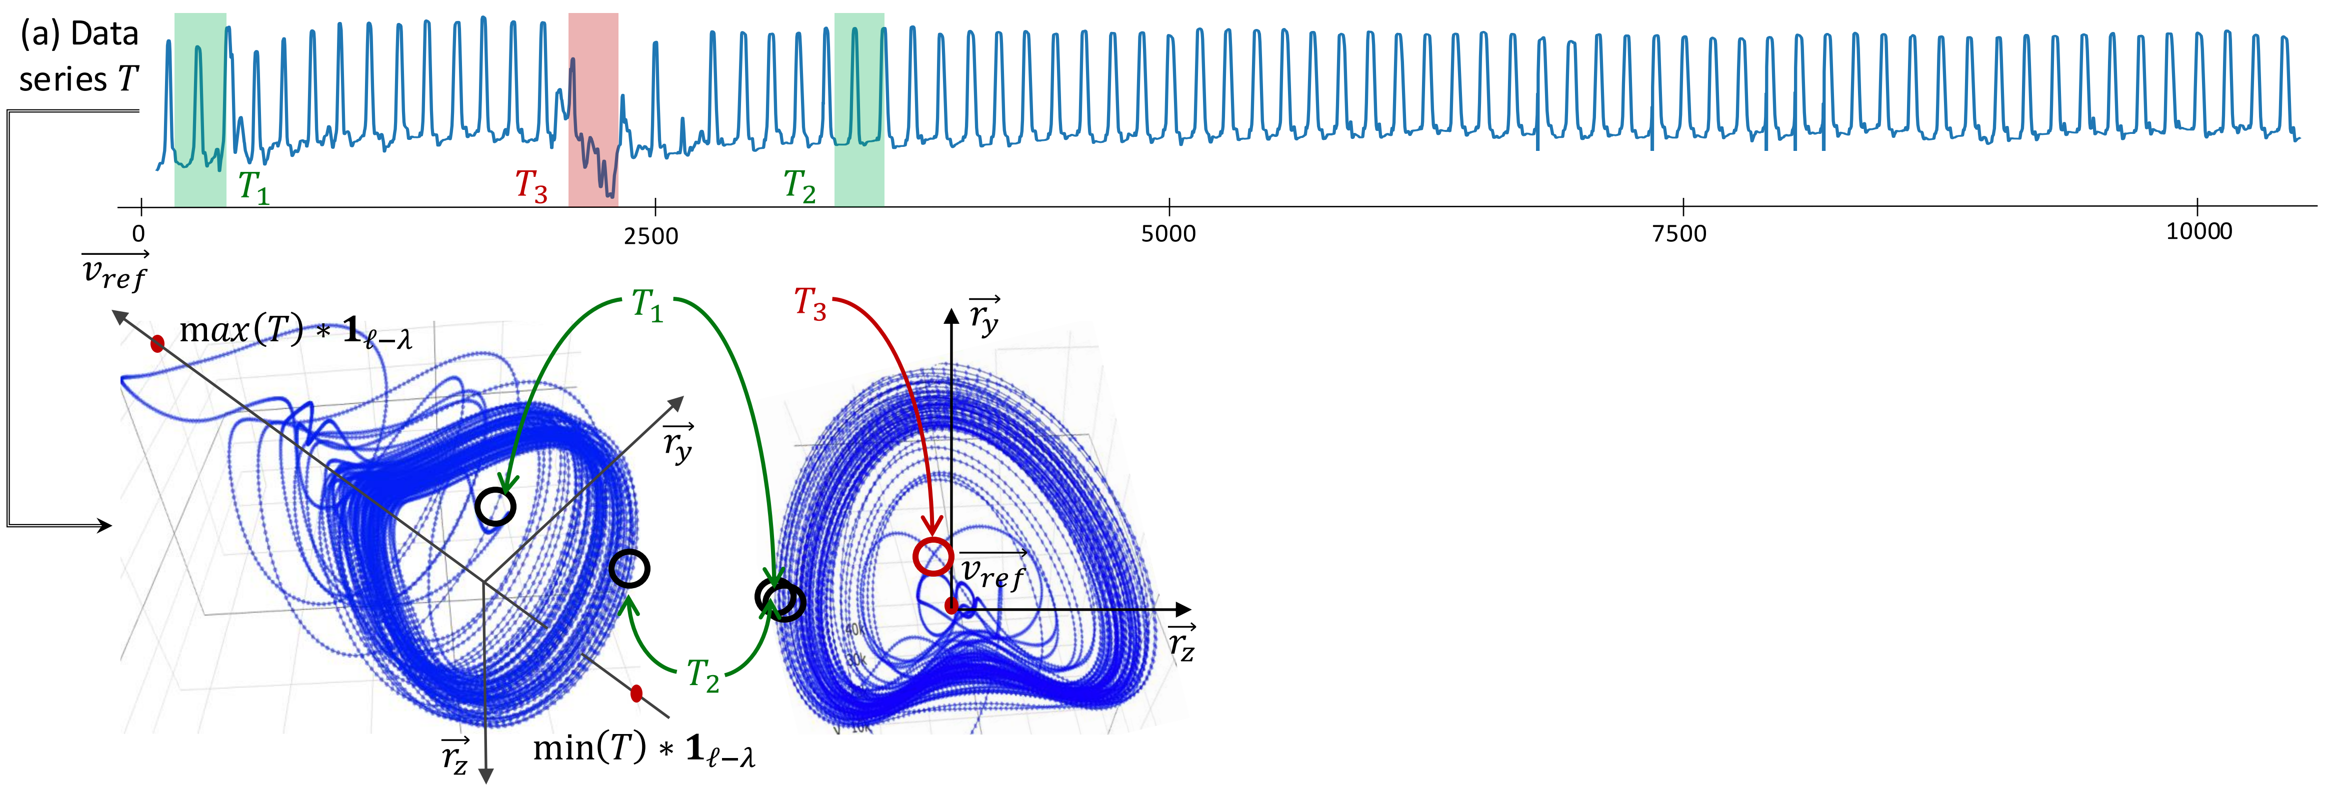

(b) First step: subsequence embedding

Series2Graph: *Computation Steps*

1 3 components of the *Principal Component Analysis* applied on all subsequences of T

2 Gaussian density estimation on each radius (among a fixed number of radius)

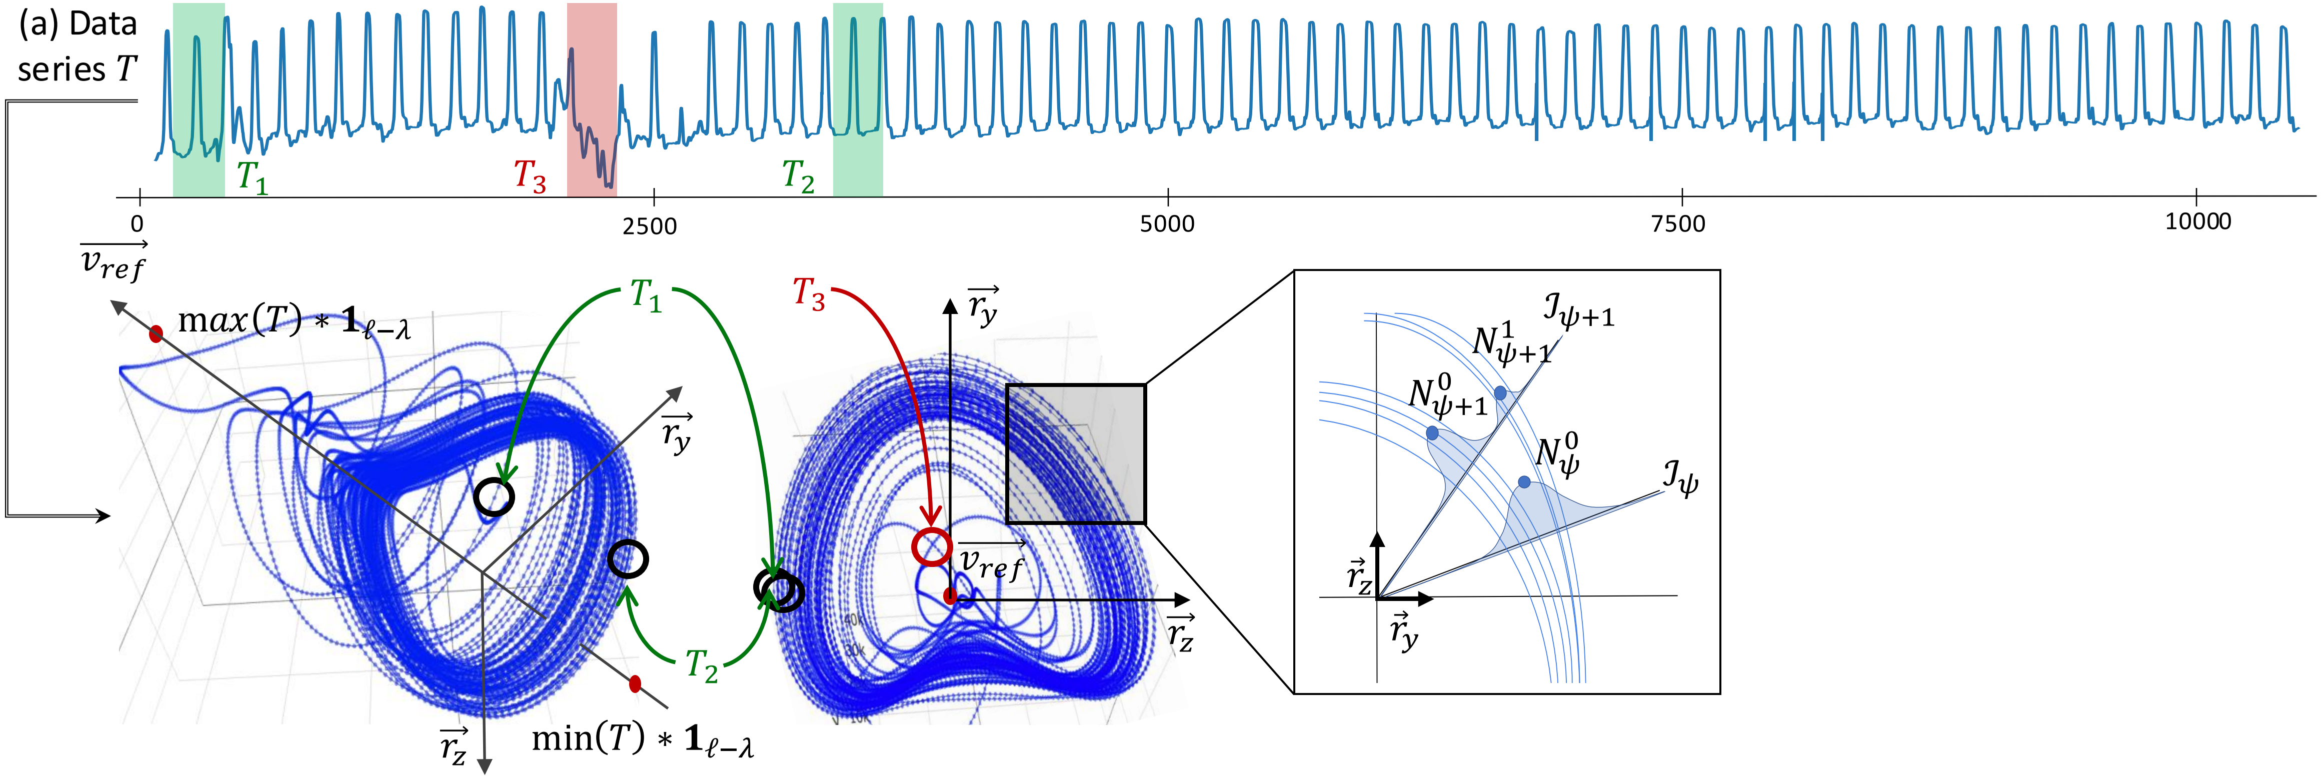

(b) First step: subsequence embedding

(c) Second step: node creation

Series2Graph: *Computation Steps*

1. 3 components of the *Principal Component Analysis* applied on all subsequences of T
2. Gaussian density estimation on each radius (among a fixed number of radius)
3. Assign each subsequence to a node and set an edge for each transition between nodes

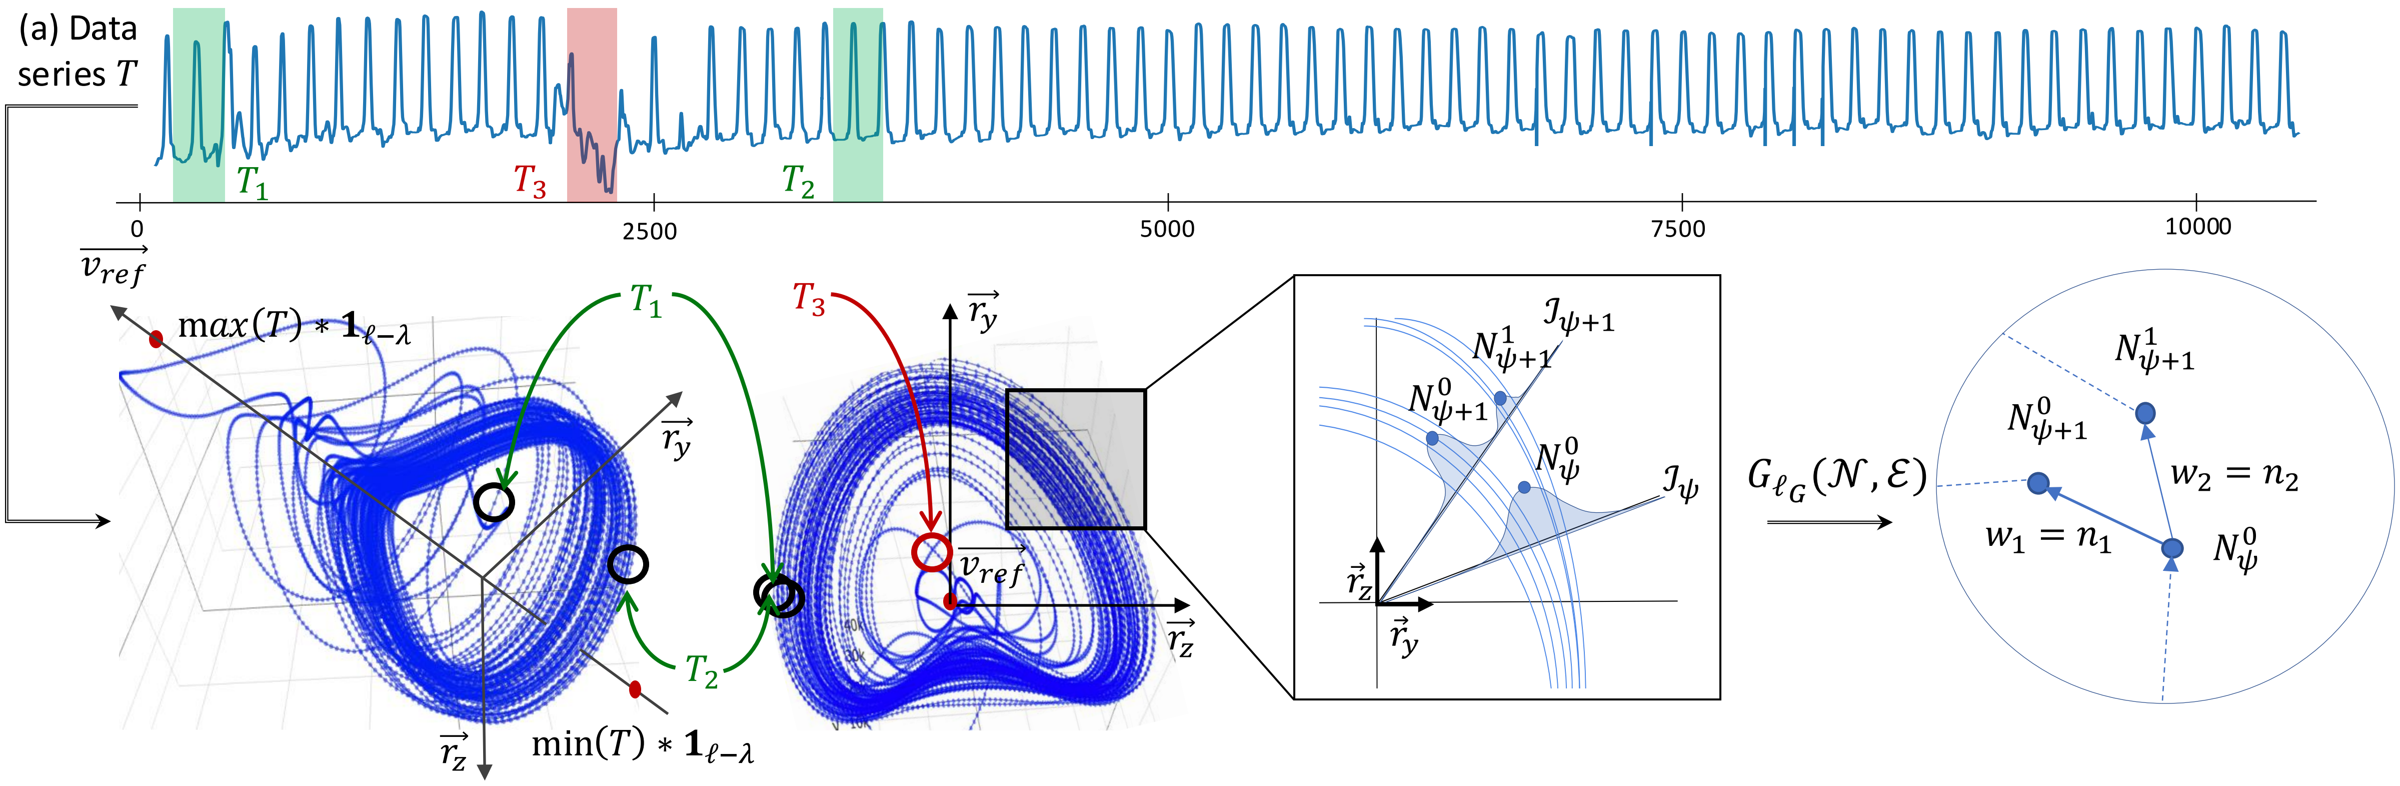

(b) First step: subsequence embedding

(c) Second step: node creation

(d) Third step: edge creation

Series2Graph: *An Example*

Snippet of SED time series

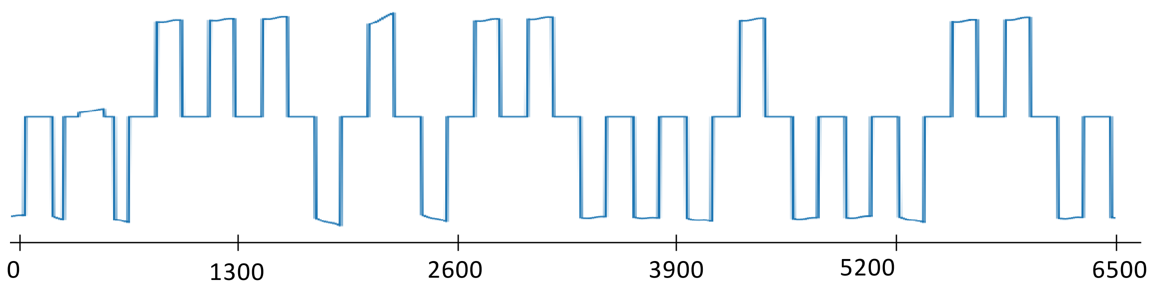

Series2Graph: *An Example*

Snippet of SED time series

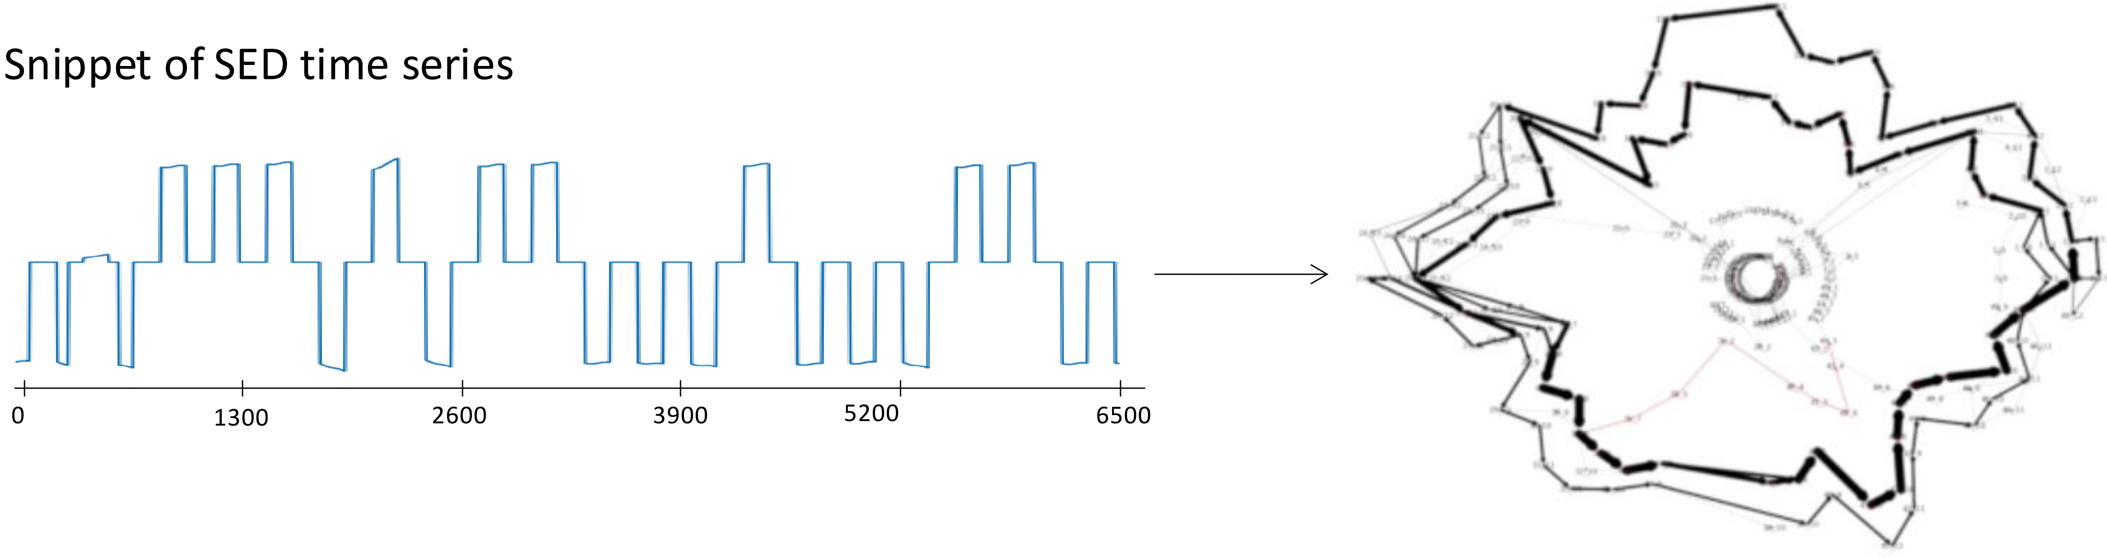

Series2Graph: *An Example*

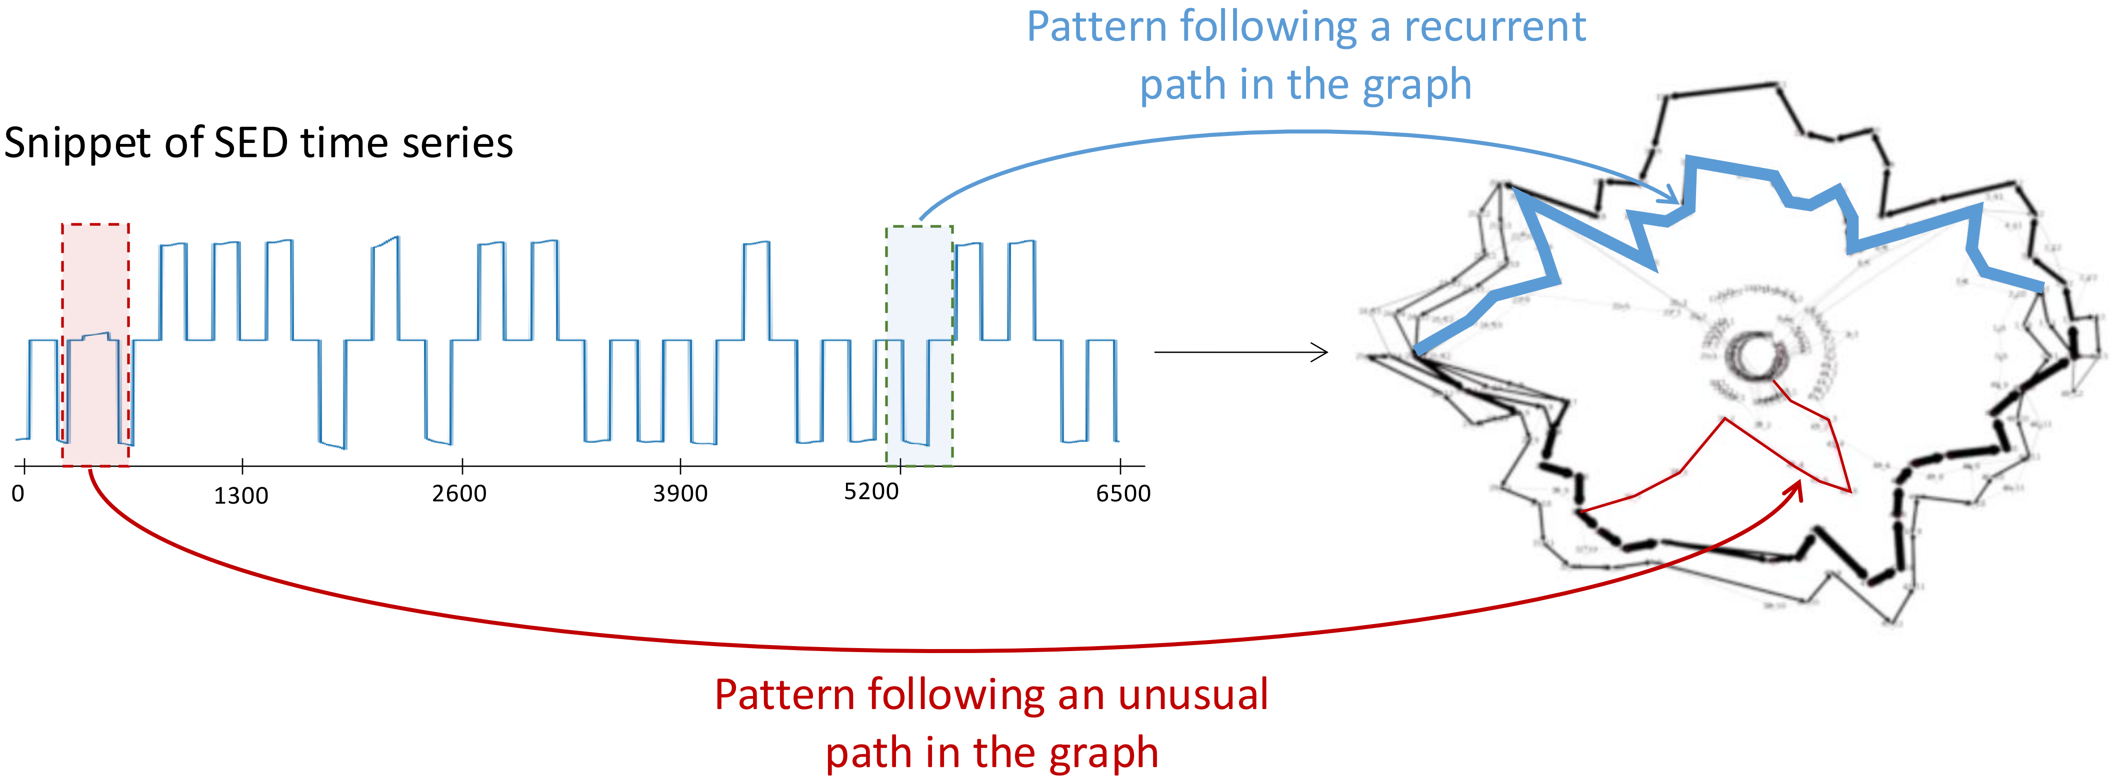

Series2Graph:
An interactive tool

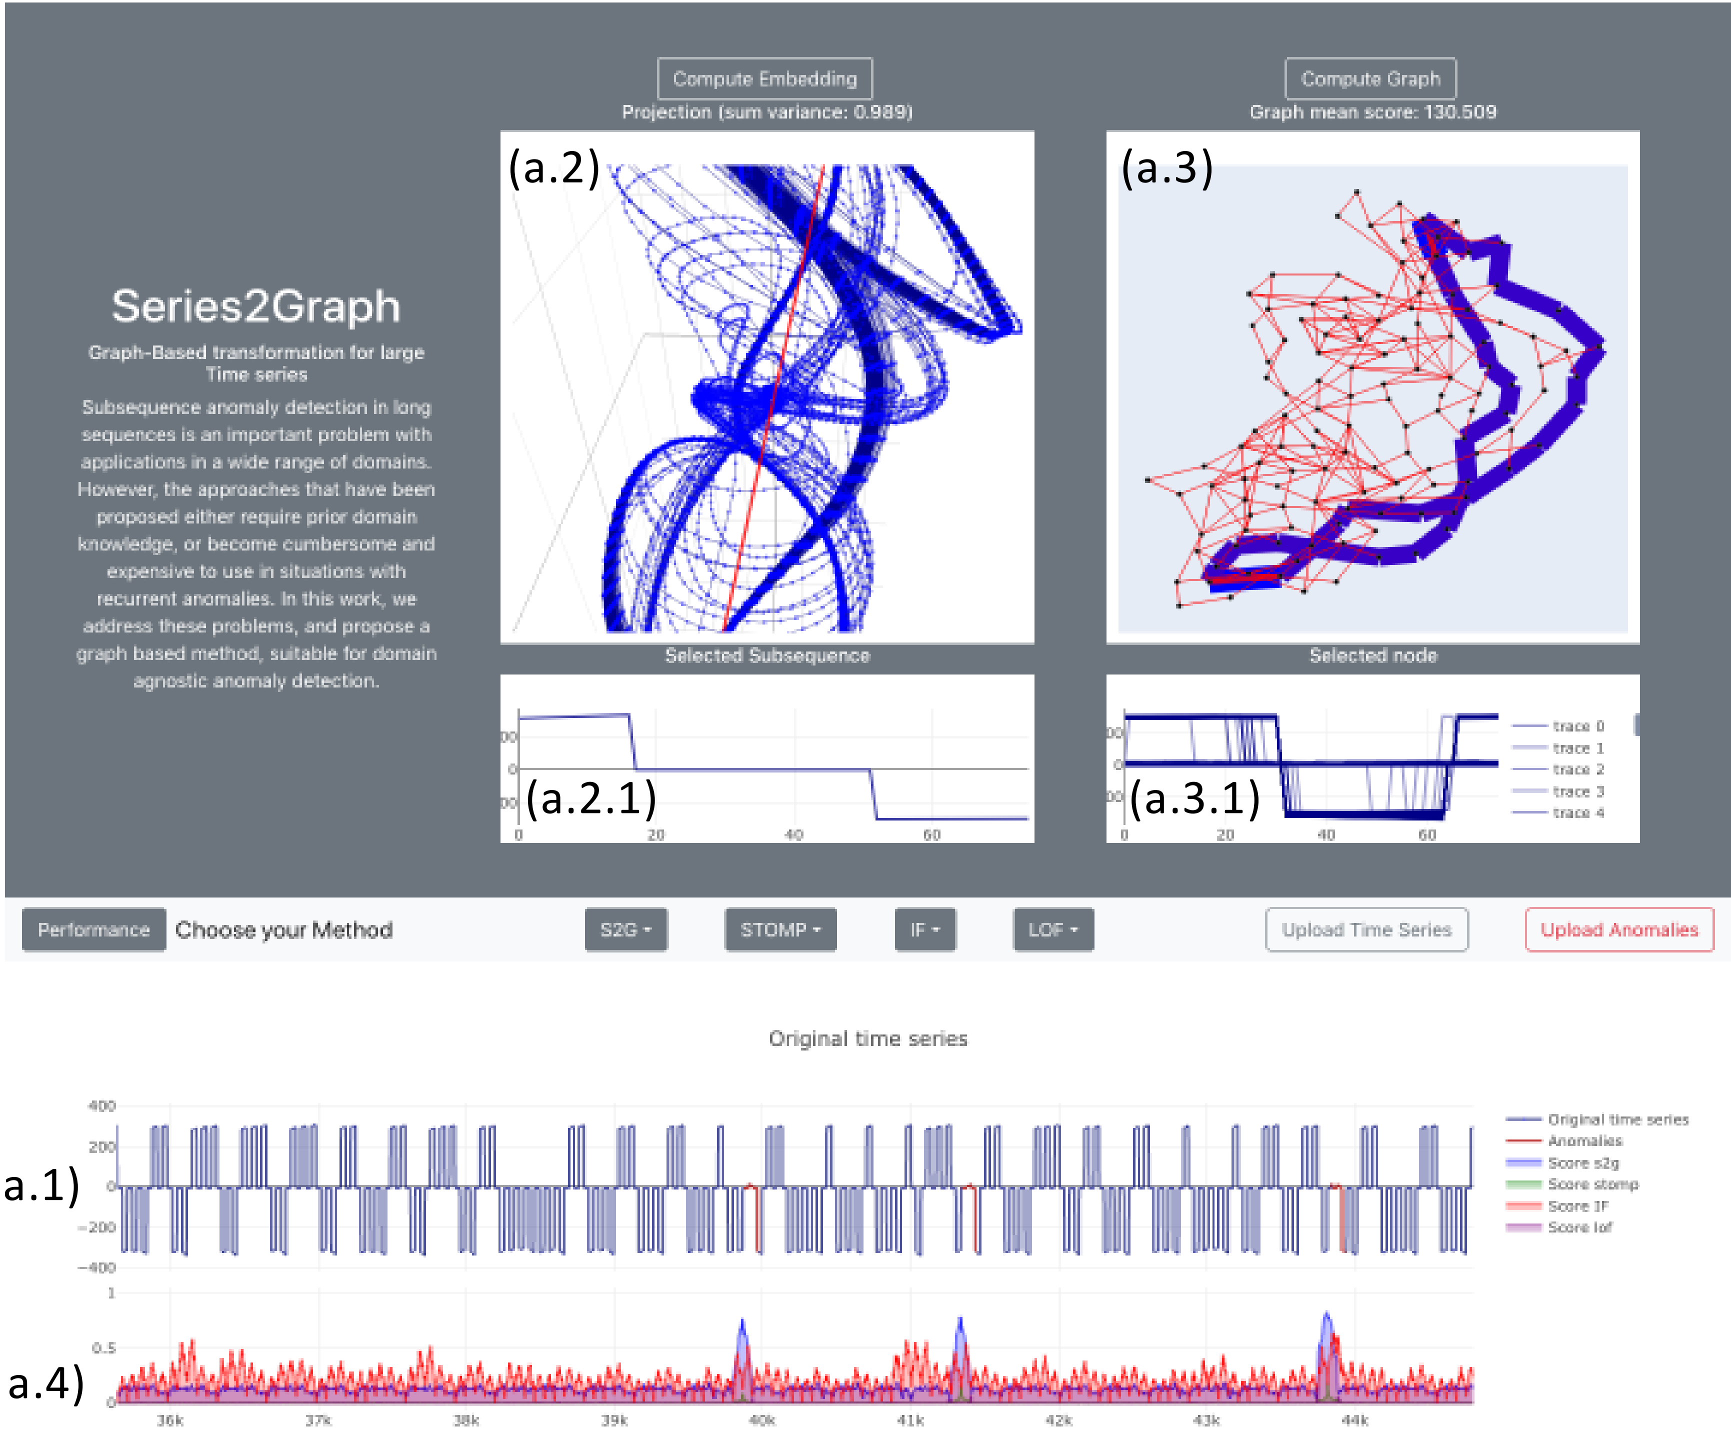

GraphAn: S2G User interface [10]

Series2Graph:
To conclude

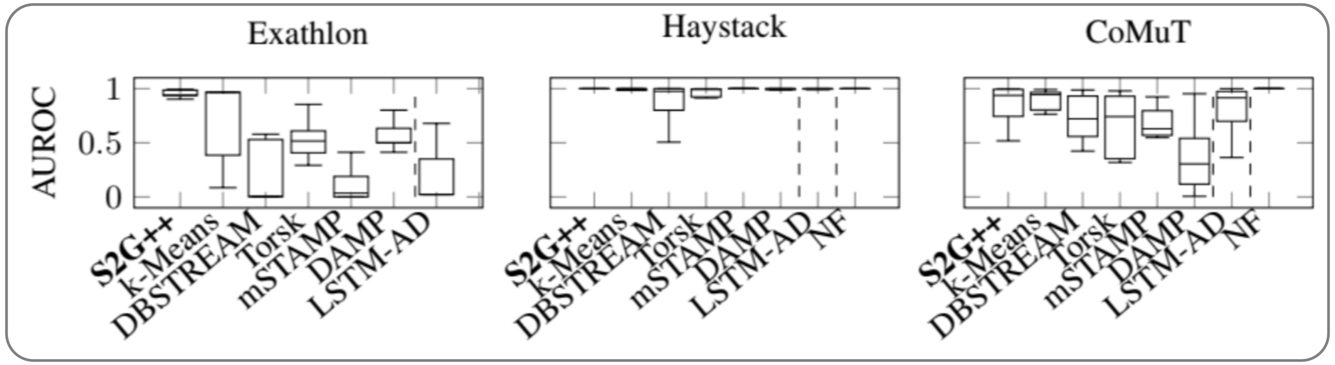

Series2Graph++: Multivariate extension of S2G [11]

In summary:

- We proposed a user interface to explore the resulting graph [10]
- Series2Graph extensions have been proposed [11,12]

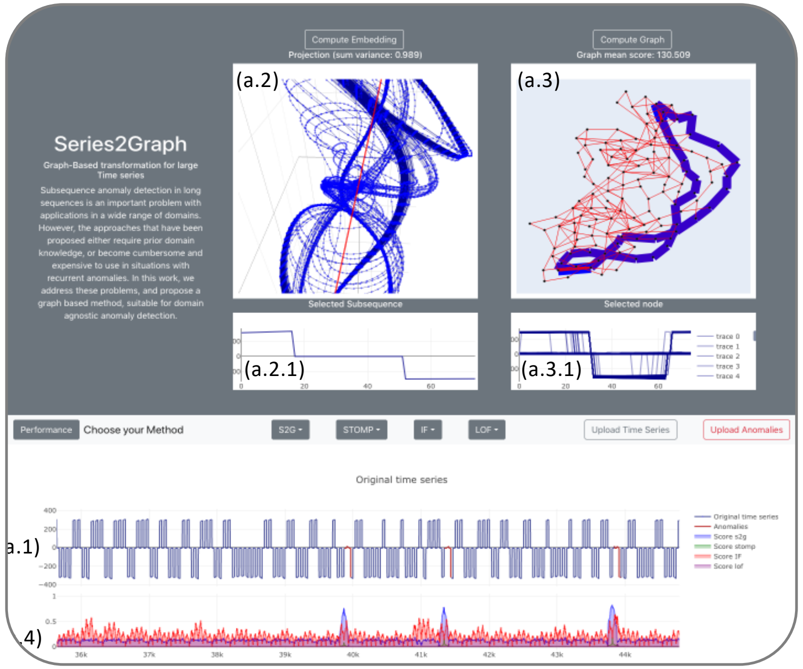

GraphAn: S2G User interface [10]

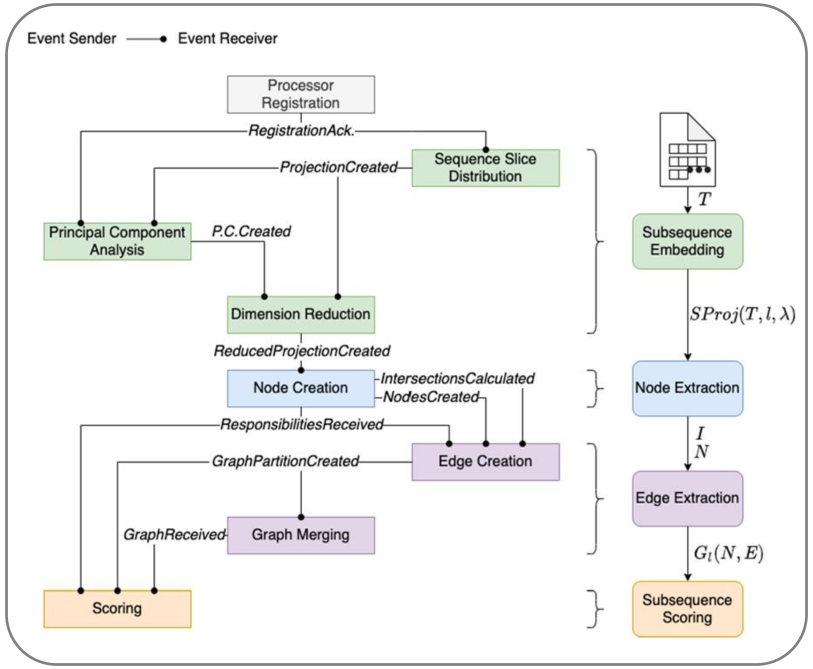

DADS: Distributed version of S2G [12]

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?

Time series Database

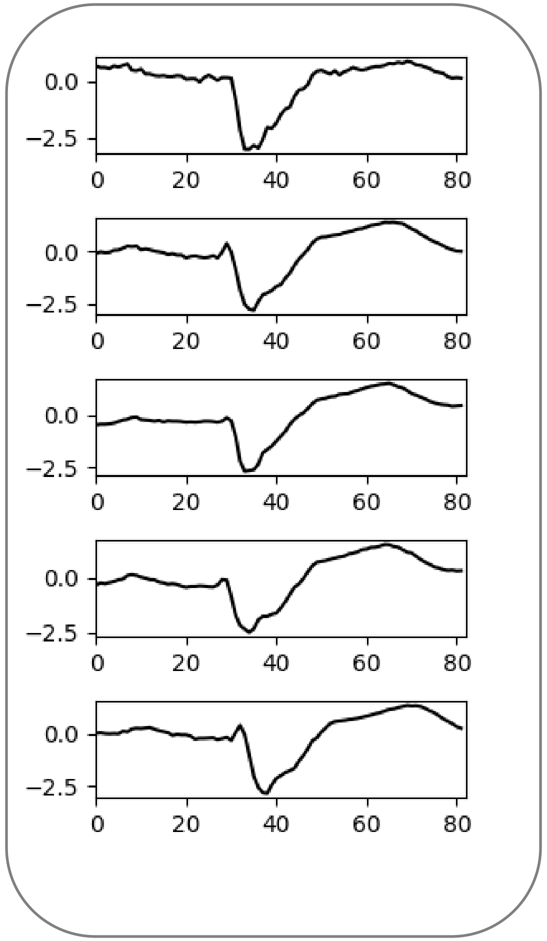

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?

Time series Database

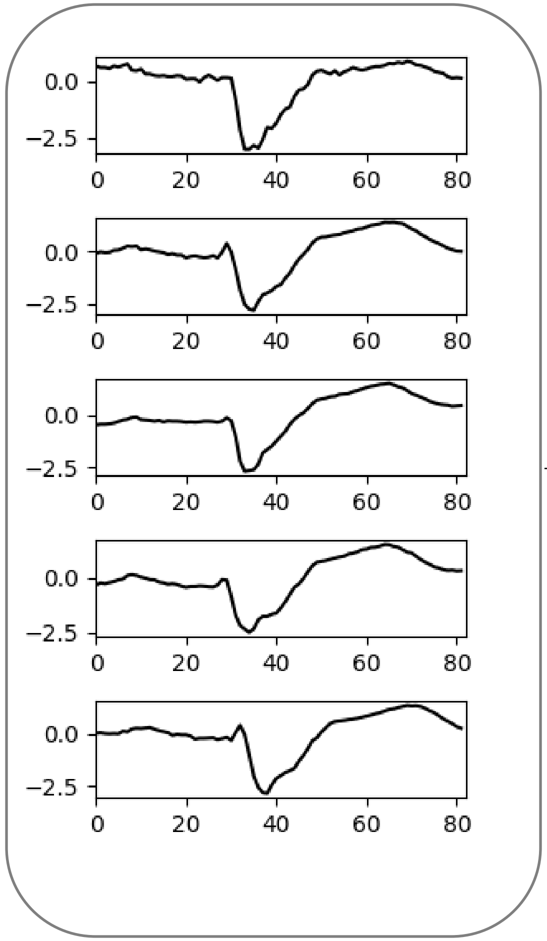

Graph embedding per time series

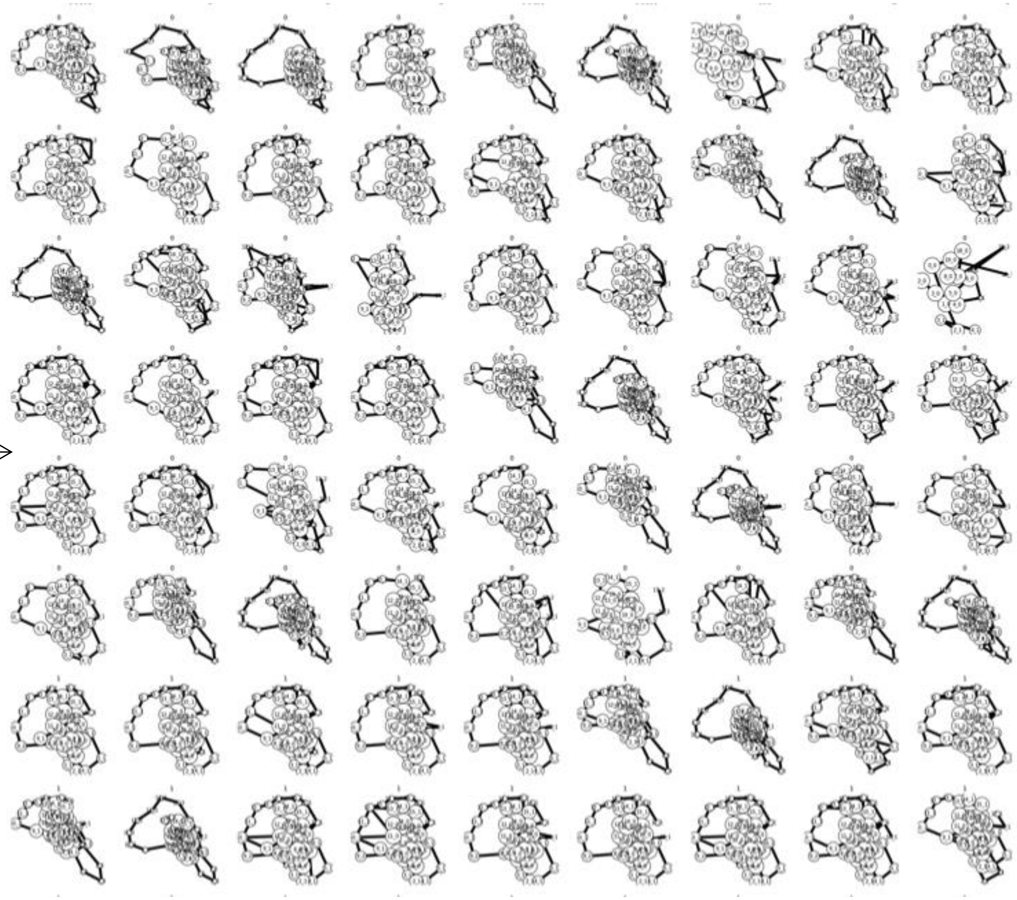

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?
- Is a unique graph meaningful for a set of time series?

Time series Database

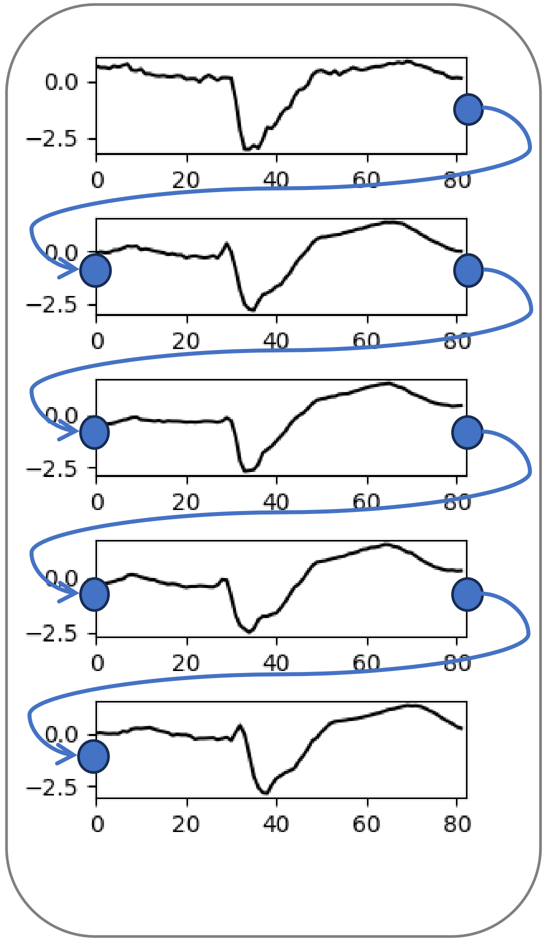

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?

- Is a unique graph meaningful for a set of time series?

Time series Database

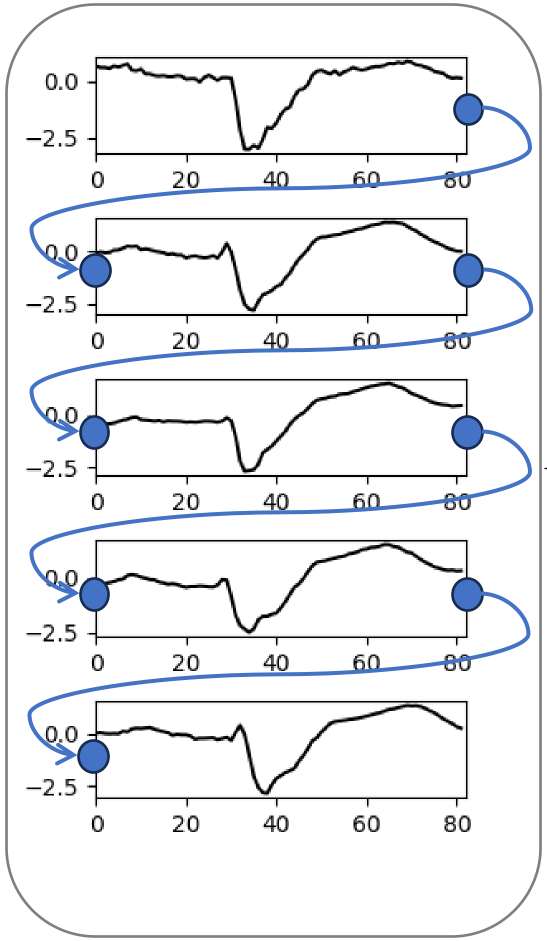

Graph embedding of the database

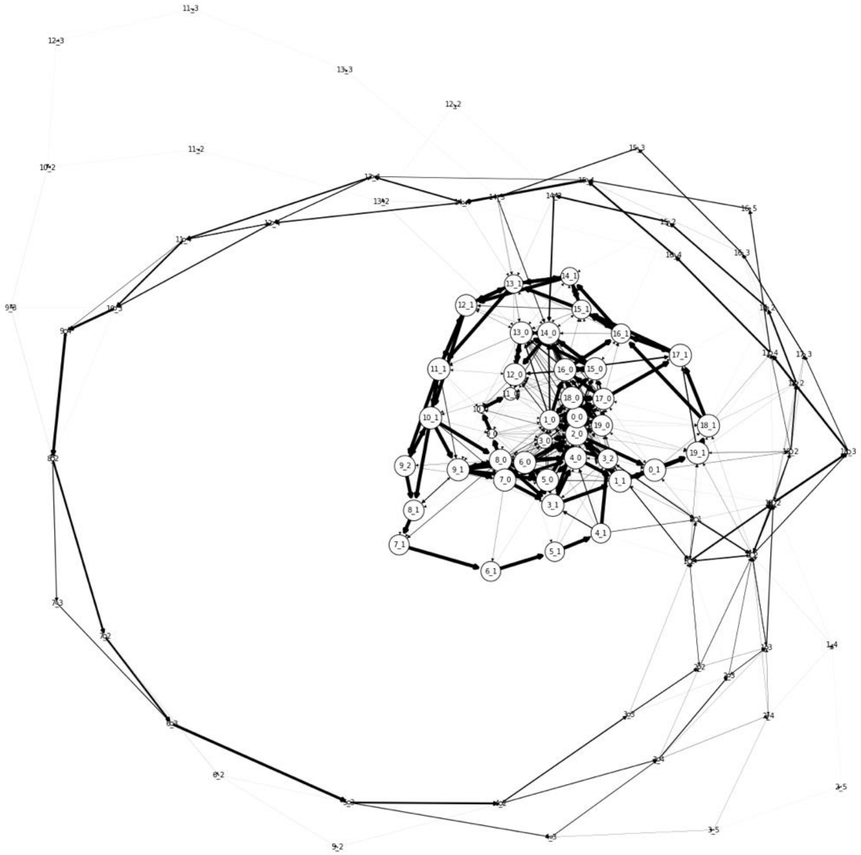

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?
- Is a unique graph meaningful for a set of time series?

Time series Database

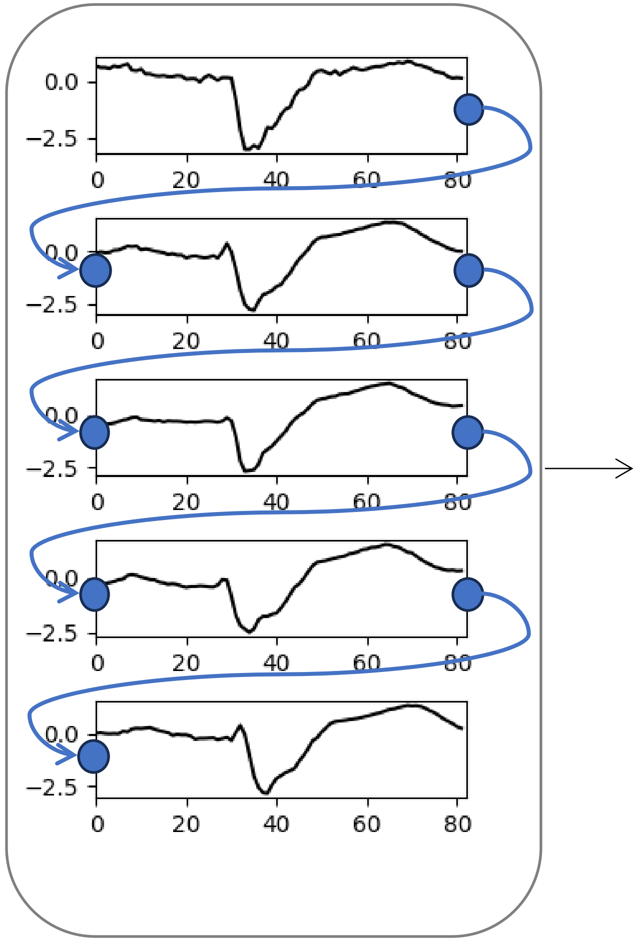

Graph embedding of the database

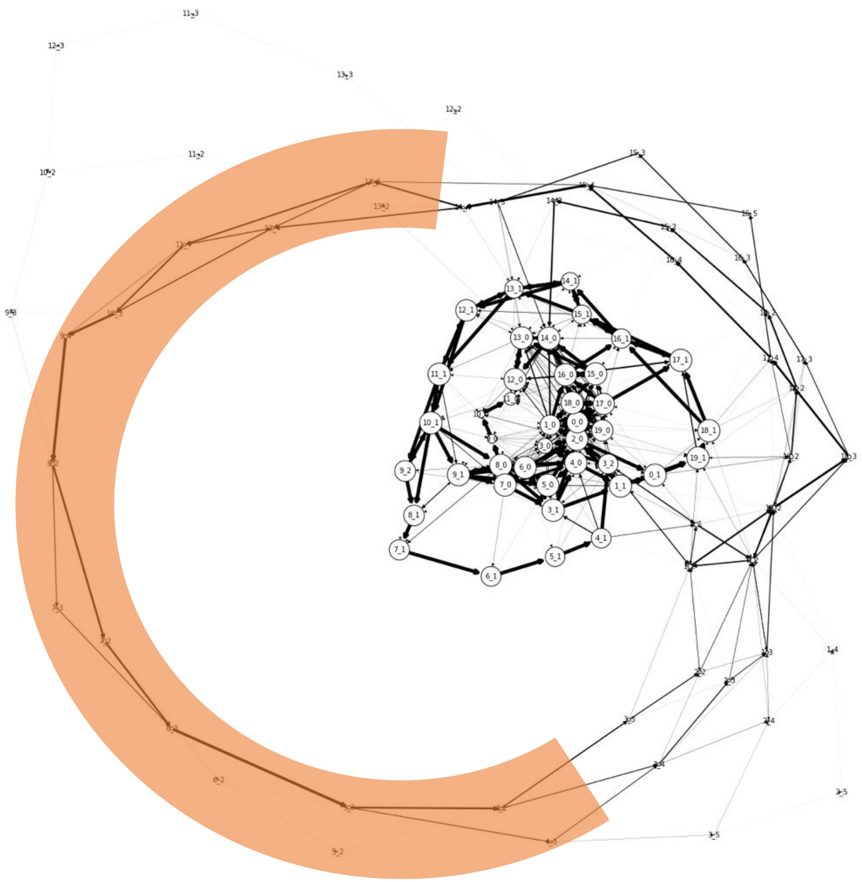

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?

- Is a unique graph meaningful for a set of time series?

Time series Database

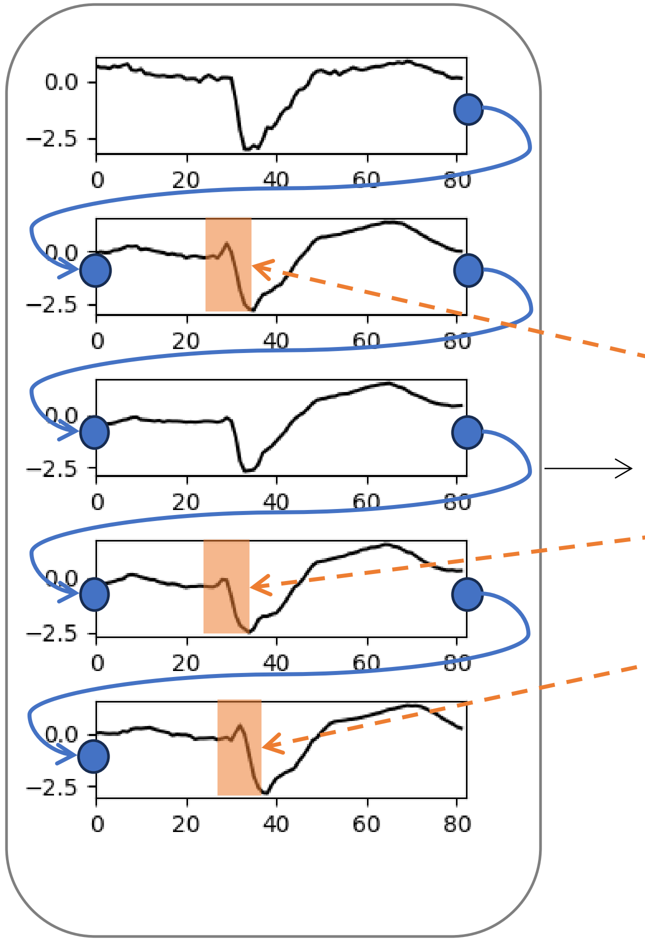

Graph embedding of the database

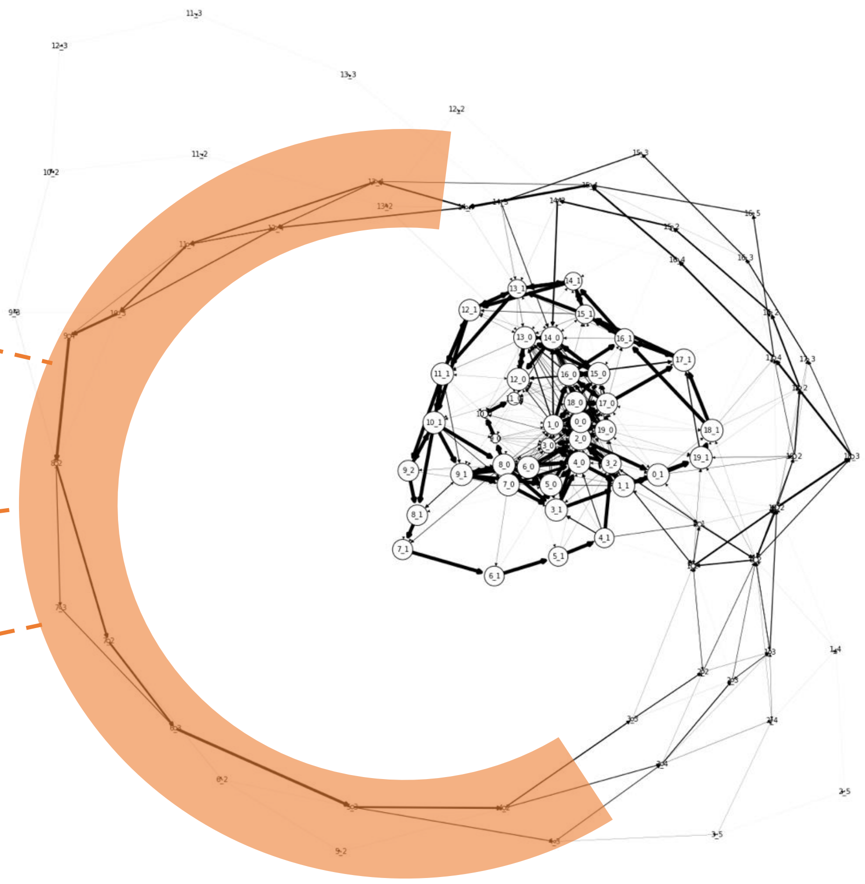

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?

- Is a unique graph meaningful for a set of time series?

Trace dataset (Class 2) [13]

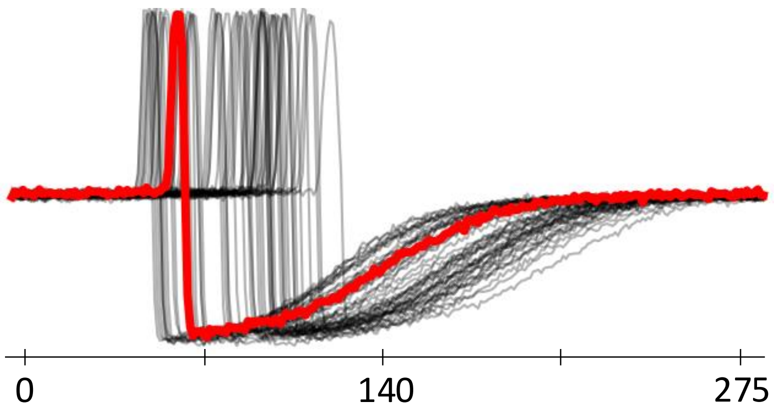

Trace dataset (Class 1) [13]

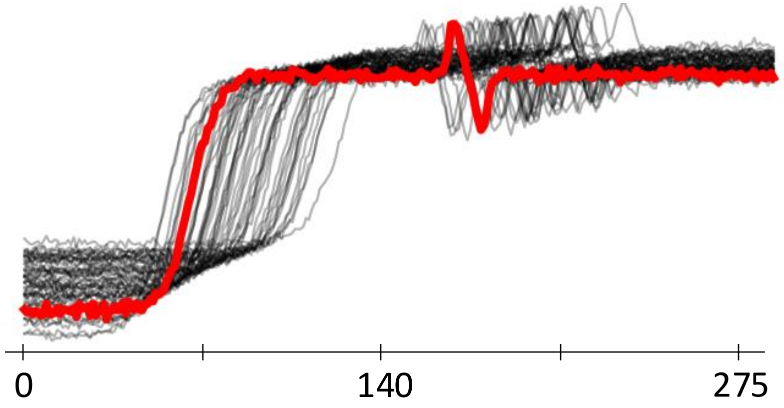

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?

- Is a unique graph meaningful for a set of time series?

Trace dataset (Class 2) [13]

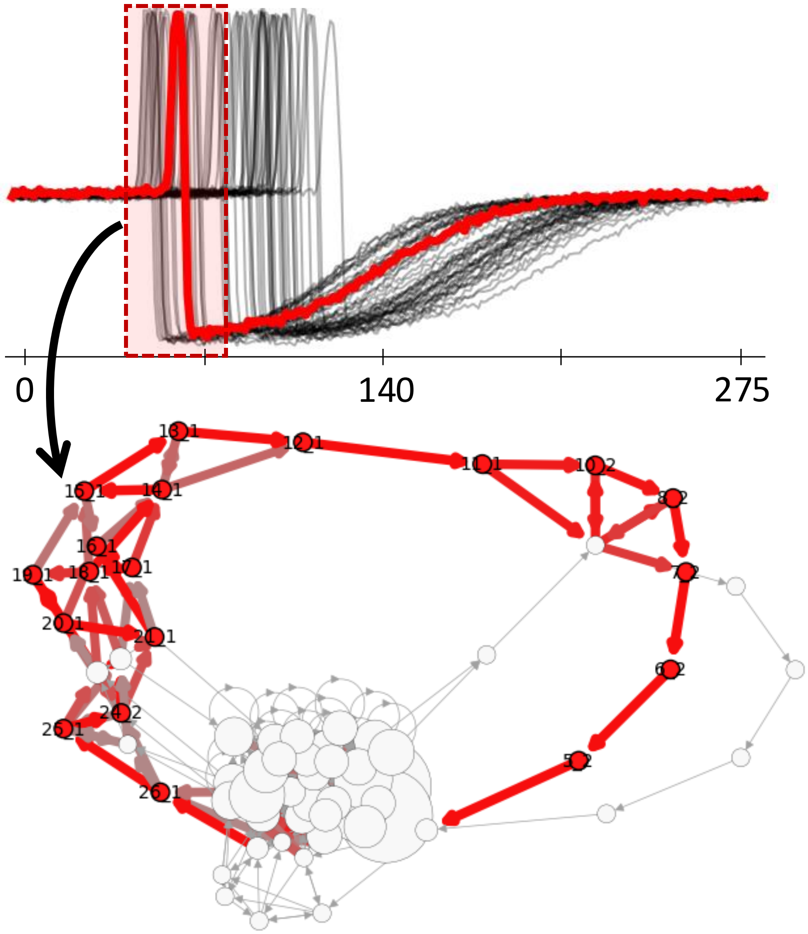

Trace dataset (Class 1) [13]

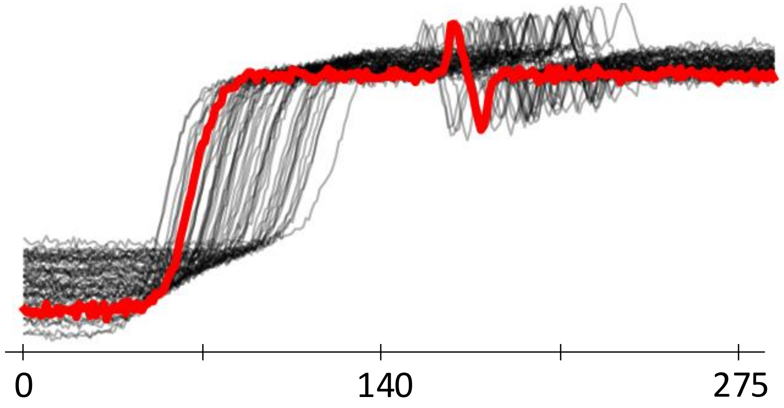

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?

- Is a unique graph meaningful for a set of time series?

Trace dataset (Class 2) [13]

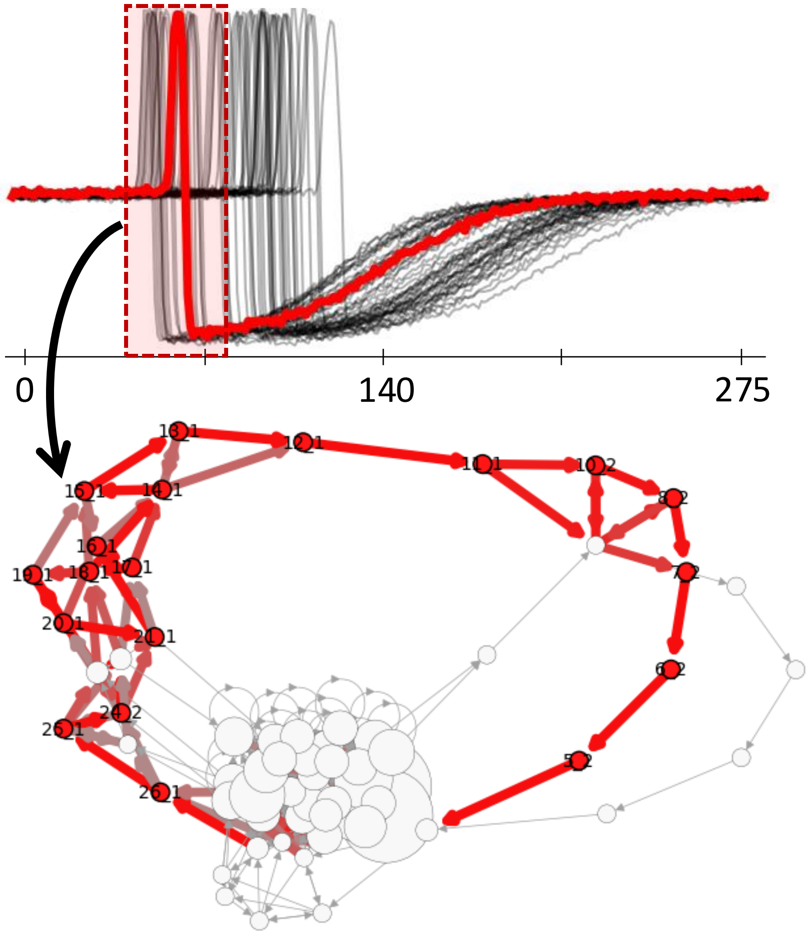

Trace dataset (Class 1) [13]

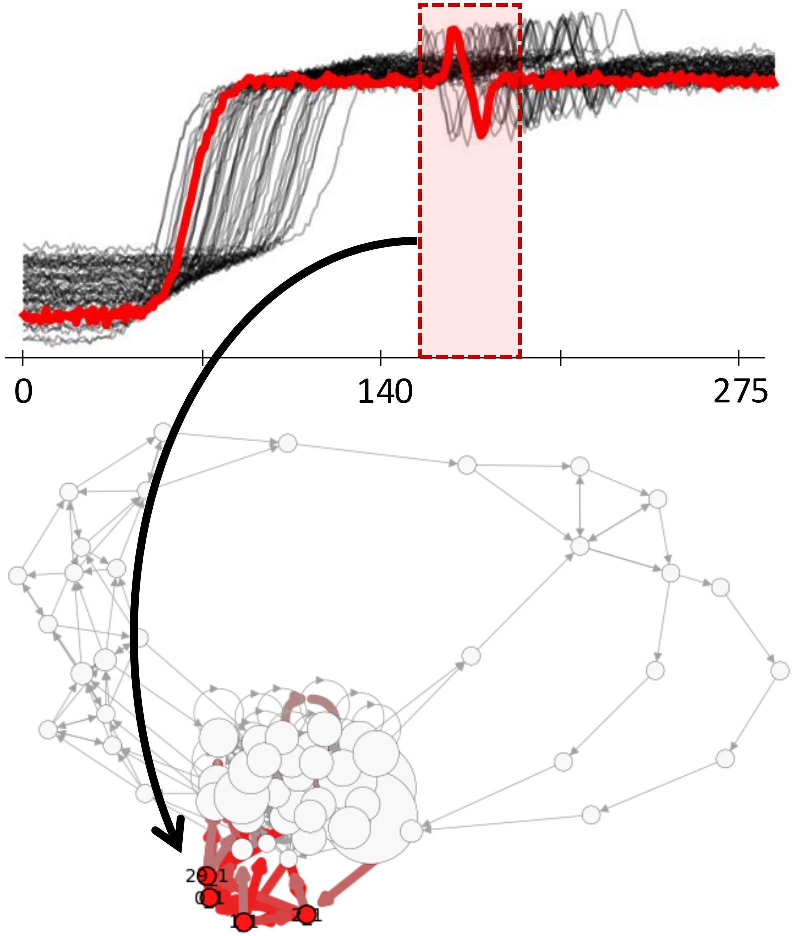

Series2Graph: *What next?*

Several research directions

- Can the graph structure of Series2Graph help identify different time series types?
- Is a unique graph meaningful for a set of time series?
- Can we use this graph to perform multiple analytics?

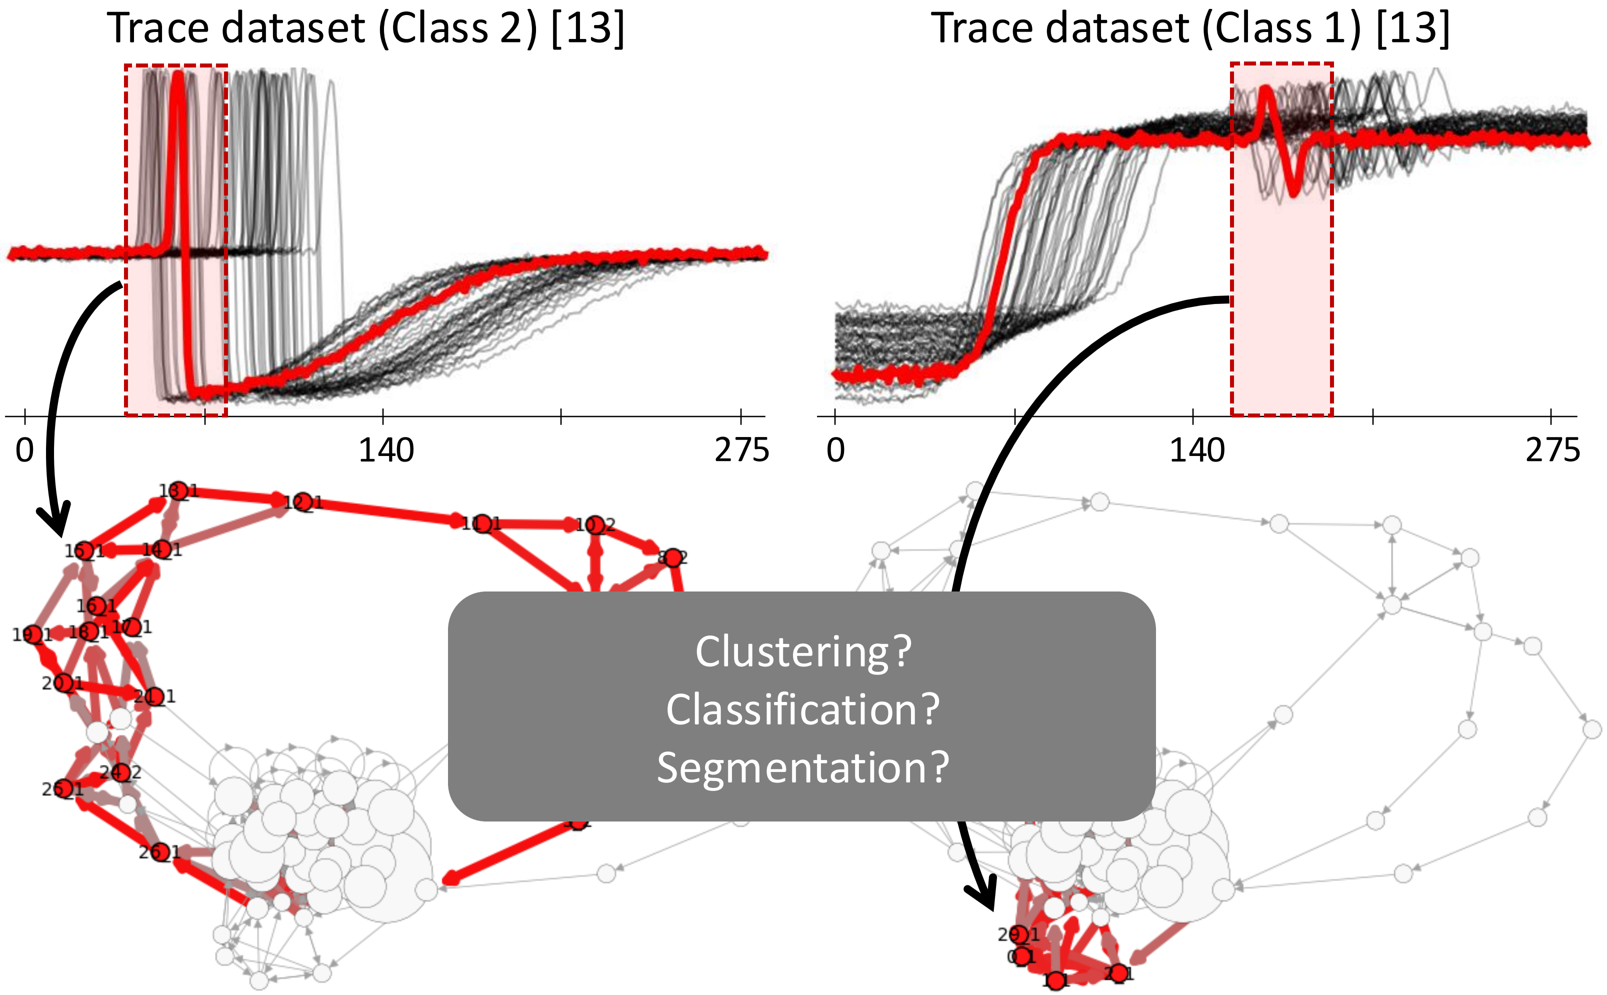

Series2Graph: *What next?*

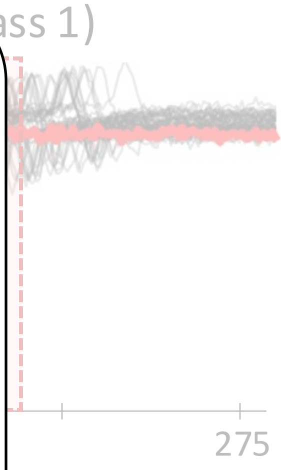

Series2Graph:

Graph-based Subsequence Anomaly Detection in Time Series

Paul Boniol and Themis Palpanas.

Paper
(VLDB 2020)

https://www.vldb.org/pvldb/vol13/p1821-boniol.pdf

GitHub Repositories

TSB-UAD	DADS	S2Gpp
TheDatumOrg/TSB-UAD	HPI-Information-Systems/DADS	HPI-Information-Systems/S2Gpp

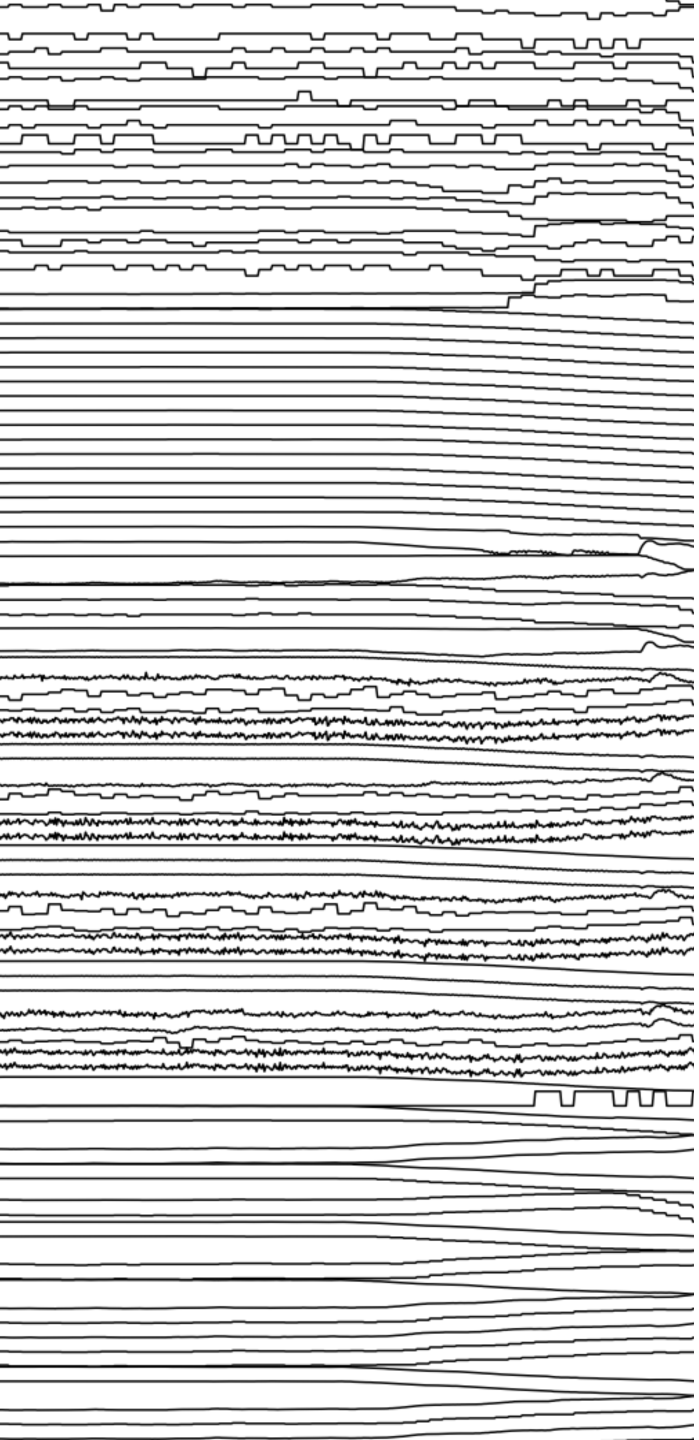

IV. Automated Solutions

How to pick automatically the best method?

Automated Solution: *Background*

Motivation:

- No one-size-fits-all model: How can we *automatically* identify the best anomaly detector given a time series?

Detection accuracy (VUS-PR) for 6 anomaly detectors across different datasets in TSB-UAD [14]

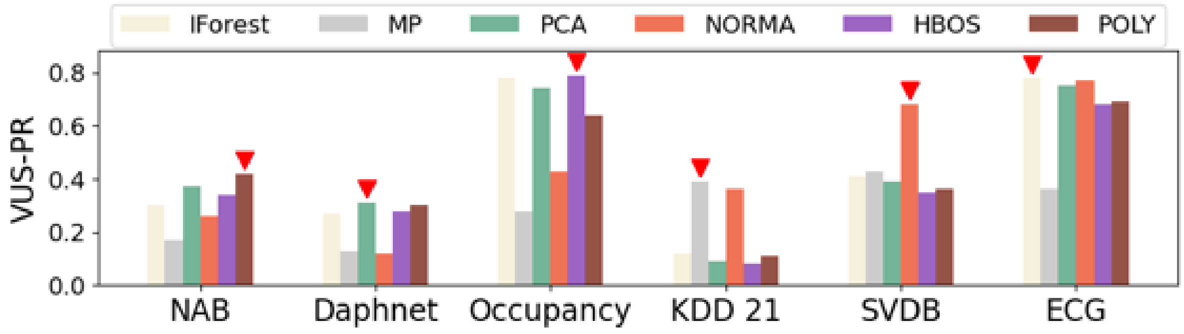

Automated Solution: *Taxonomy*

(a) Model Selection:

Selecting the best anomaly detector from a predefined candidate model set.

- *(a.1) Internal Evaluation*
- *(a.2) Meta-learning-based*

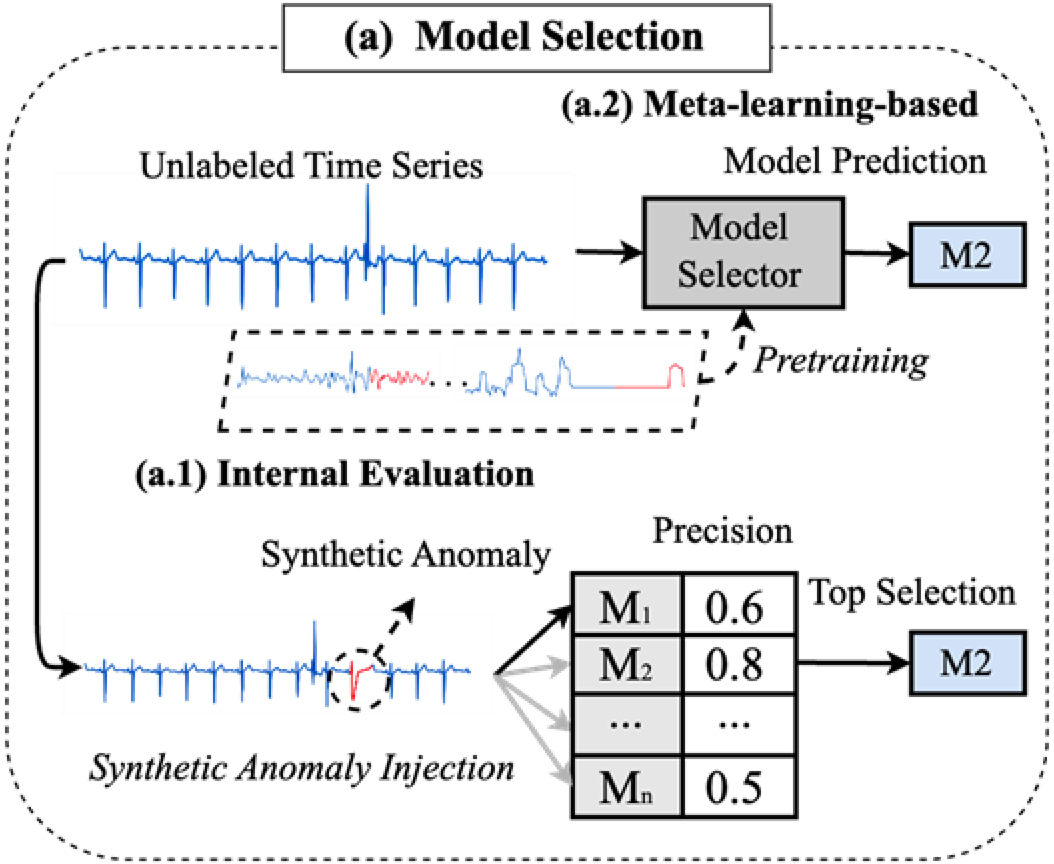

Automated Solution: *Taxonomy*

(a) Model Selection:

Selecting the best anomaly detector from a predefined candidate model set.

- *(a.1) Internal Evaluation*
- *(a.2) Meta-learning-based*

(b) Model Generation:

Creating an entirely new model for the given time series based on the candidate mode set

- *(b.1) Ensembling-based*
- *(b.2) Pseudo-label-based*

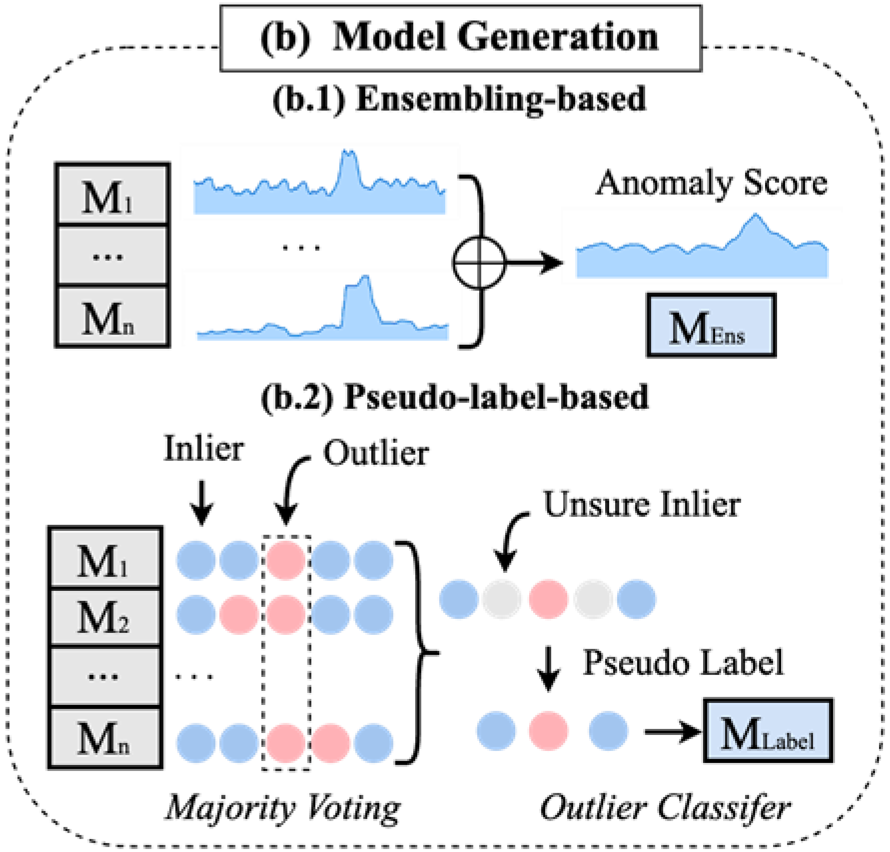

Automated Solution: *Internal Evaluation*

Definition: Evaluate the effectiveness of a model without any reliance on external information

- Stand-alone: Clustering Quality, EM&MV, Synthetic anomaly injection

- Collective: Model Centrality, Rank Aggregation

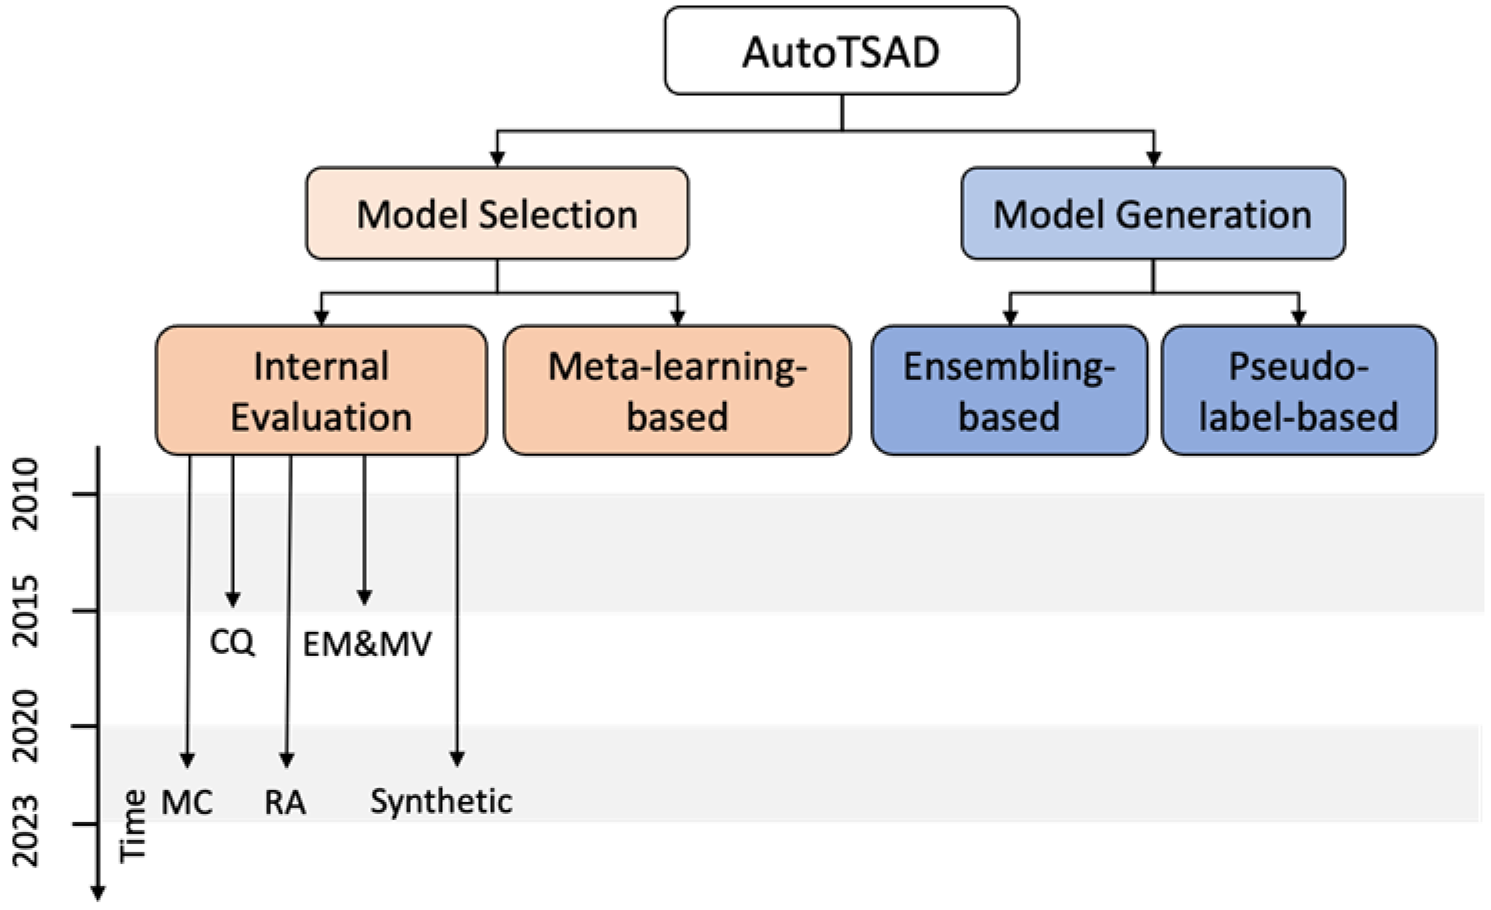

Automated Solution: *Internal Evaluation*

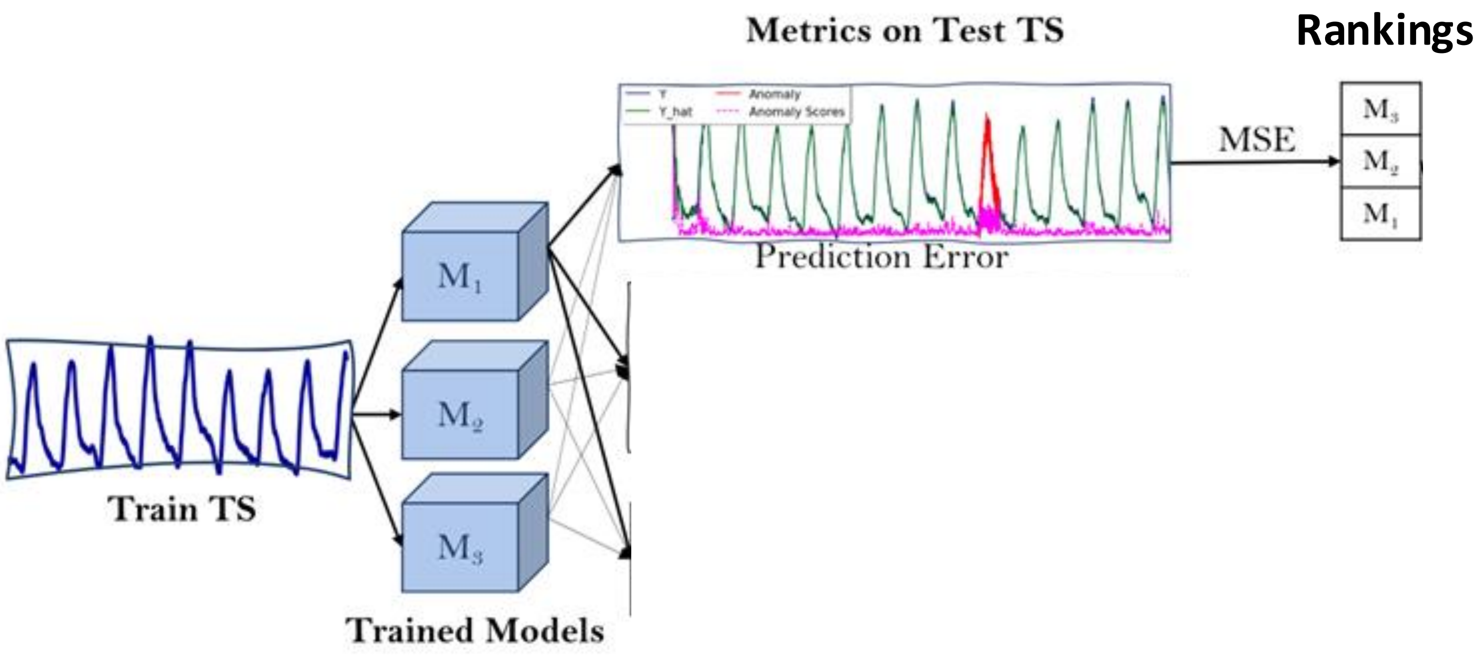

Image from [15]: Internal Evaluation workflow.

Automated Solution: *Internal Evaluation*

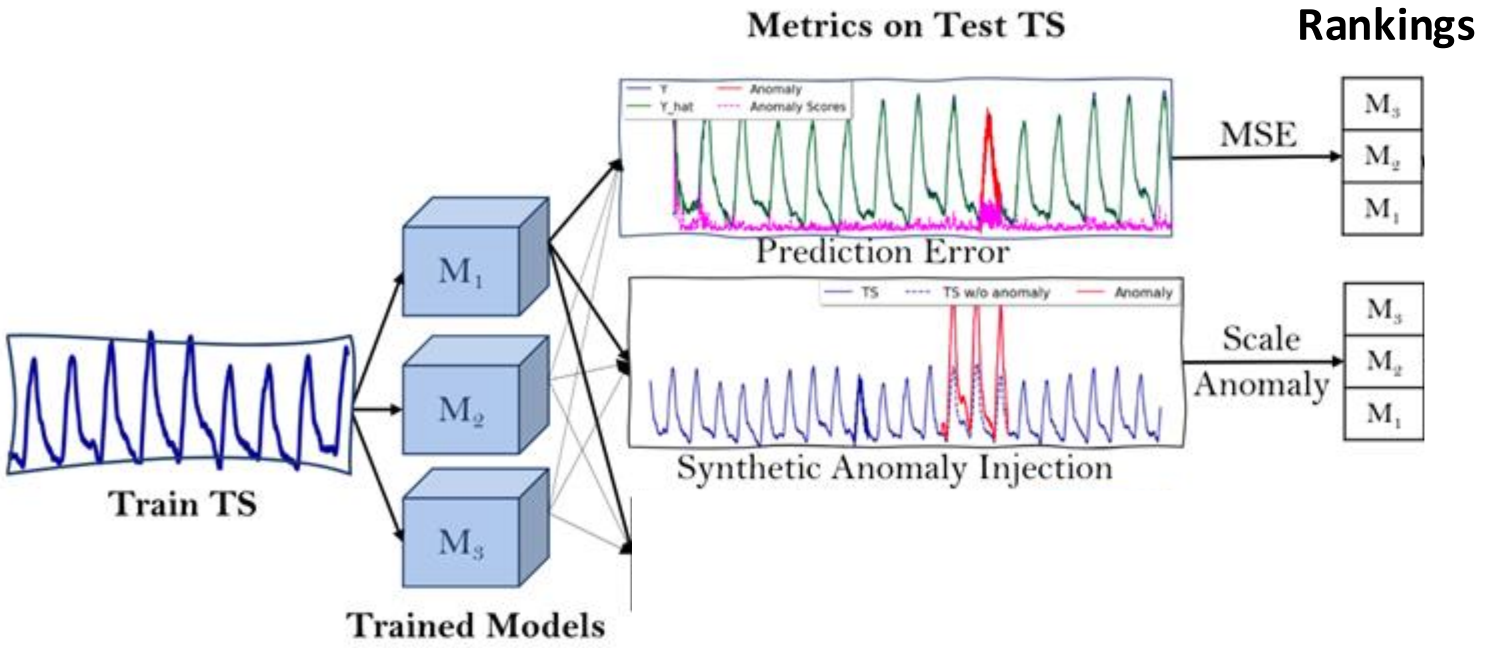

Image from [15]: Internal Evaluation workflow.

Automated Solution: *Internal Evaluation*

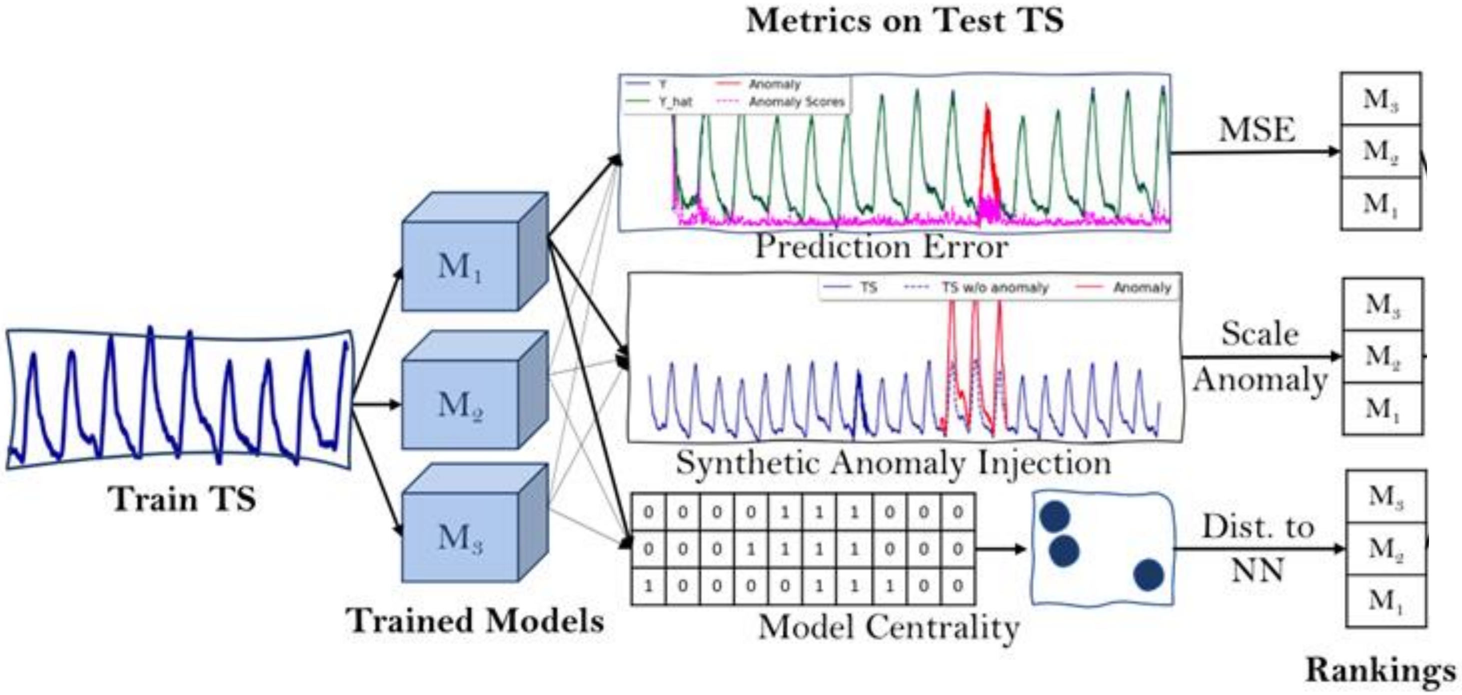

Image from [15]: Internal Evaluation workflow.

Automated Solution: *Internal Evaluation*

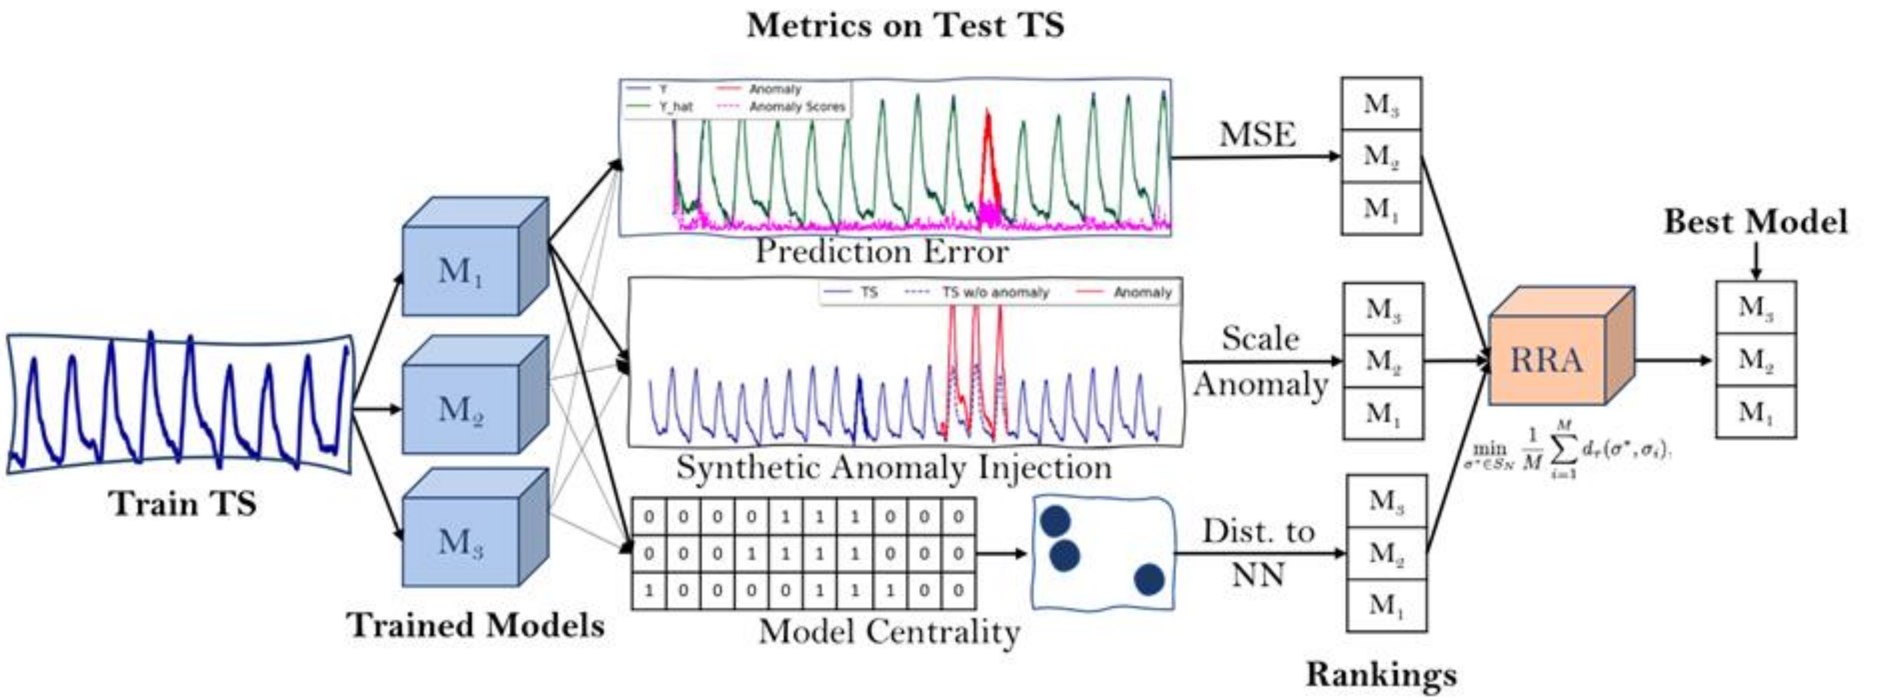

Image from [15]: Internal Evaluation workflow.

Automated Solution: *Ensembling-based*

Definition: Integrate predictions from the candidate model set

- Full: OE

- Selective: HITS, IOE

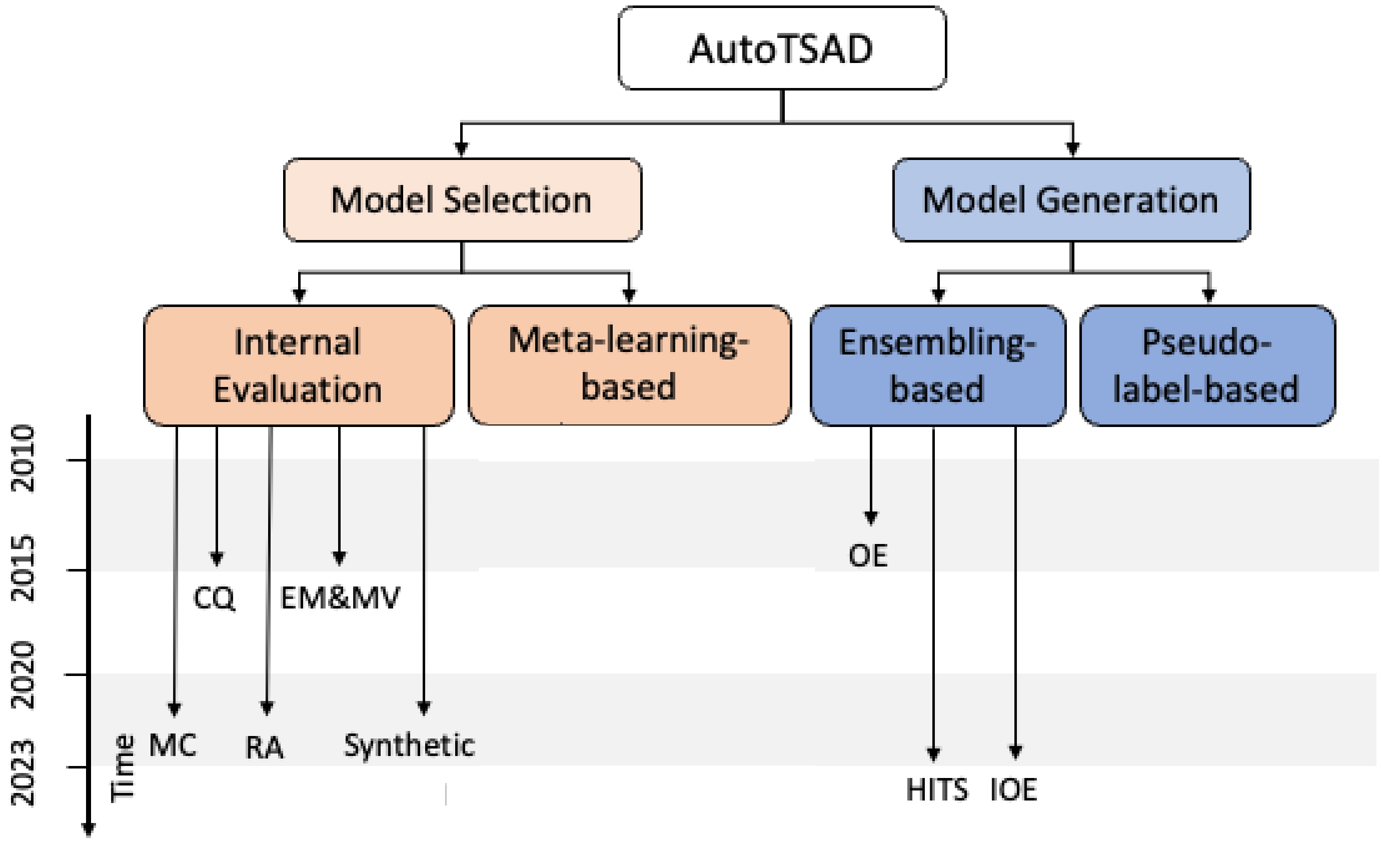

Automated Solution: *Ensembling-based*

Definition: Integrate predictions from the candidate model set

- Full: OE
- Selective: HITS, IOE

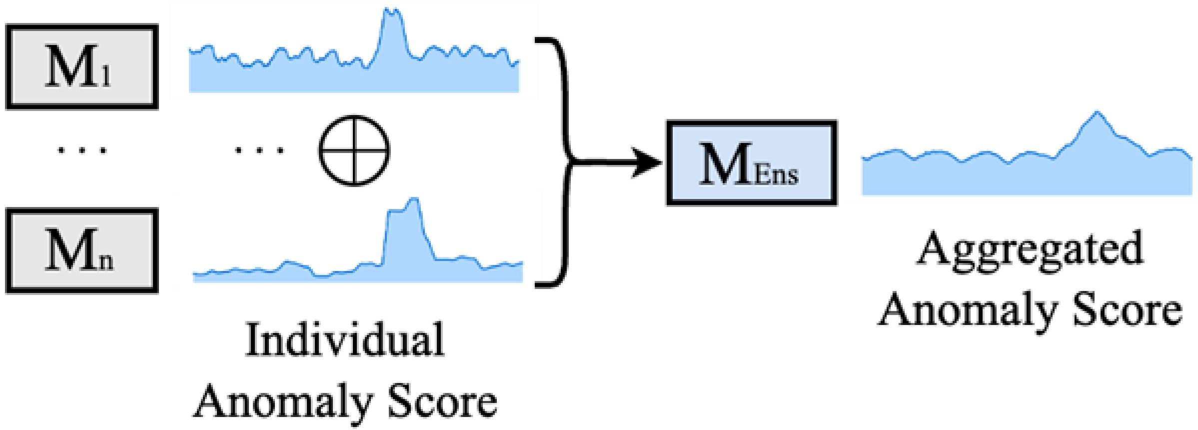

Automated Solution: *Pseudo-label-based*

Definition: Generate pseudo-labels to transform the unsupervised anomaly detection problem into a supervised framework

- AutoOD: Augment, Clean
- Booster: UADB

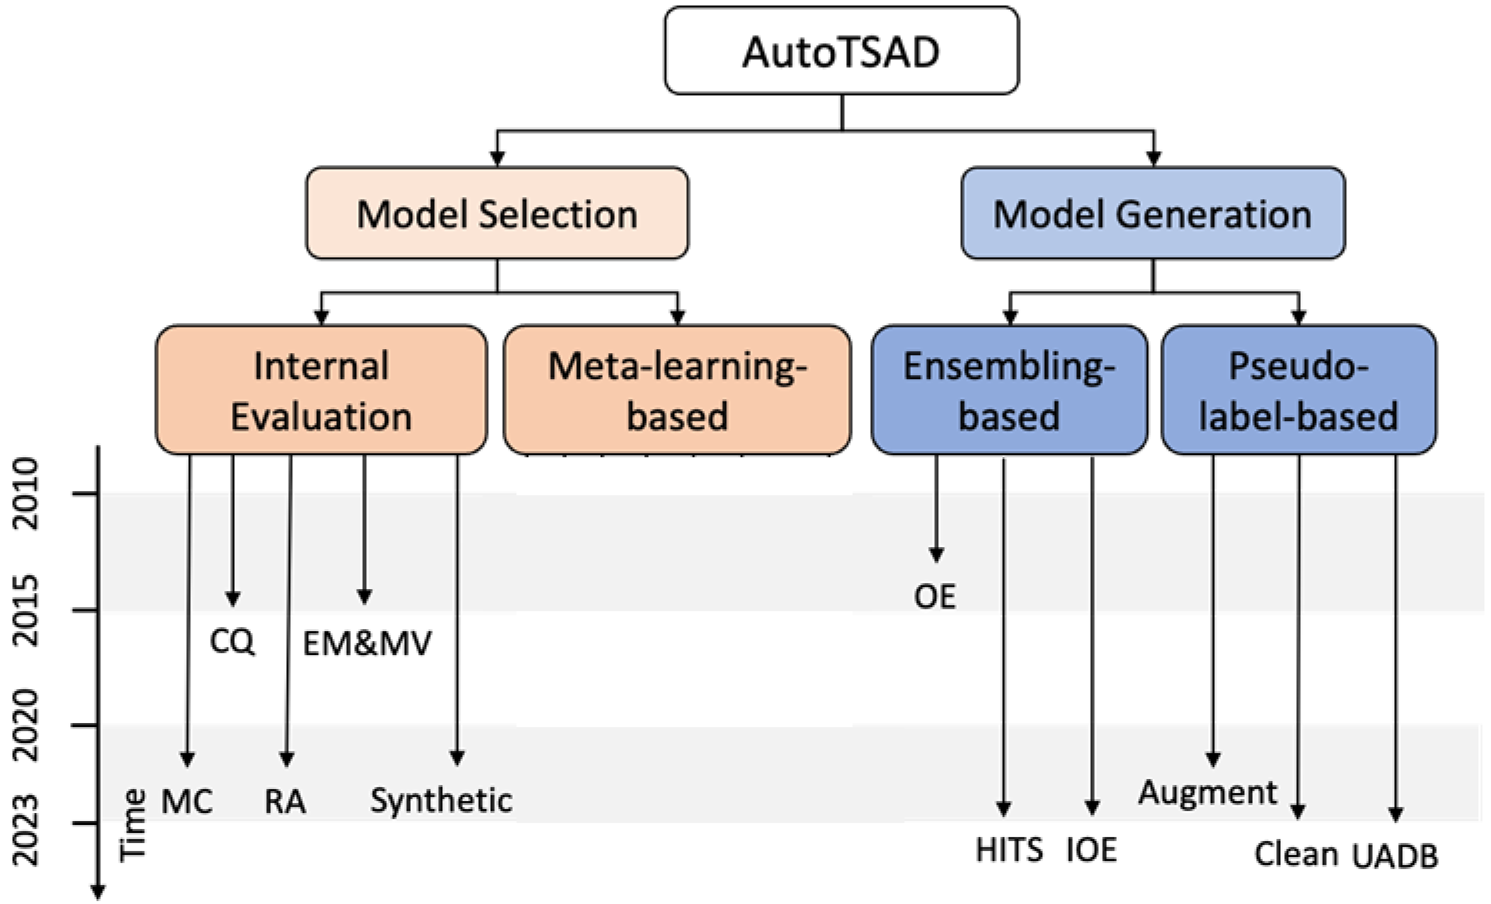

Automated Solution: *Pseudo-label-based*

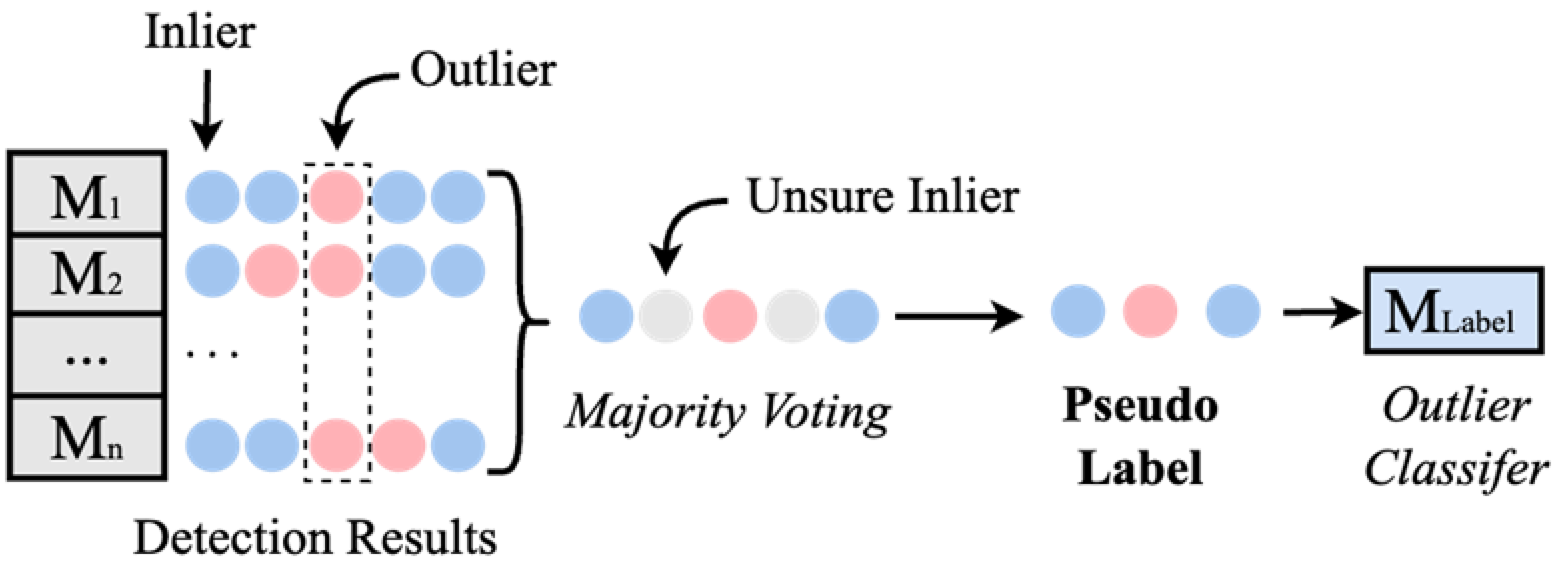

Pseudo-label-based Method Framework [16].

Automated Solution: *Meta-learning-based*

Definition: Using insights from historical labeled datasets to select the best model for new data

- Classification: Auto-Selector, MSAD
- Regression: RG, UReg, Cfact
- Nearest Neighbor: kNN
- Other Optimization: ISAC, MetaOD

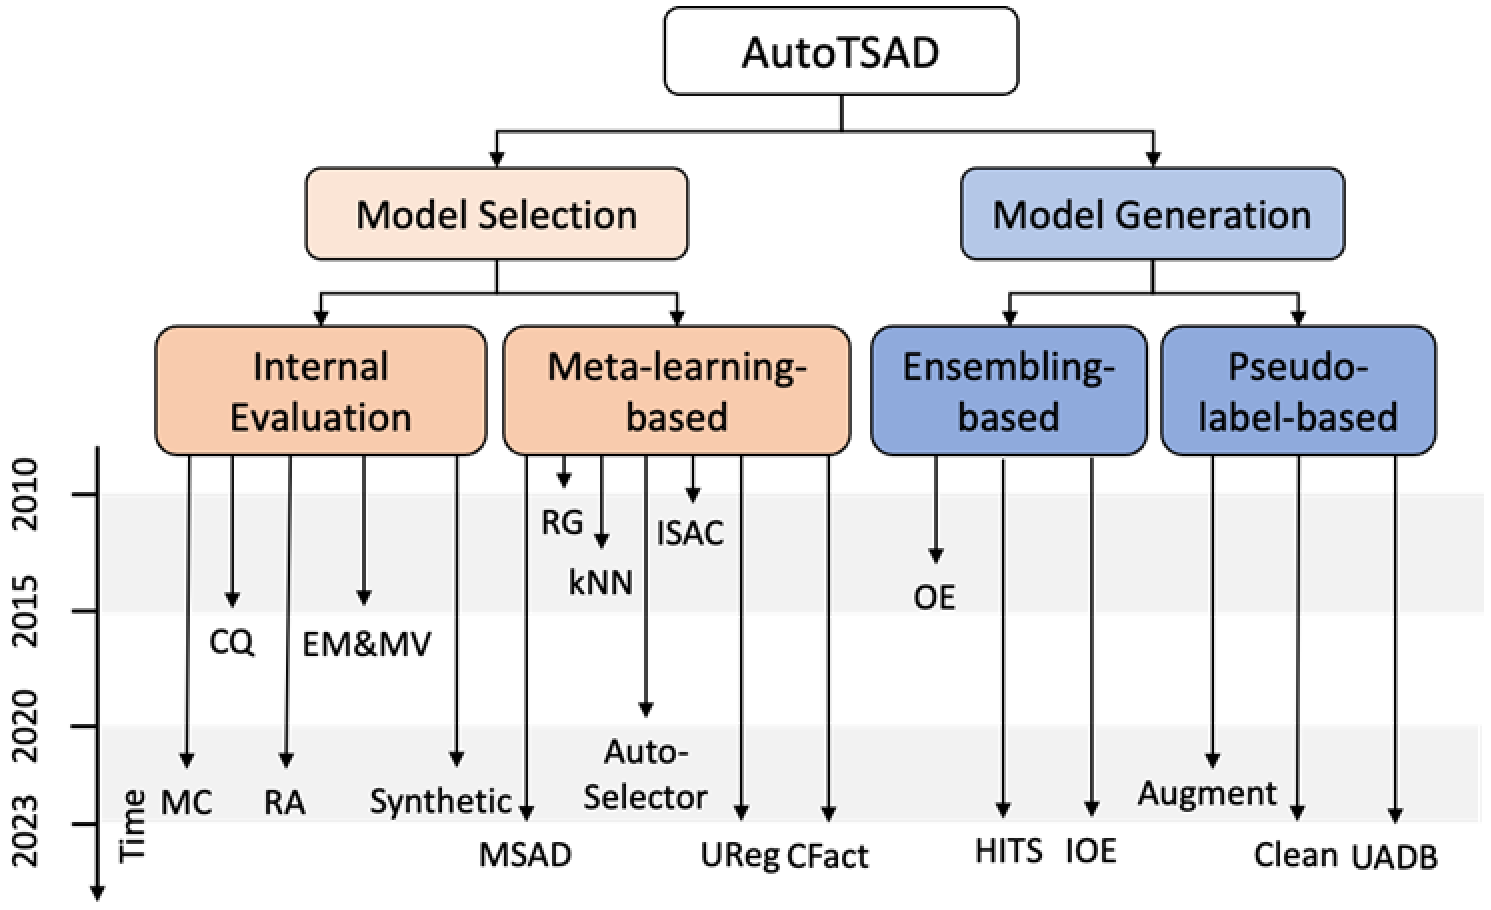

Automated Solution: *Meta-learning-based*

Definition: Using insights from historical labeled datasets to select the best model for new data.

- Other Optimization: ISAC, MetaOD

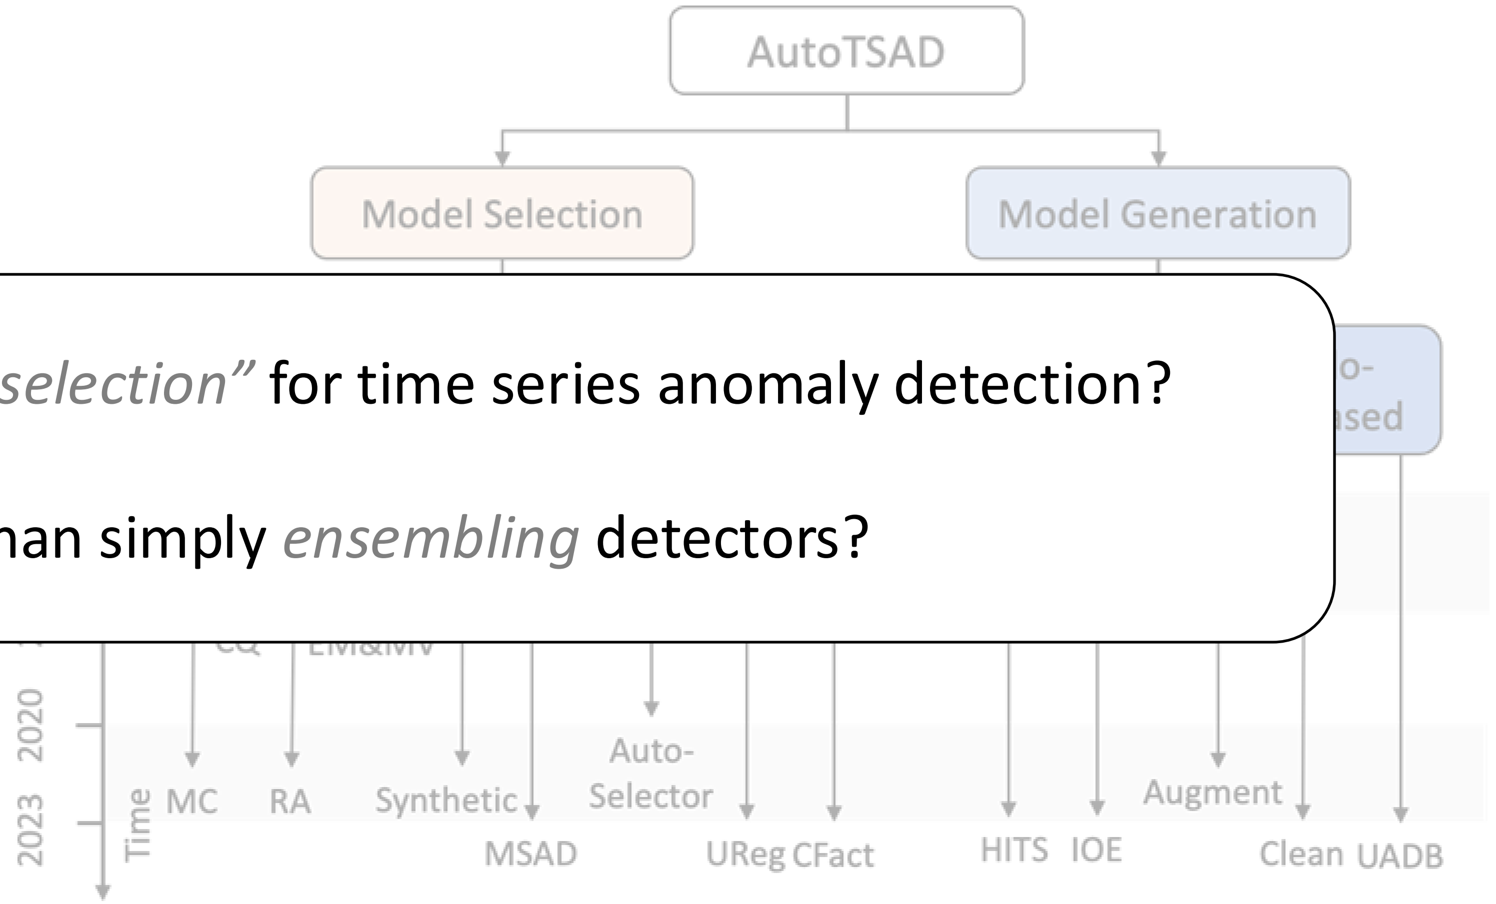

How can we do *"model selection"* for time series anomaly detection?

Is it better than simply *ensembling* detectors?

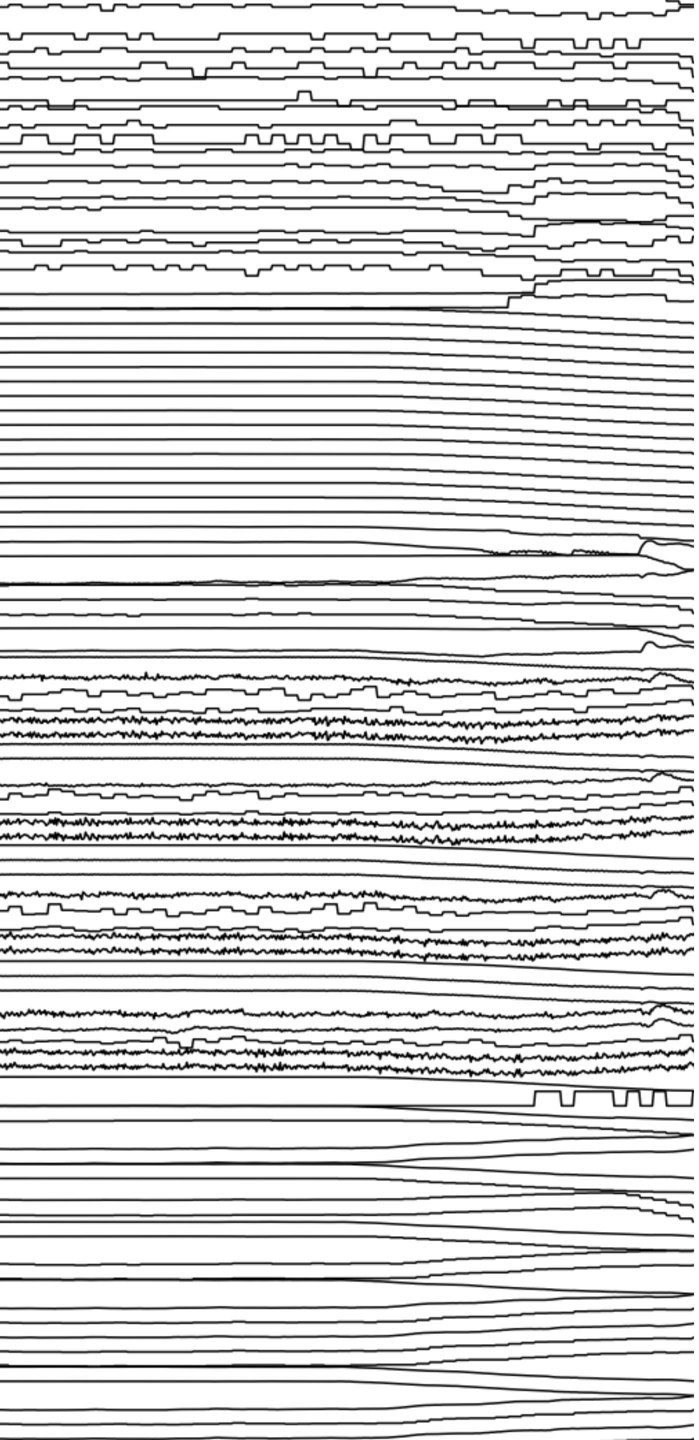

IV. MSAD

Model Selection for Anomaly Detection

MSAD: *Ensembling versus Model Selection*

Ensembling is proposed as a mitigation strategy to the previous limitation [17]

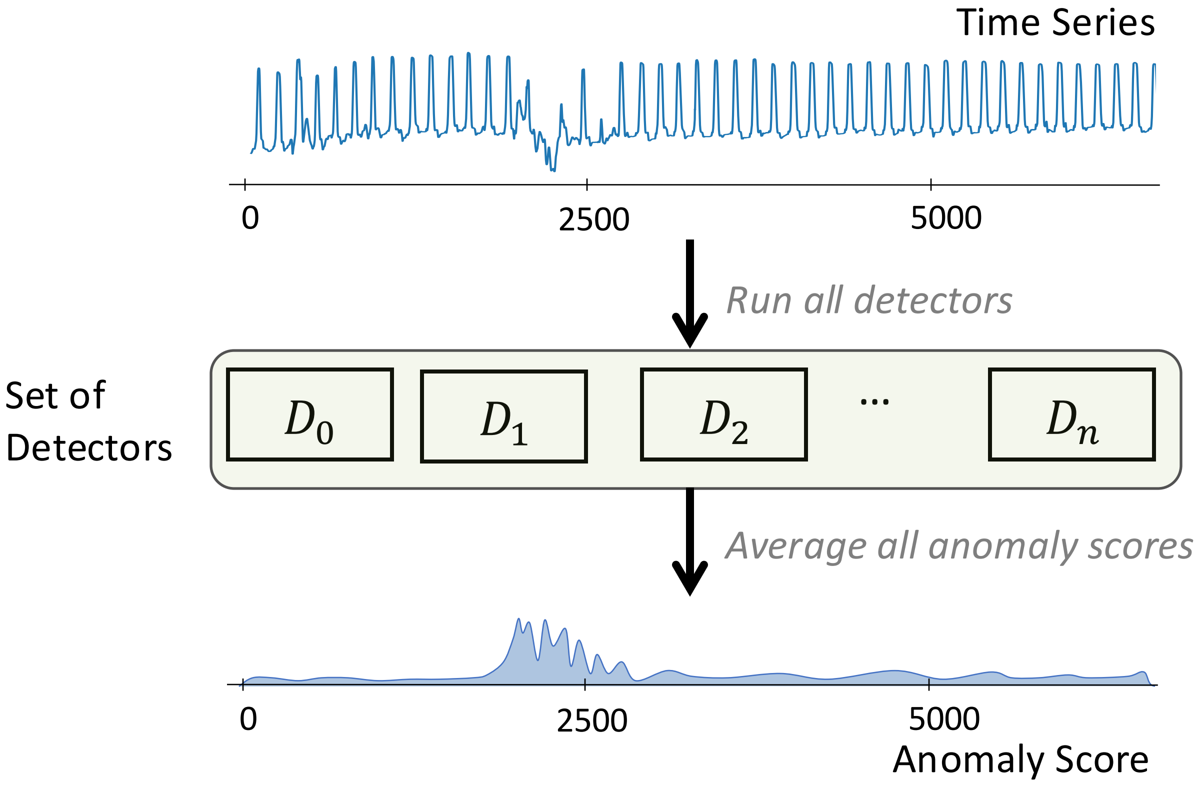

MSAD: *Ensembling versus Model Selection*

Ensembling is proposed as a mitigation strategy to the previous limitation [17]

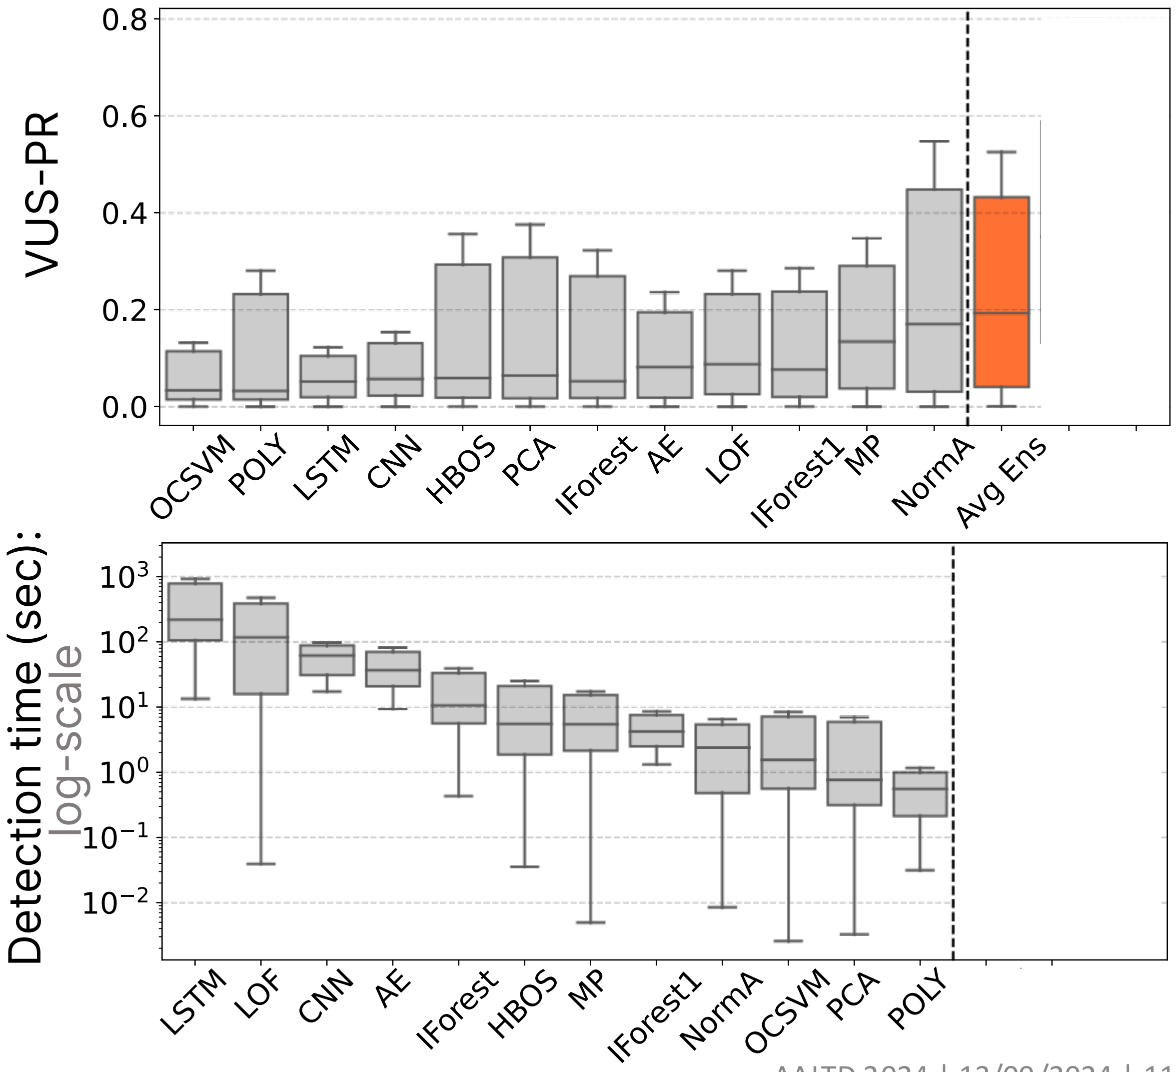

MSAD: *Ensembling versus Model Selection*

Ensembling is proposed as a mitigation strategy to the previous limitation [17]

… But is problematic in terms of execution time

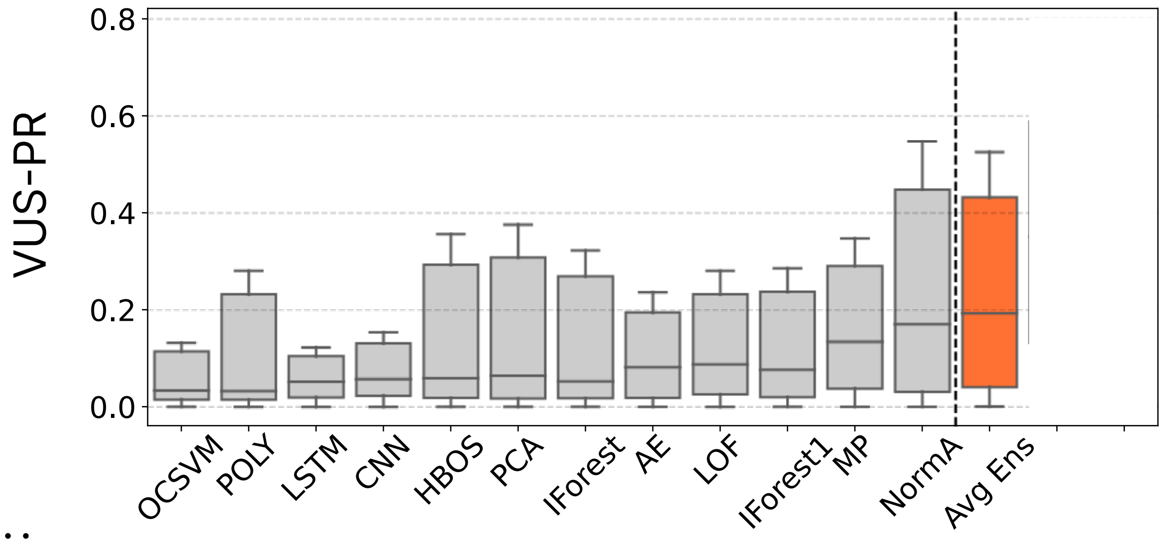

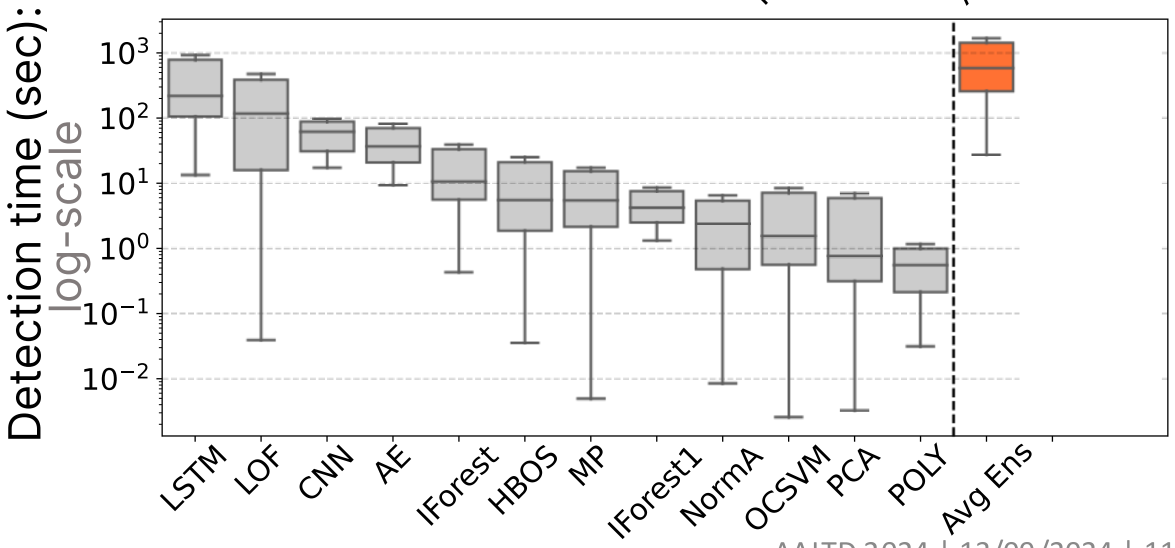

MSAD: *Ensembling versus Model Selection*

Ensembling is proposed as a mitigation strategy to the previous limitation [17]

… But is problematic in terms of execution time

Model Selection (MS) is a solution to reduce the execution time

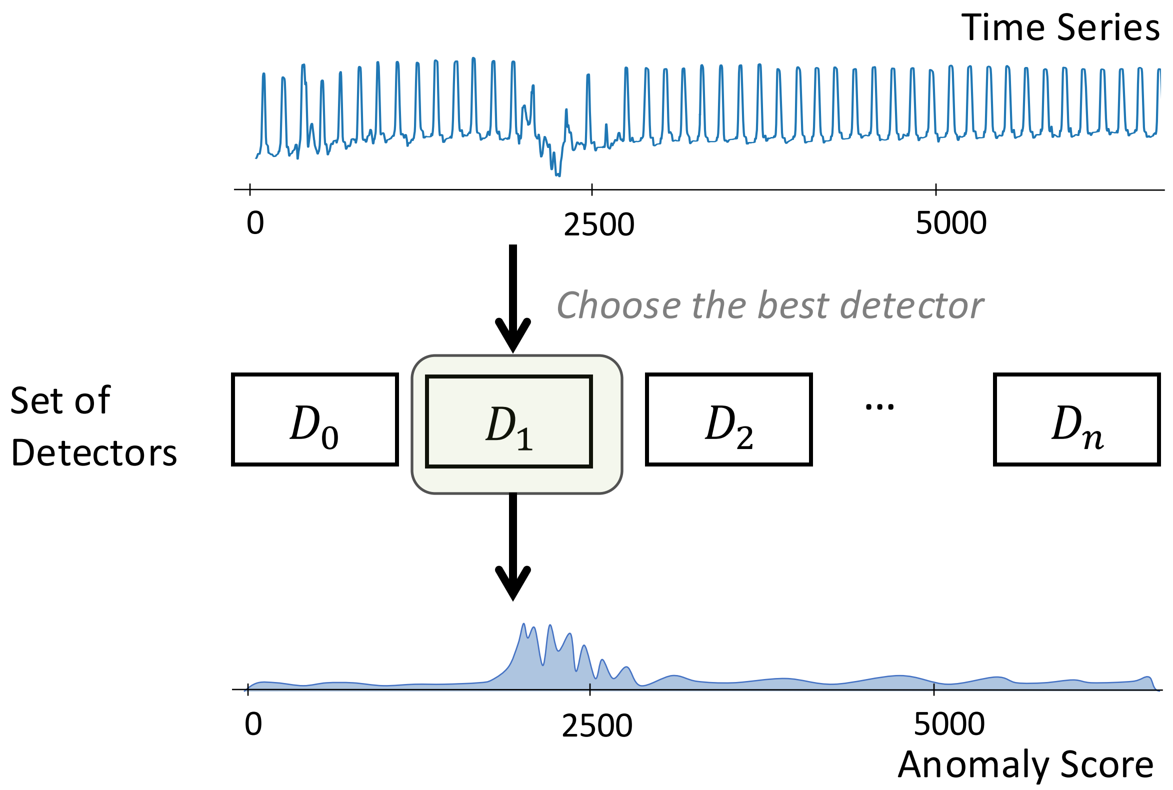

MSAD: *Ensembling versus Model Selection*

Ensembling is proposed as a mitigation strategy to the previous limitation [17]

… But is problematic in terms of execution time

Model Selection (MS) is a solution to reduce the execution time

The best possible achievable performances (Oracle) is motivating

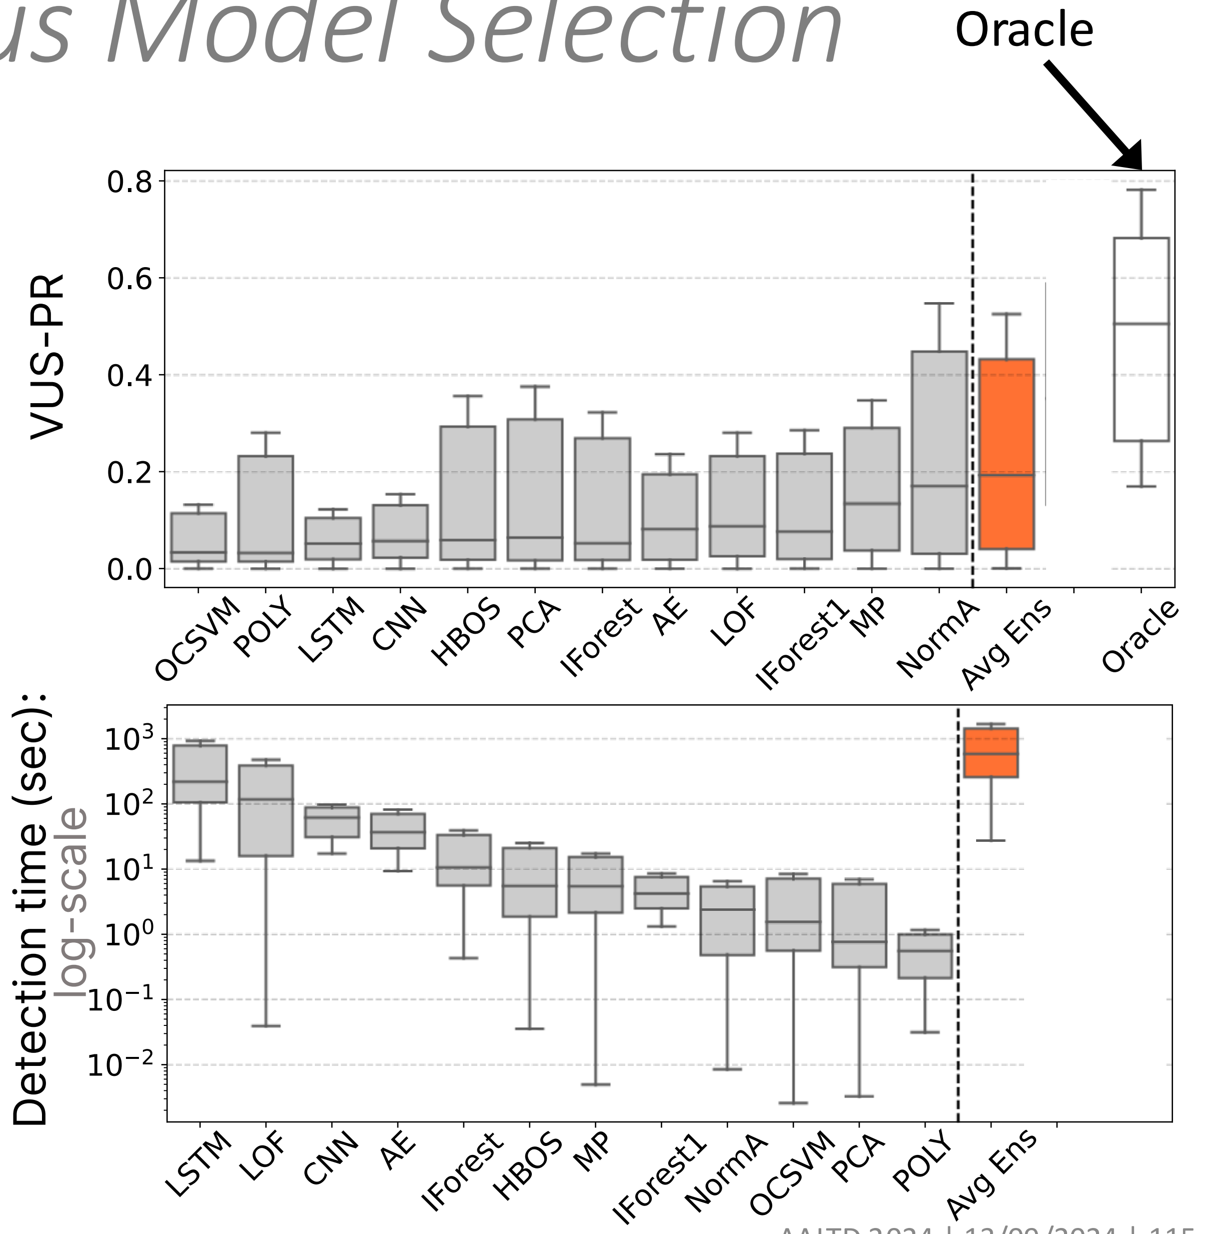

MSAD: *Ensembling versus Model Selection*

Ensembling is proposed as a mitigation strategy to the previous limitation (Aggarwal, C., C., et al. SIGKDD 2015)

The best possible achievable performances (Oracl

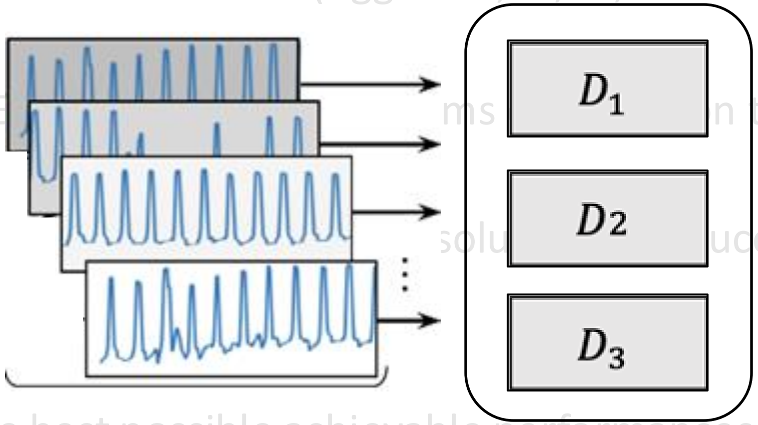

Time Series For Training

Candidate Model Set

	D_1	D_2	D_3	...
TS_1	0.5	0.7	0.9	
TS_2	0.6	0.4	0.7	
TS_3	0.5	0.8	0.6	
...				

Performance Matrix

	D_1	D_2	D_3	...
TS_1	0	0	1	
TS_2	0	0	1	
TS_3	0	1	0	
...				

Label

Time series classification methods could be a solution

MSAD: *Classification for Model Selection*

Research Questions (RQs)

1. What is the best approach:
    1. Individual Detectors
    2. Average Ensembling (Avg Ens)
    3. Model Selection (MS)

2. What is the best input: Time Series Features OR Raw Values?

3. What-if model selection is tested on completely new datasets?

MSAD: *Experimental Pipeline*

(a) Time series T

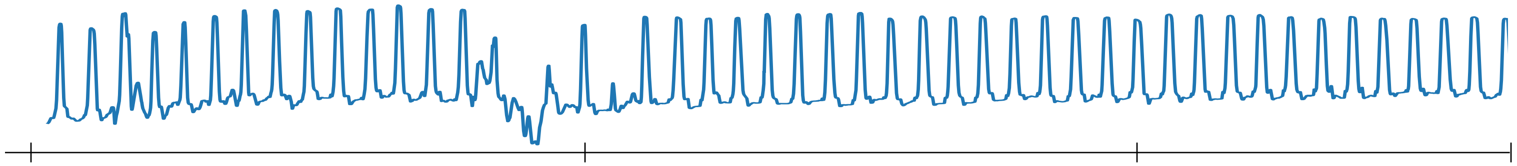

Step 1: Acquiring Labeled Time Series

We use the TSB-UAD benchmark [14], on which we know in advance which detector is the best for each time series.

MSAD: *Experimental Pipeline*

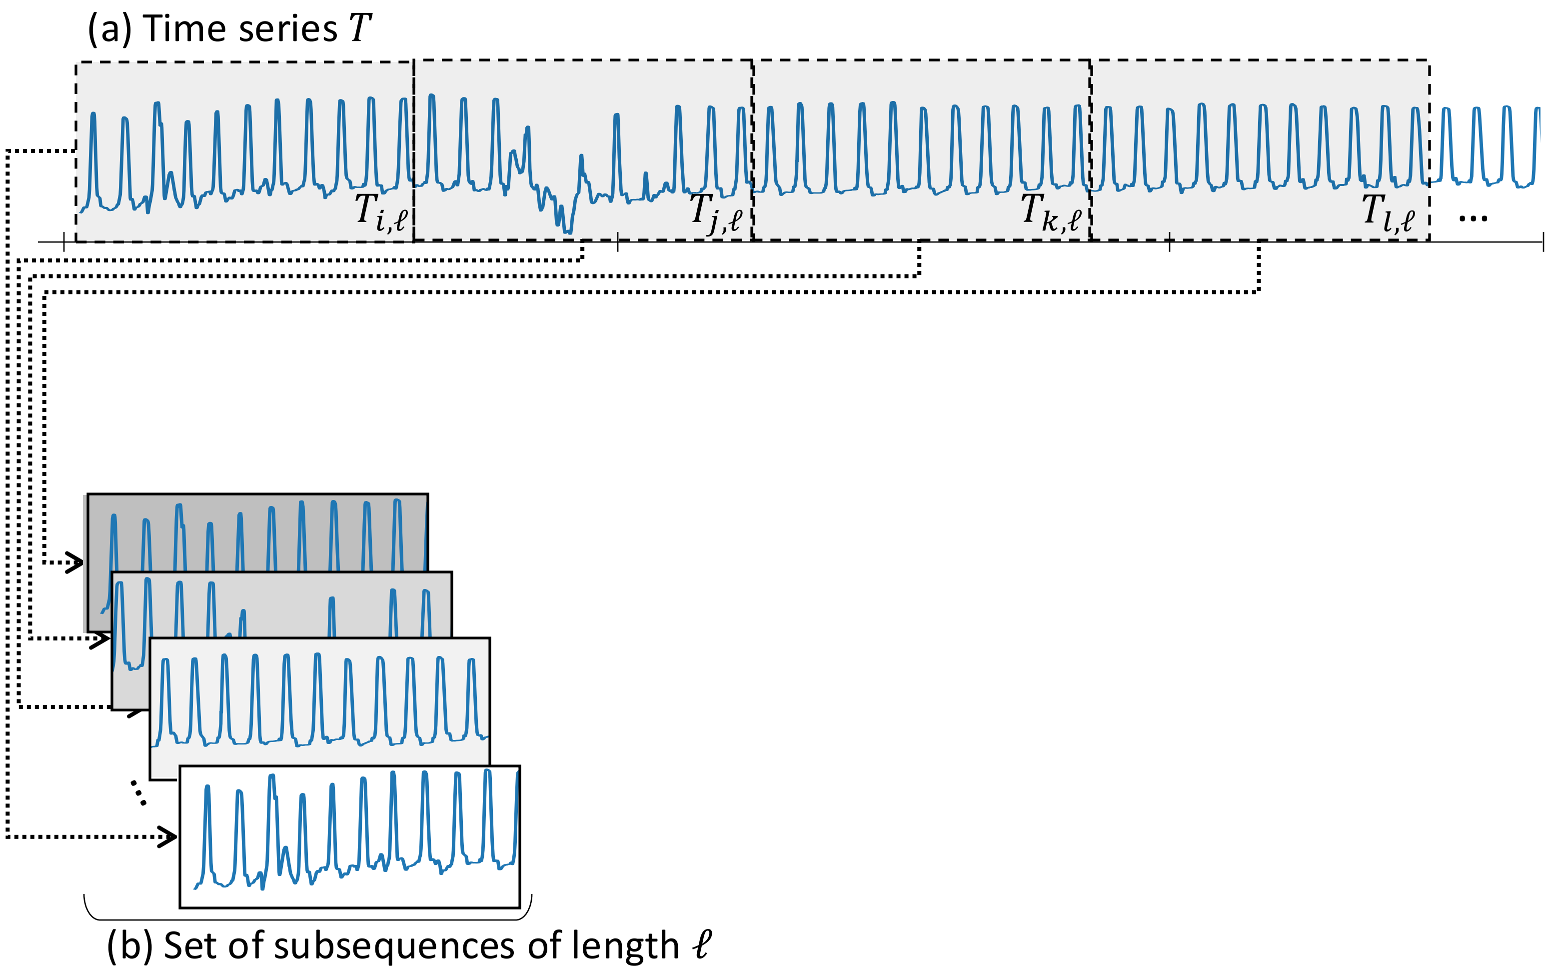

Step 2: Segmentation

We segment the time series into equal length subsequences.

Each subsequence is assigned to the same label (best detector)

MSAD: *Experimental Pipeline*

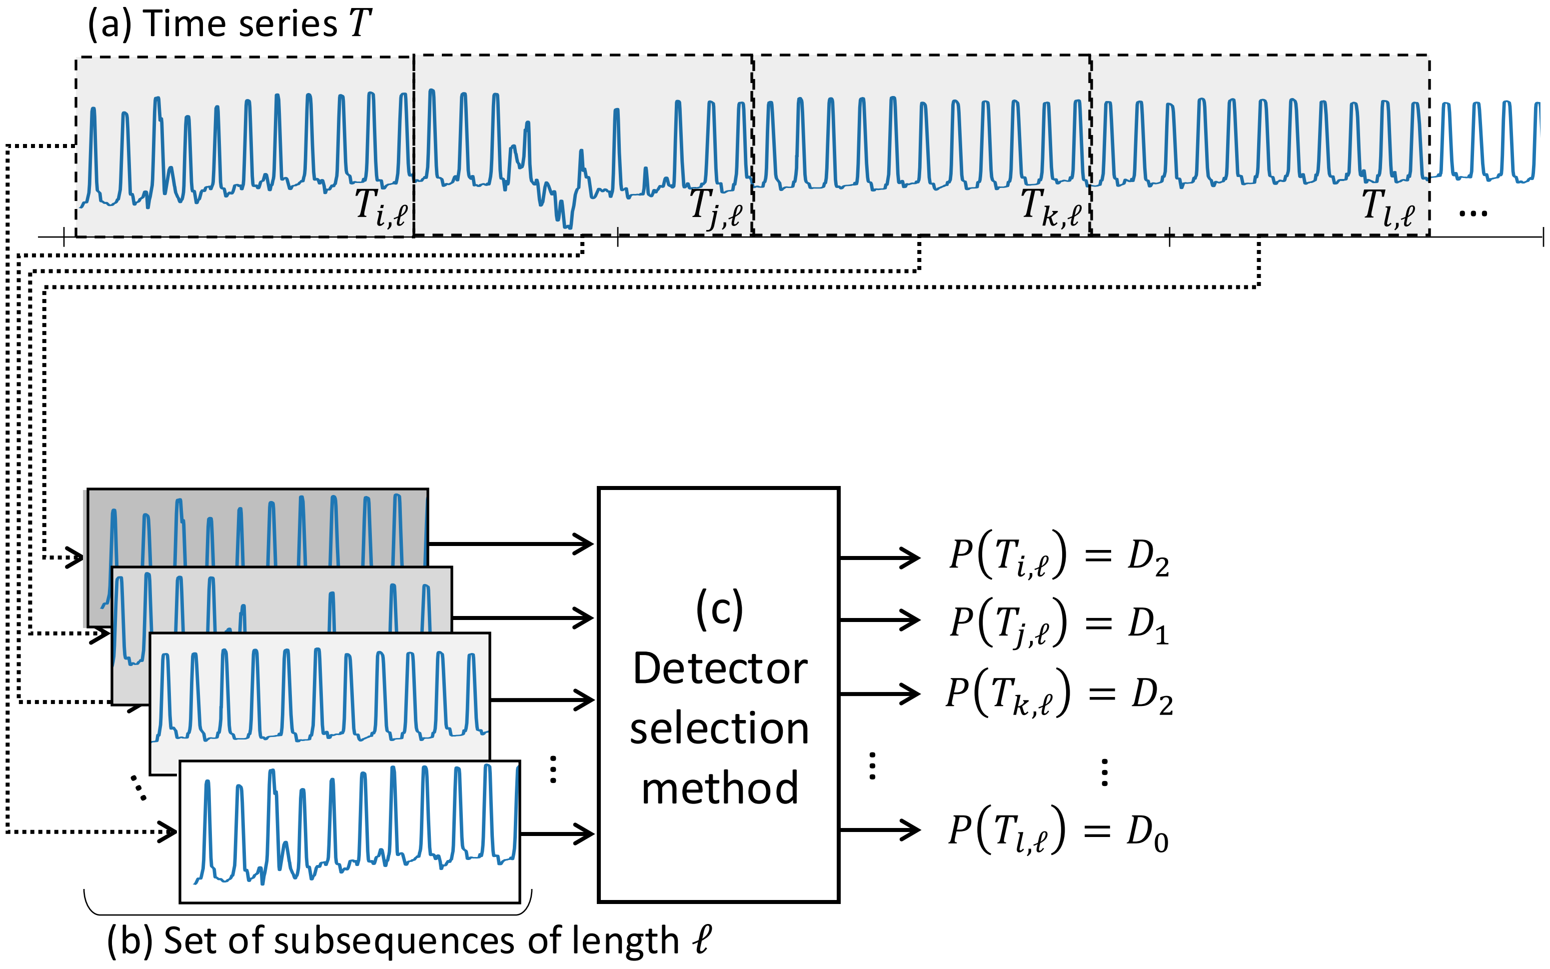

Step 3: Prediction

We train a time series classification method to predict which detector is the best (using the labels from TSB-UAD).

MSAD: *Experimental Pipeline*

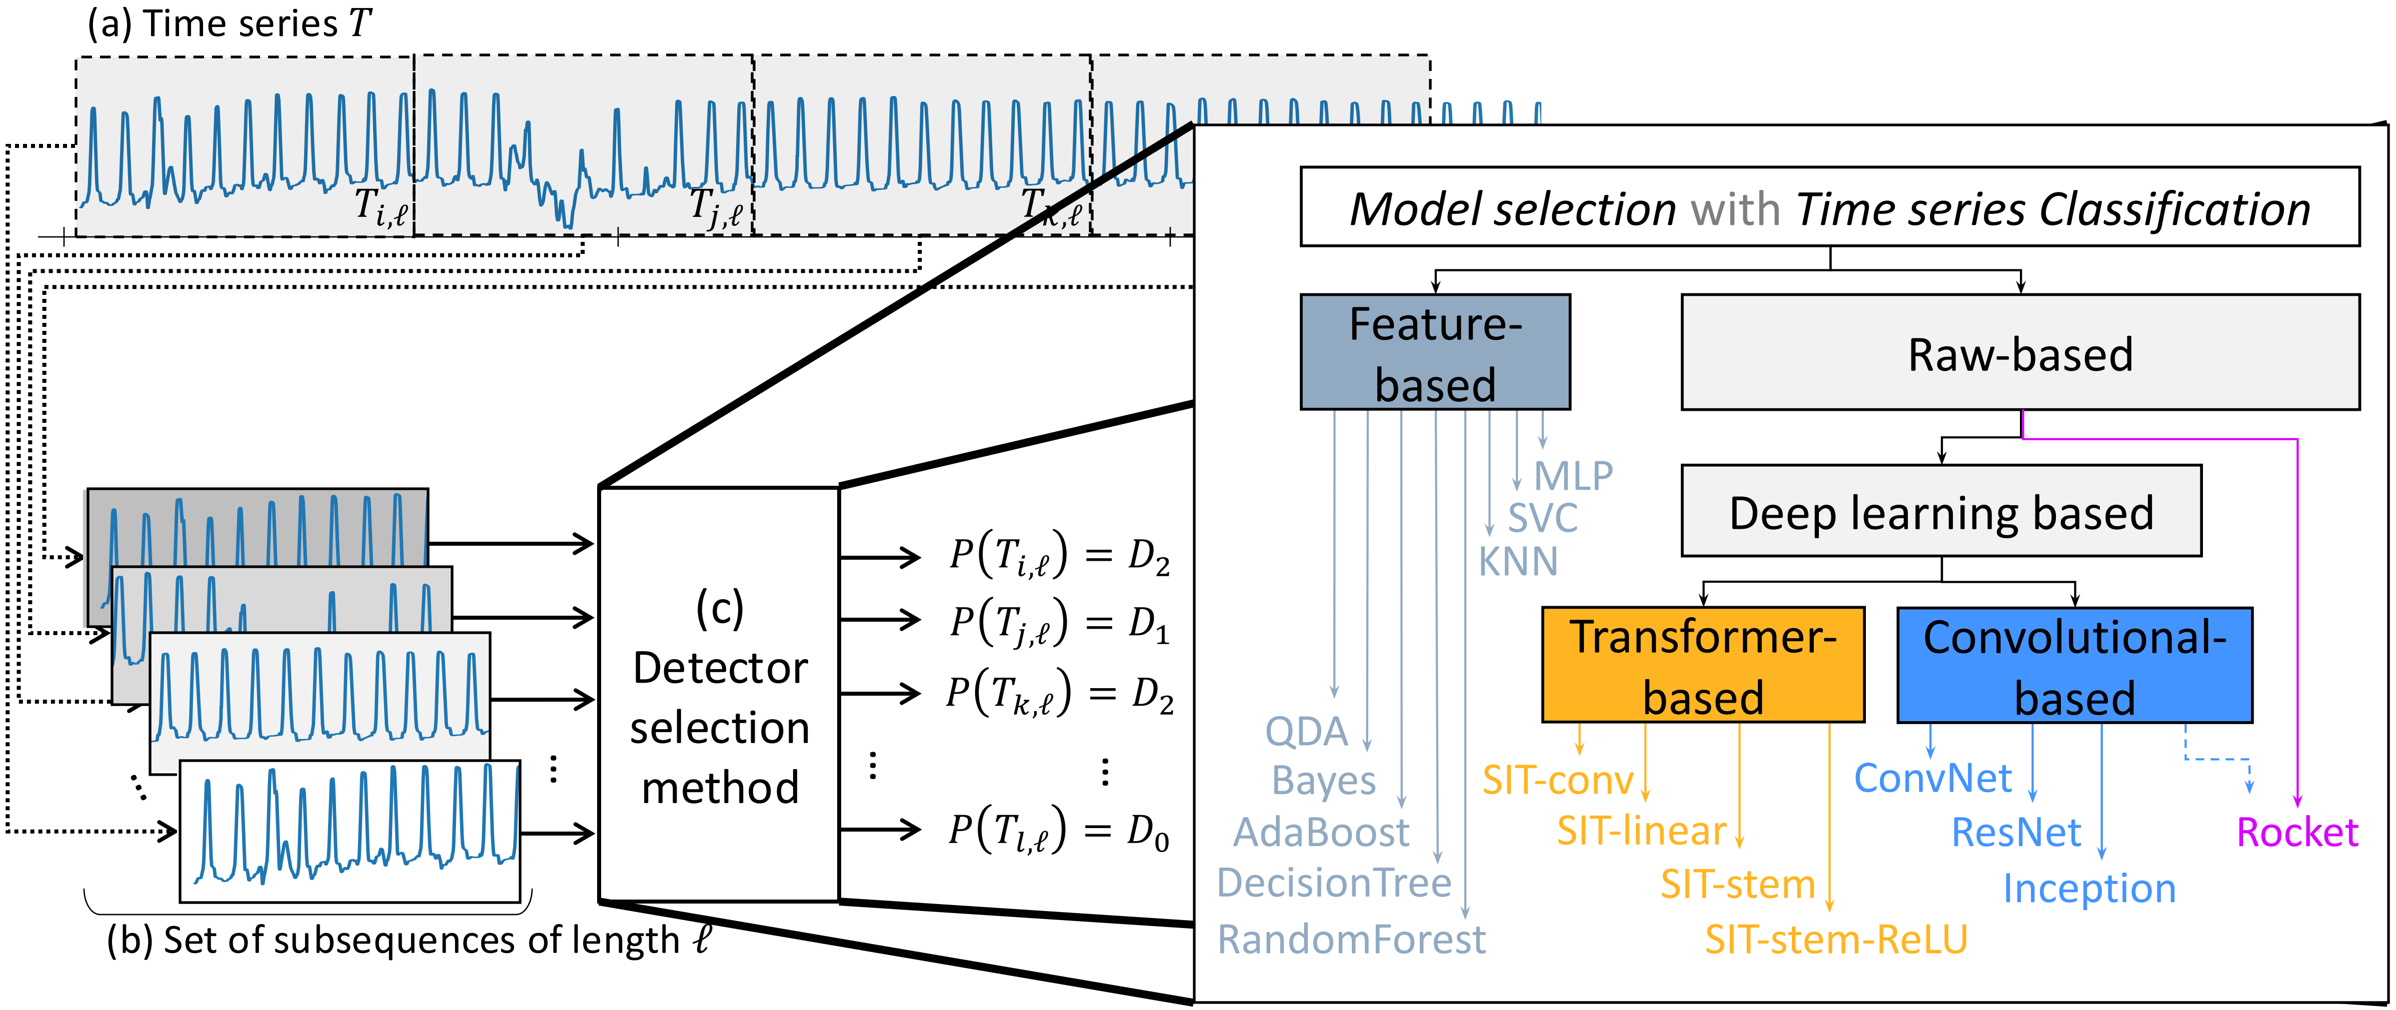

MSAD: *Experimental Pipeline*

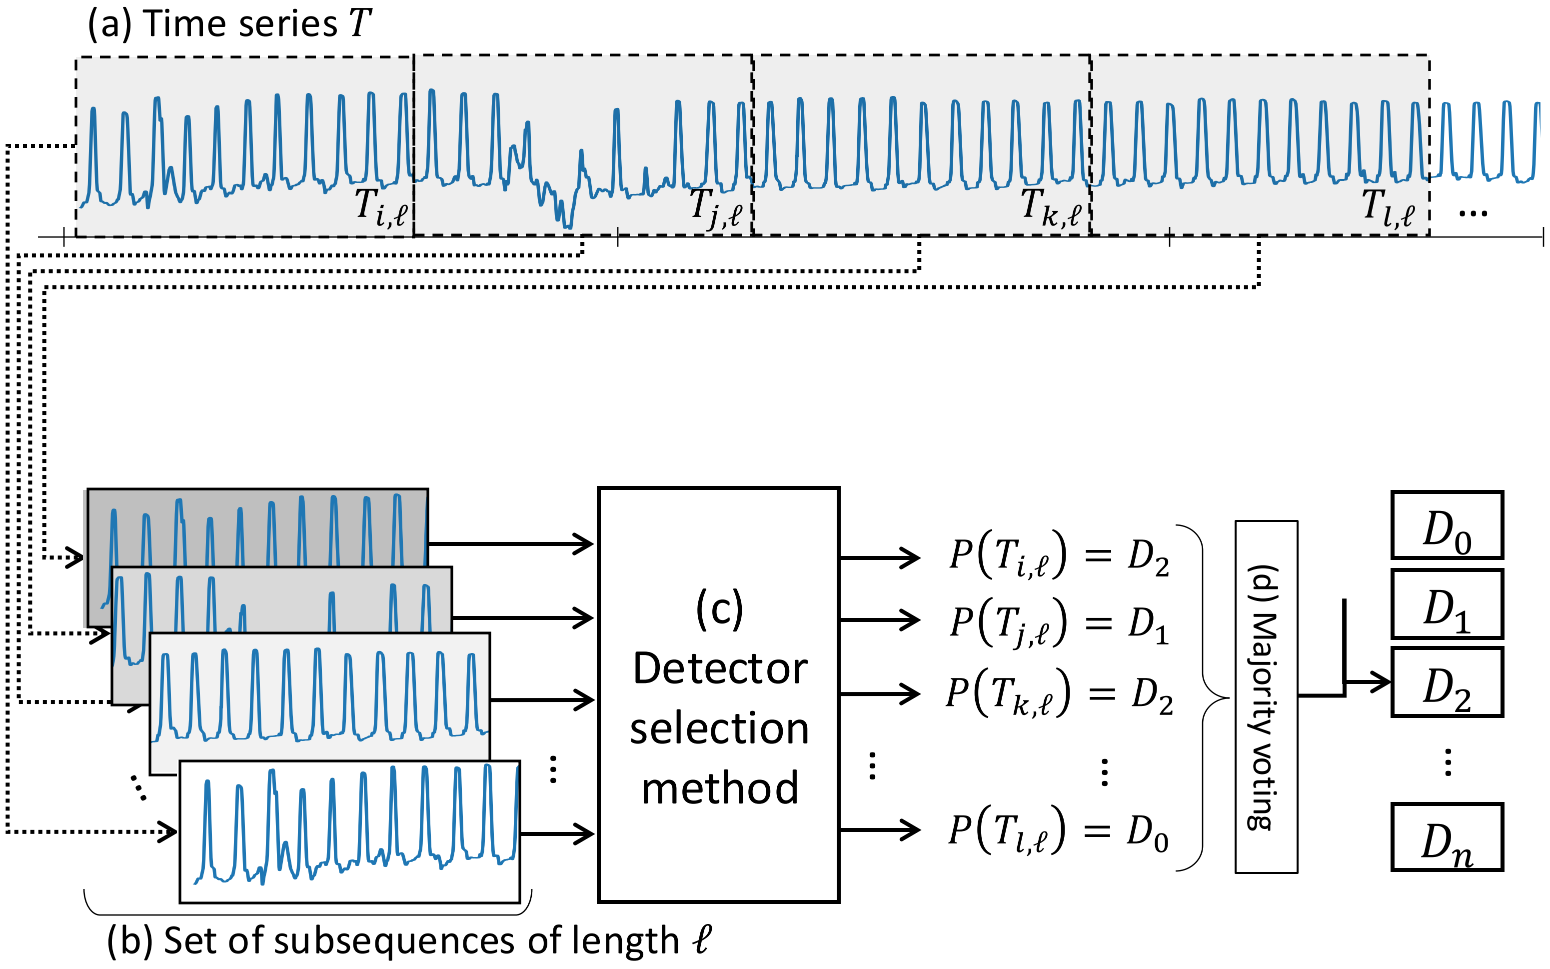

Step 4: Selection

We pick the most selected detector for all the subsequences of a time series.

MSAD: *Experimental Pipeline*

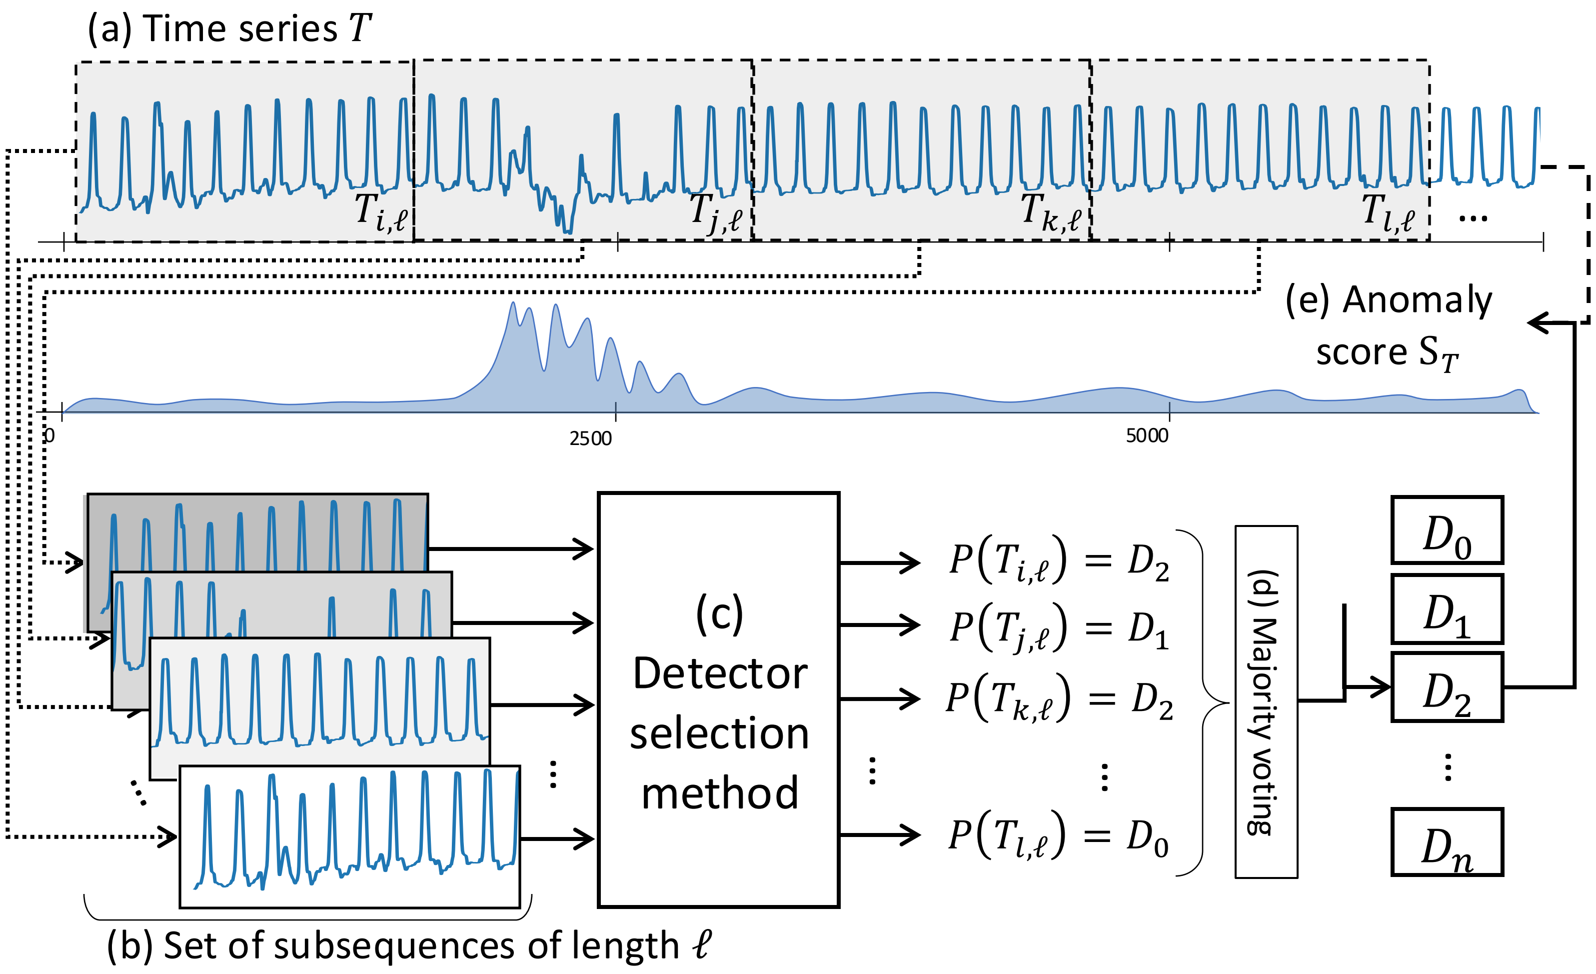

Step 5: Anomaly Score Computation

We finally compute the anomaly score using the selected detector.

MSAD: *Experimental Evaluation*

We conduct our experimental evaluation on the TSB-UAD benchmark :

MSAD: *Experimental Evaluation*

We conduct our experimental evaluation on the TSB-UAD benchmark :

16 time series classification methods:

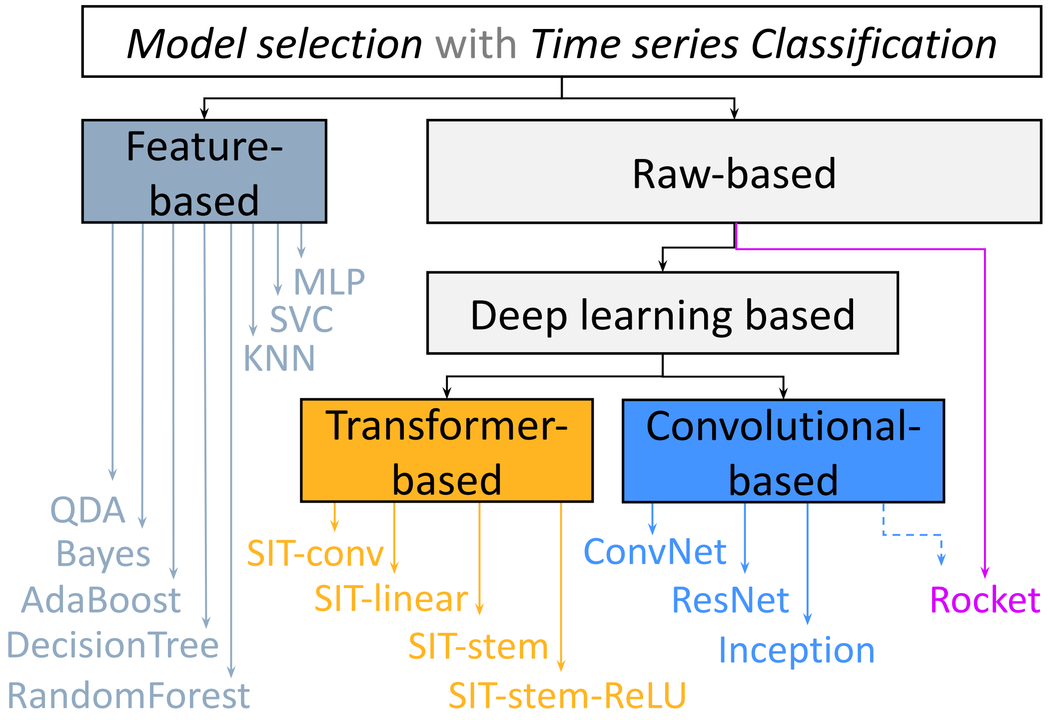

MSAD: *Experimental Evaluation*

We conduct our experimental evaluation on the TSB-UAD benchmark :

16 time series classification methods:

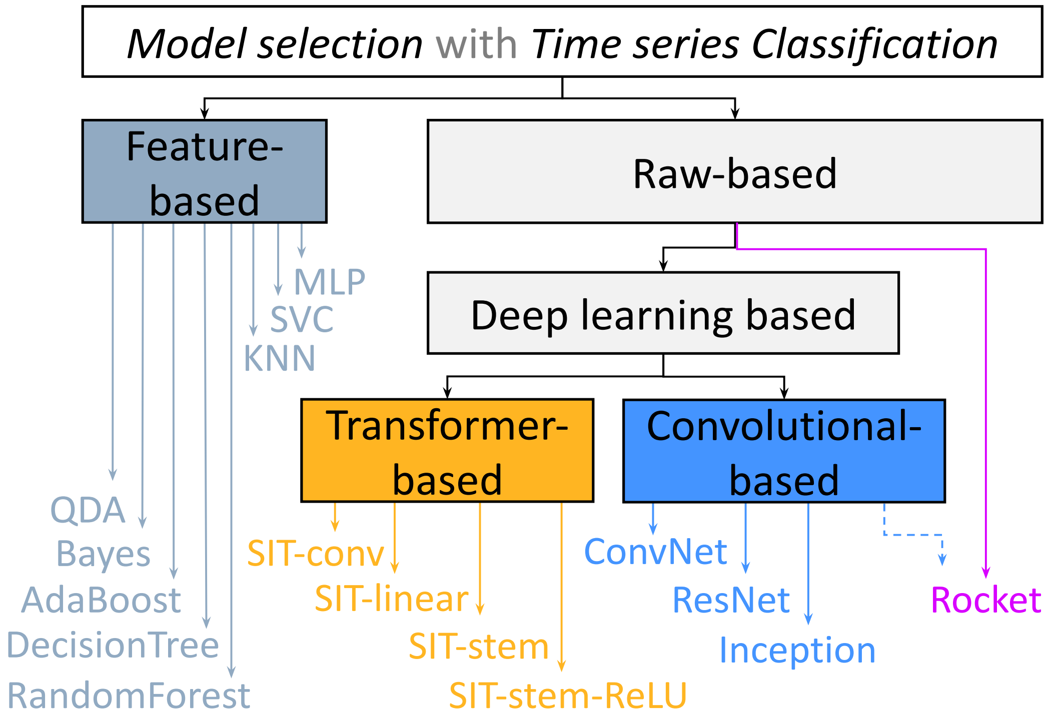

With 8 segmentation window lengths:

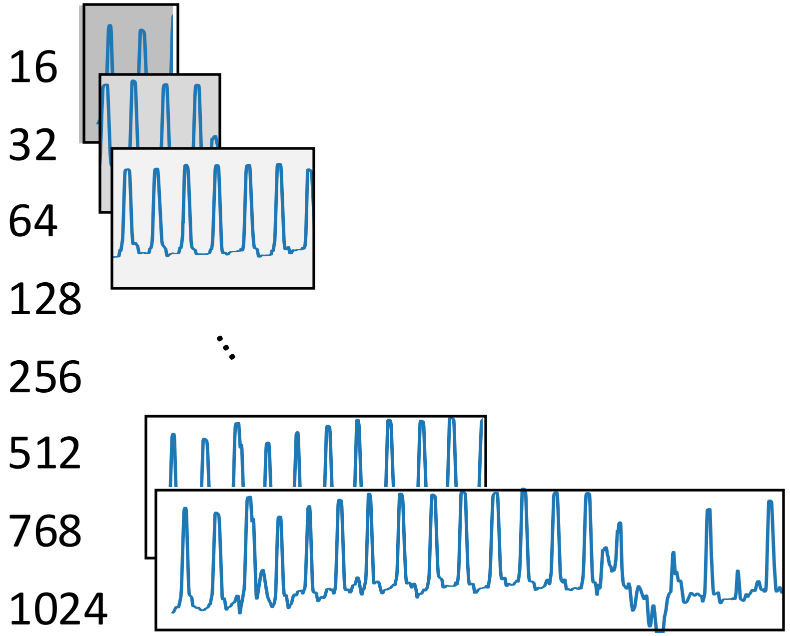

Random split (70/30) of TSB-UAD benchmark between train and test

MSAD: *Experimental Evaluation*

- **Raw values** is the best input compared to time series **features**

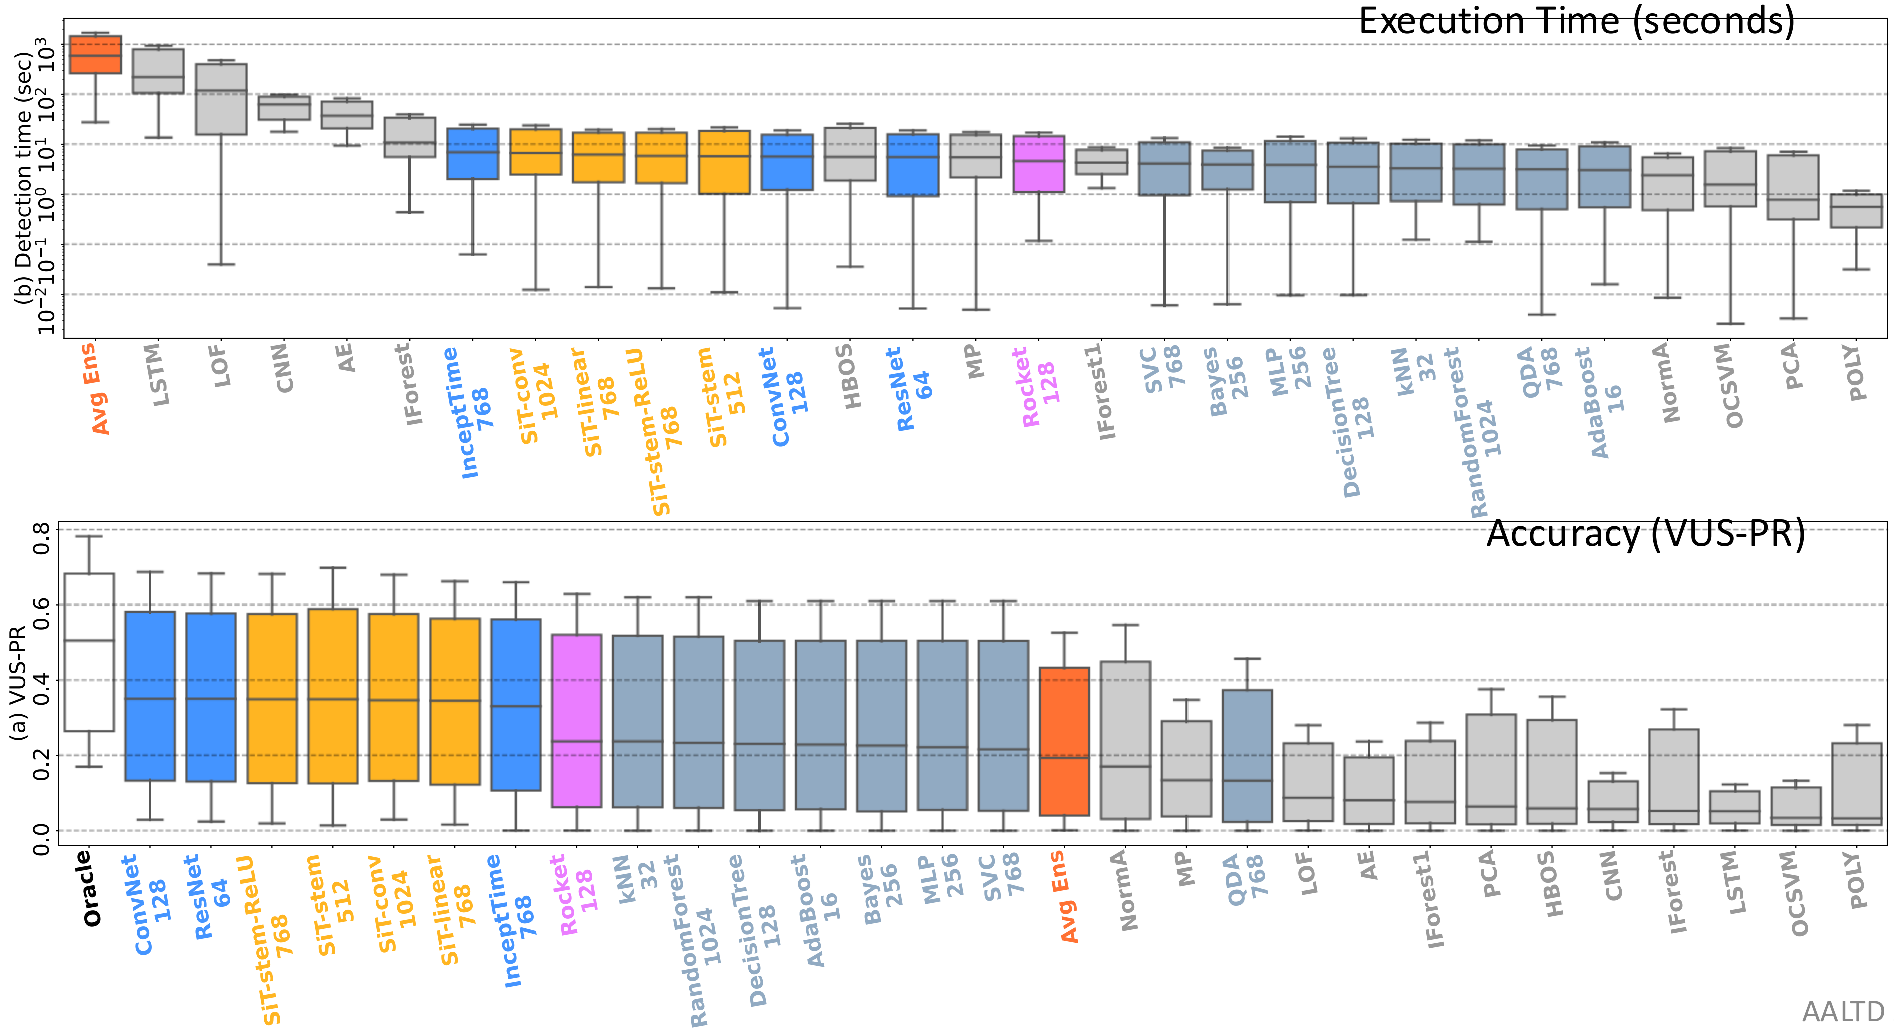

MSAD: *Experimental Evaluation*

- The window length influence is different based on the type of methods

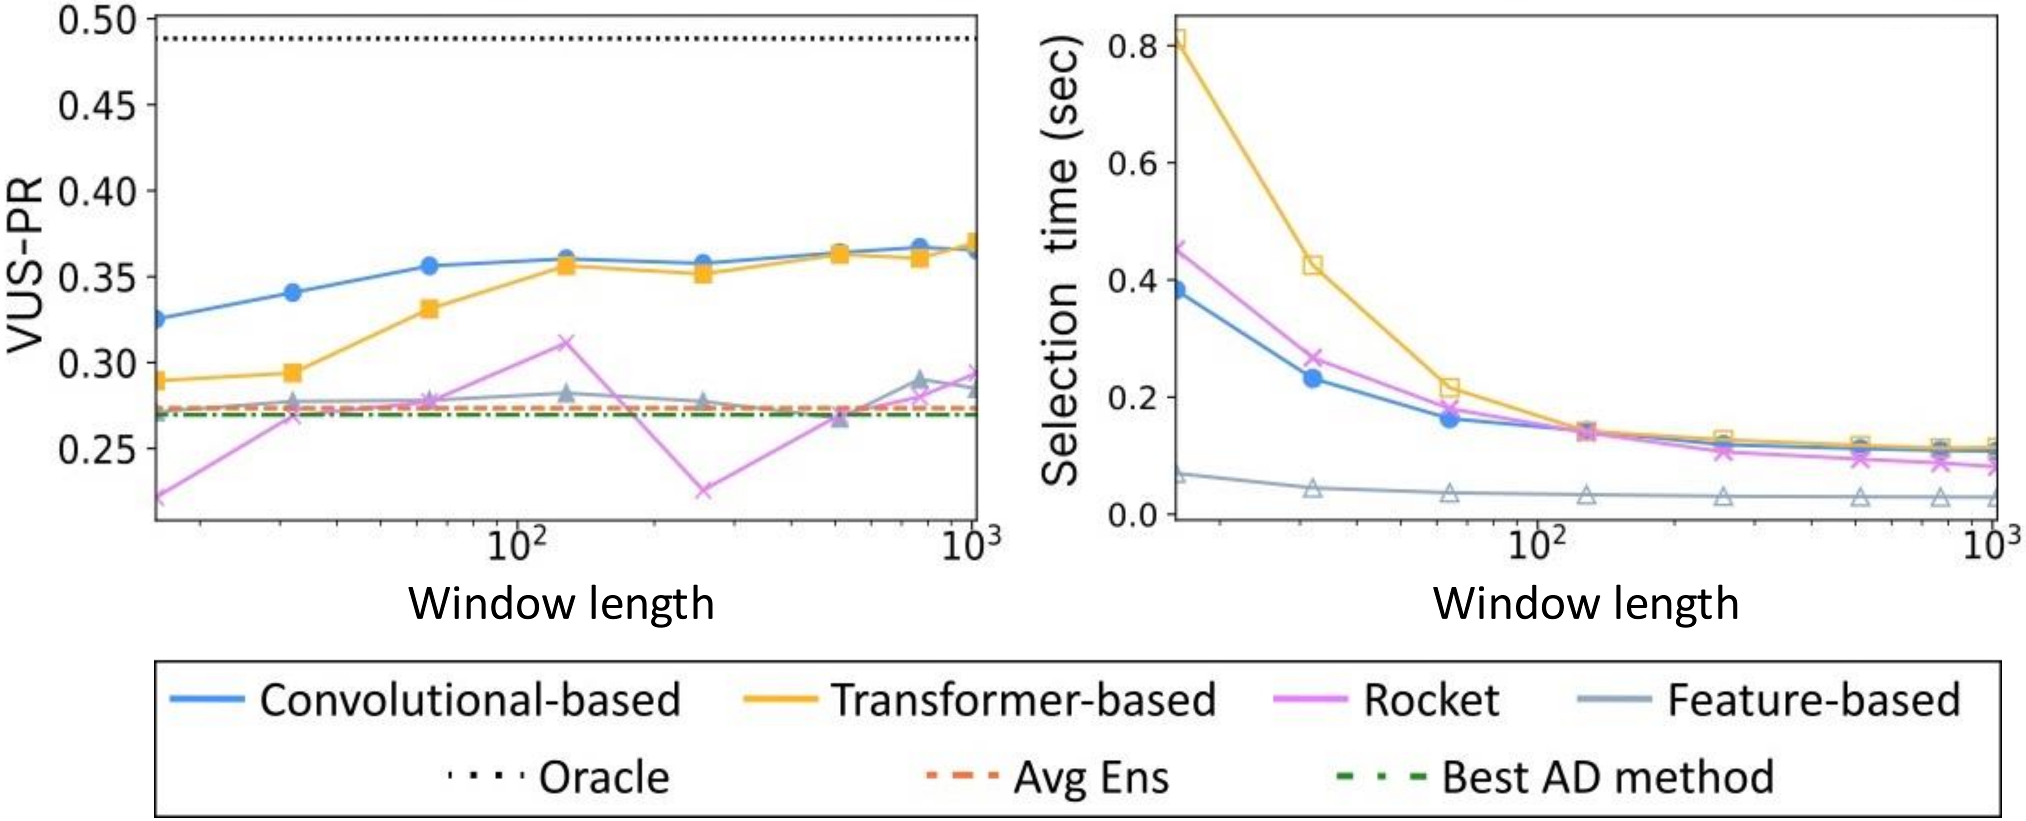

MSAD: *Experimental Evaluation*

- MS outperforms the **Individual detectors** and the **Avg Ens** in terms of accuracy
- MS outperforms **Avg Ens** in terms of execution time

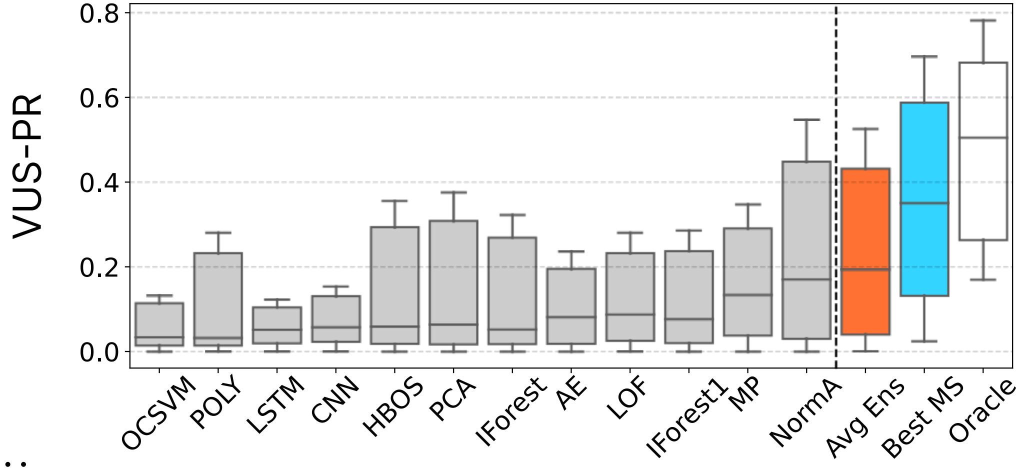

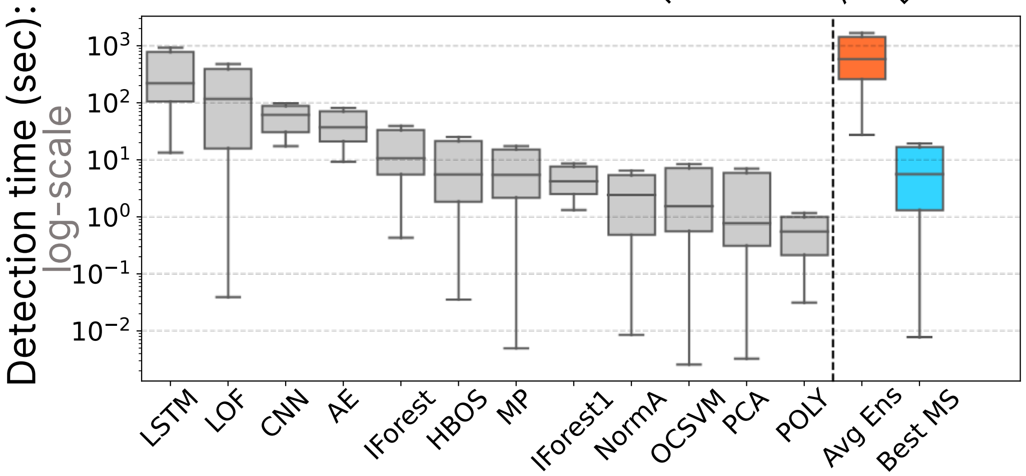

MSAD: *Experimental Evaluation*

- MS outperforms the **Individual detectors** and the **Avg Ens** in terms of accuracy

- MS outperforms **Avg Ens** in terms of execution time

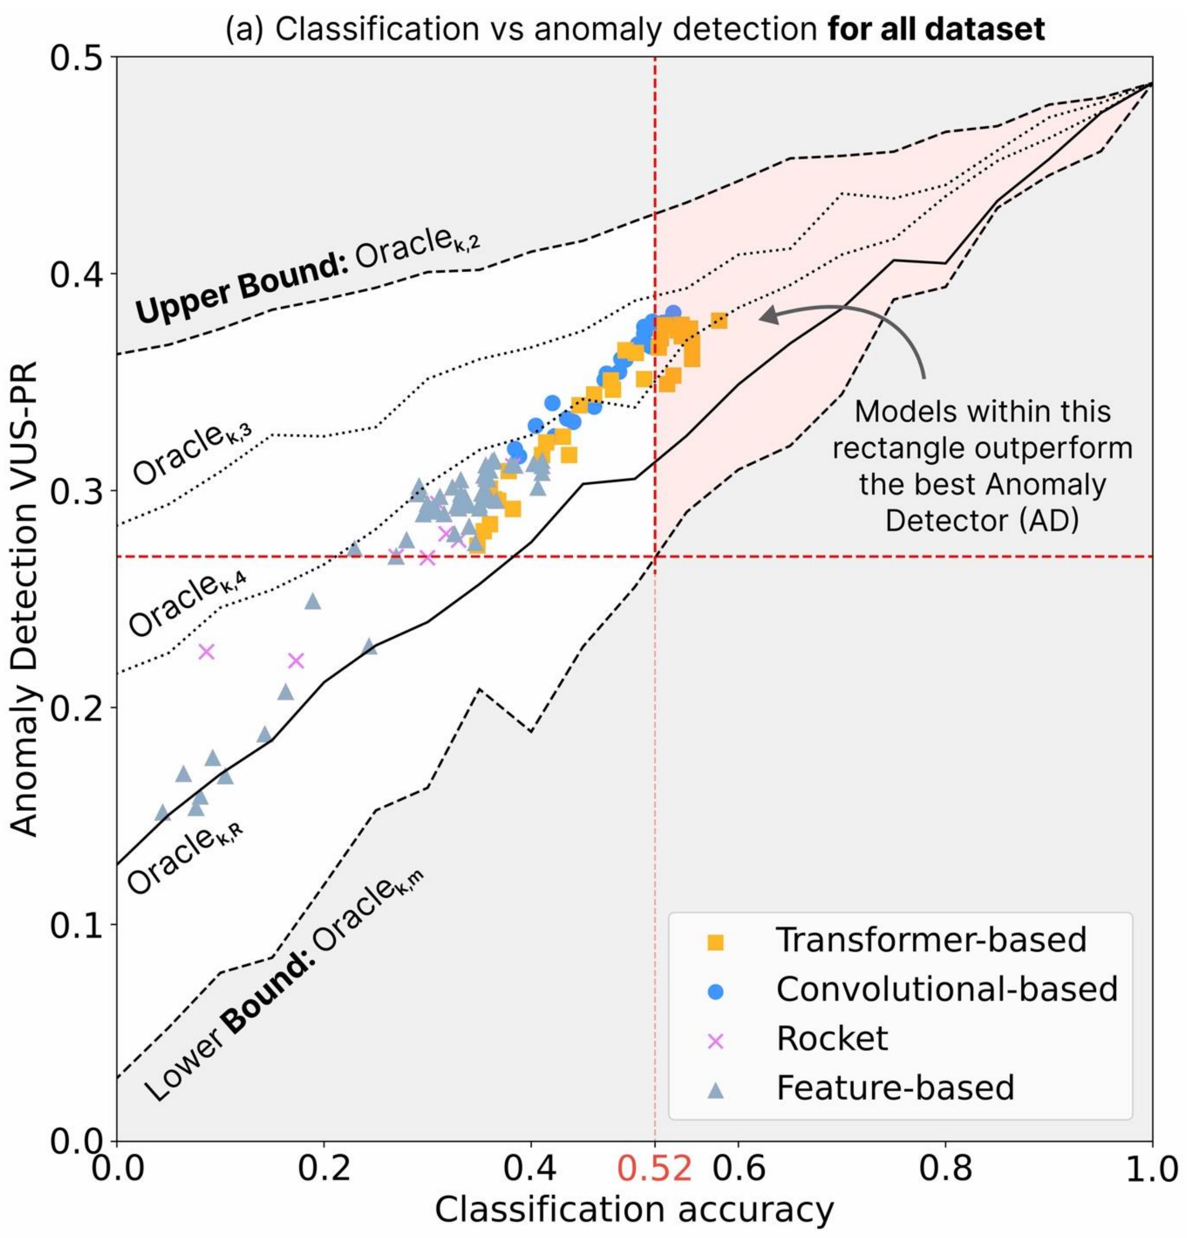

MSAD: *Experimental Evaluation*

- **MS** outperforms the **Individual detectors** and the **Avg Ens** in terms of accuracy
- **MS** outperforms **Avg Ens** in terms of execution time
- Potential improvement in terms of classification

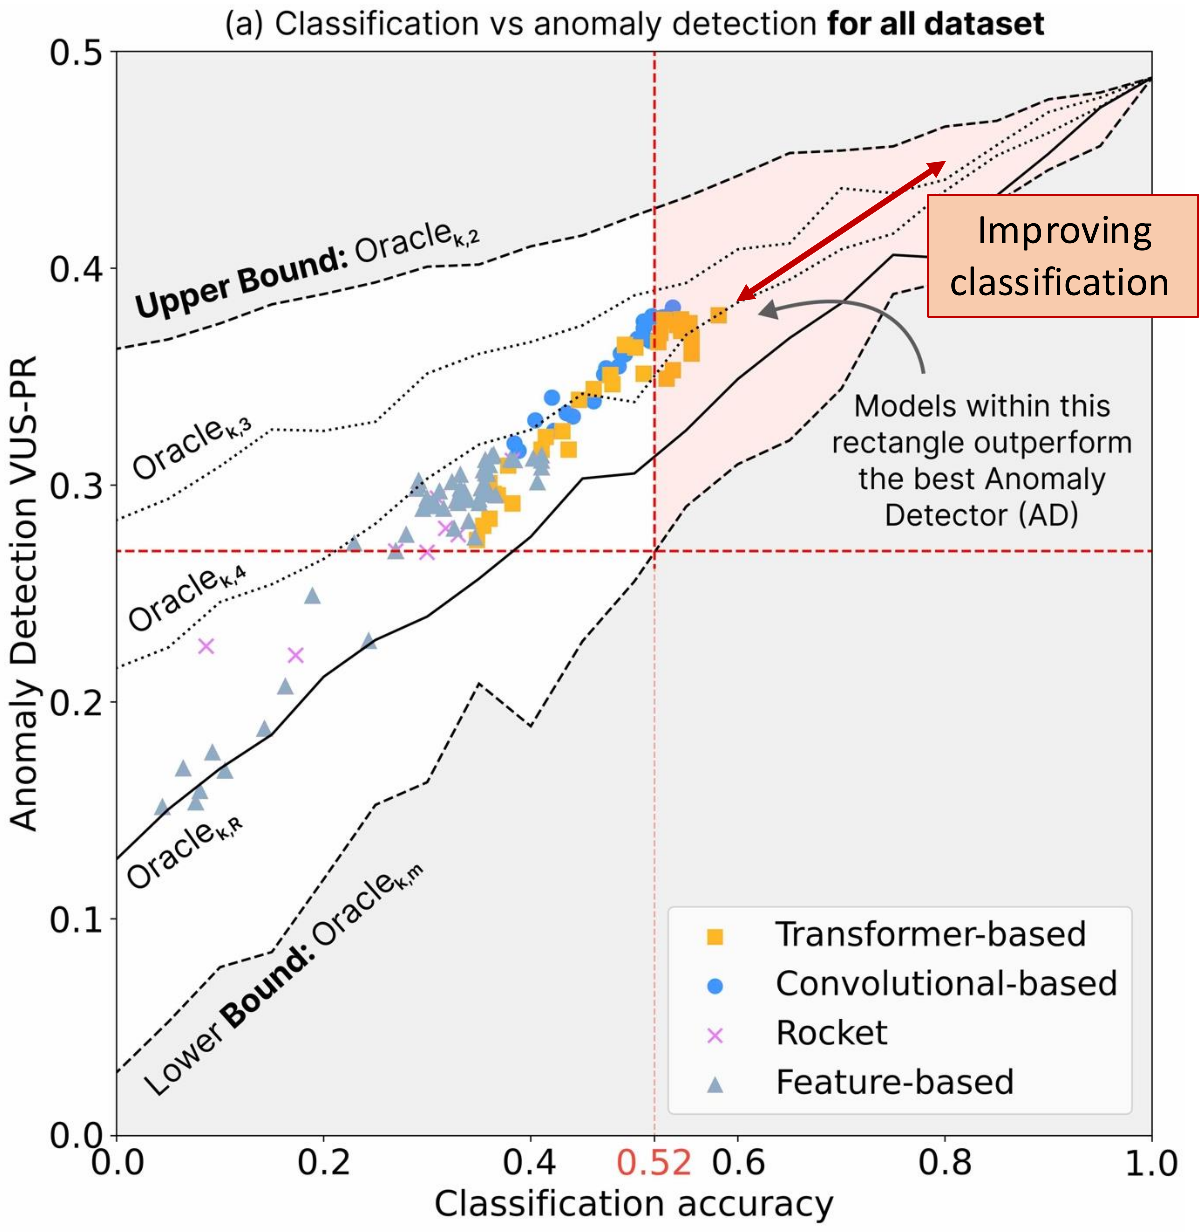

MSAD: *Experimental Evaluation*

- MS outperforms the **Individual detectors** and the **Avg Ens** in terms of accuracy
- MS outperforms **Avg Ens** in terms of execution time
- Potential improvement in terms of classification
- Potential improvement in terms of ranking detectors

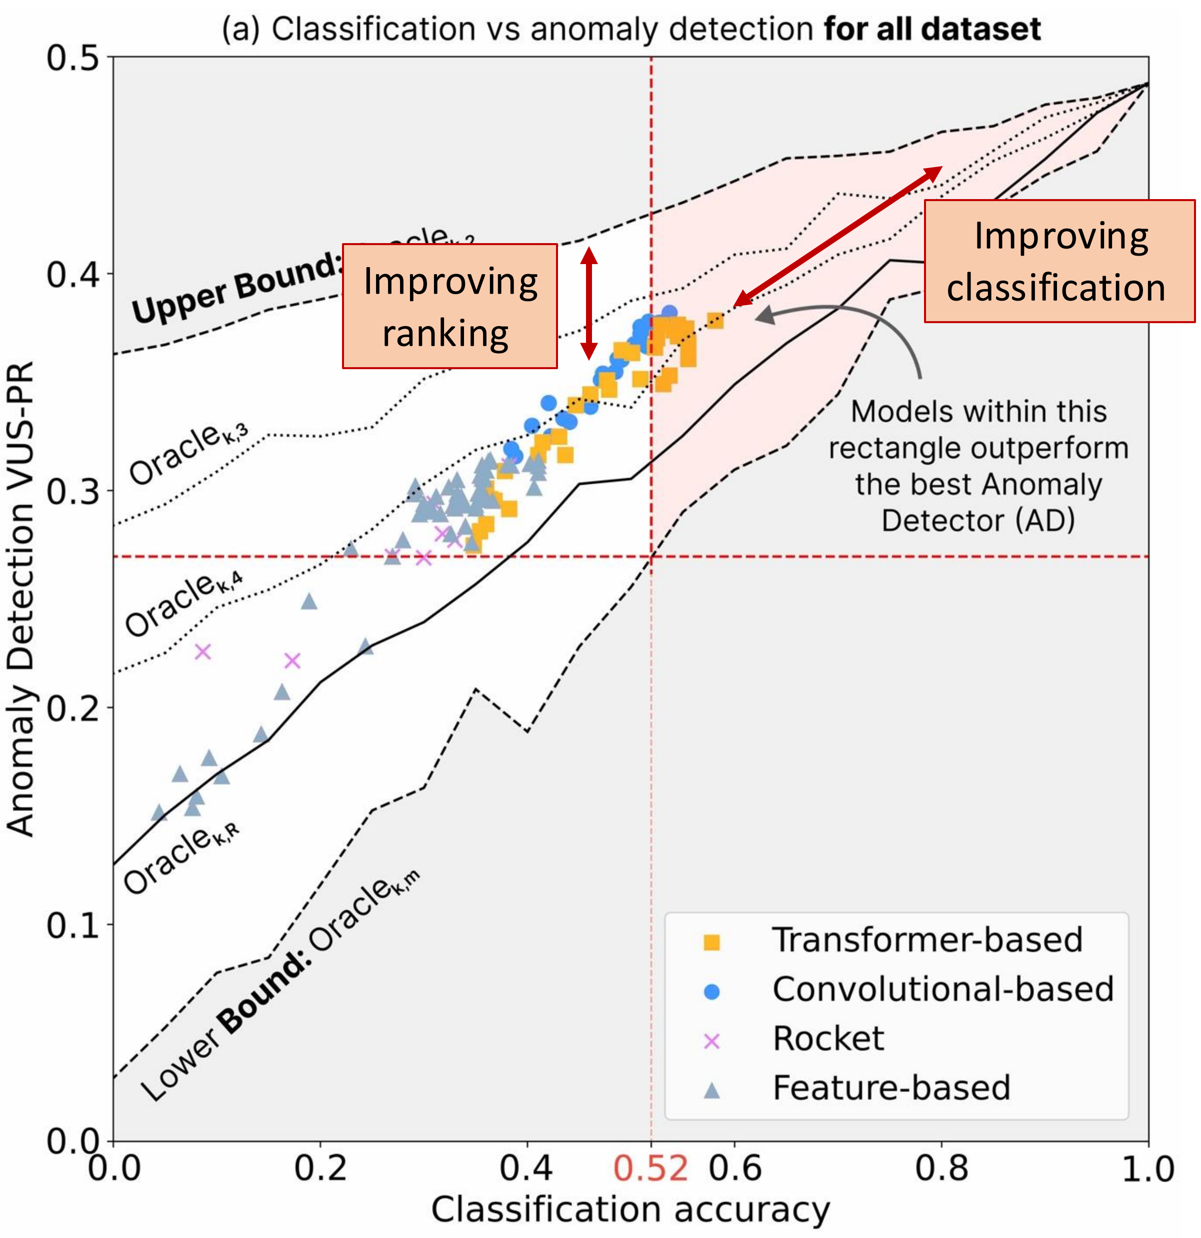

MSAD: *Experimental Evaluation*

Out-of-distribution testing: How well a model handles unfamiliar data?

MSAD: *Experimental Evaluation*

Out-of-distribution testing: How well a model handles unfamiliar data?

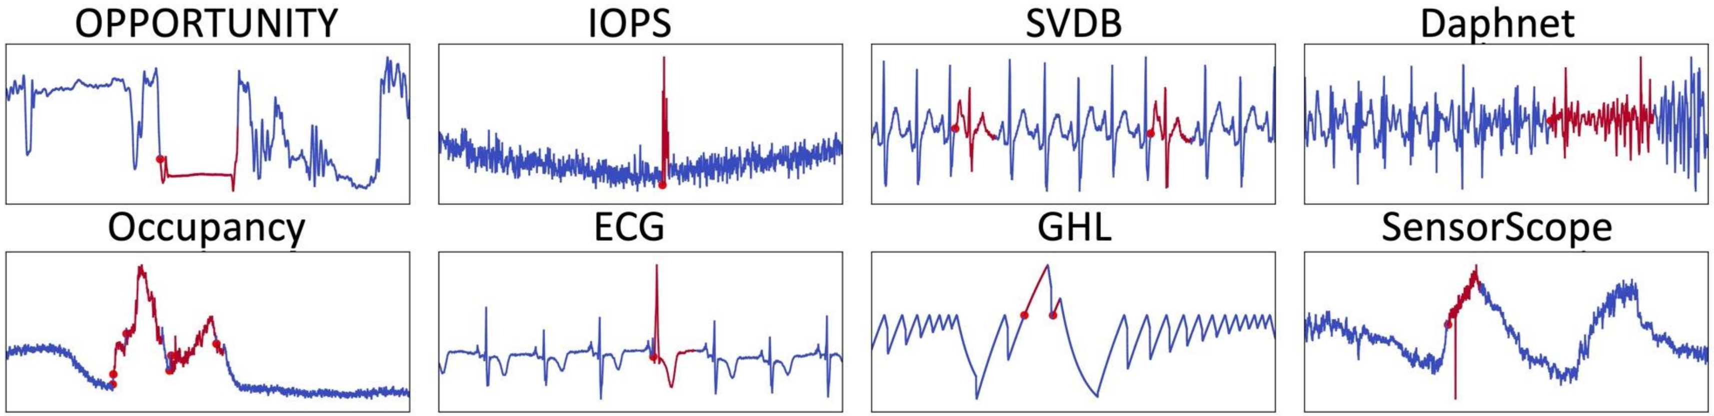

MSAD: *Experimental Evaluation*

Out-of-distribution testing: How well a model handles unfamiliar data?

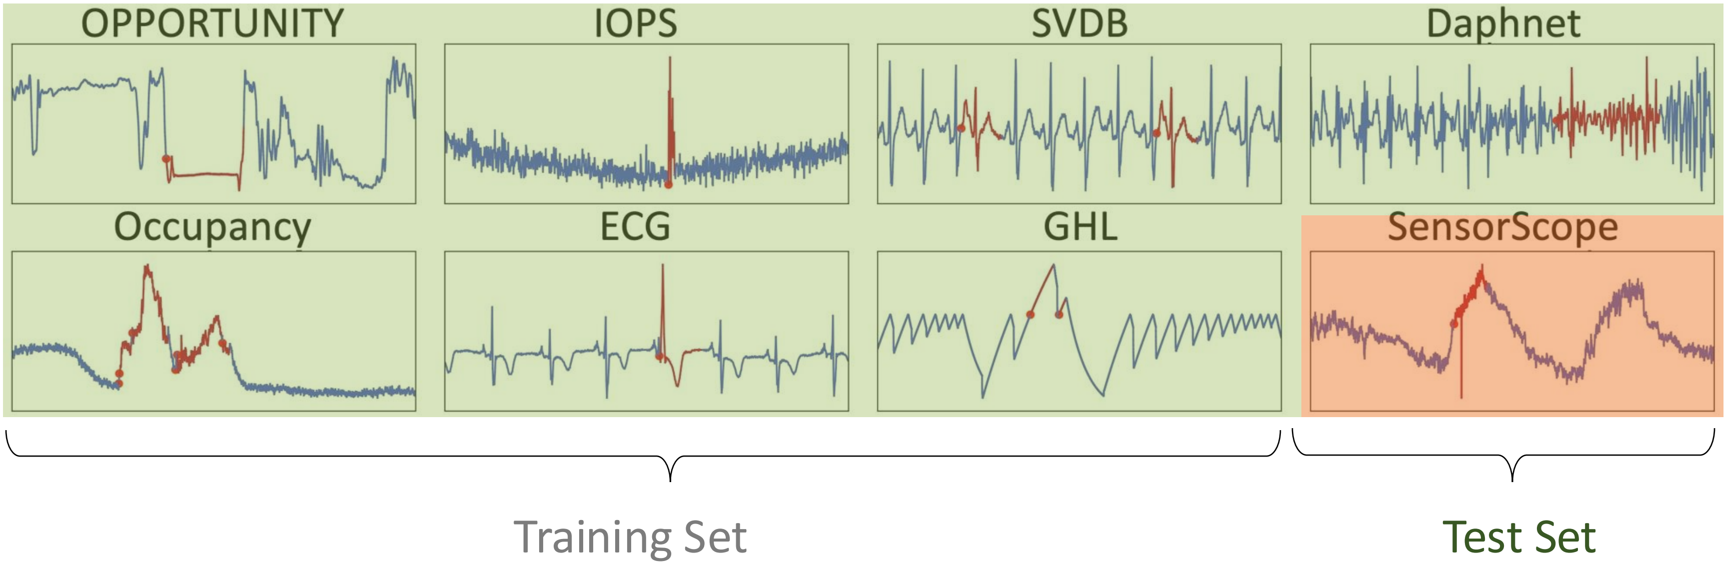

MSAD: *Experimental Evaluation*

Out-of-distribution testing: How well a model handles unfamiliar data?

MSAD: *Experimental Evaluation*

Out-of-distribution testing: How well a model handles unfamiliar data?

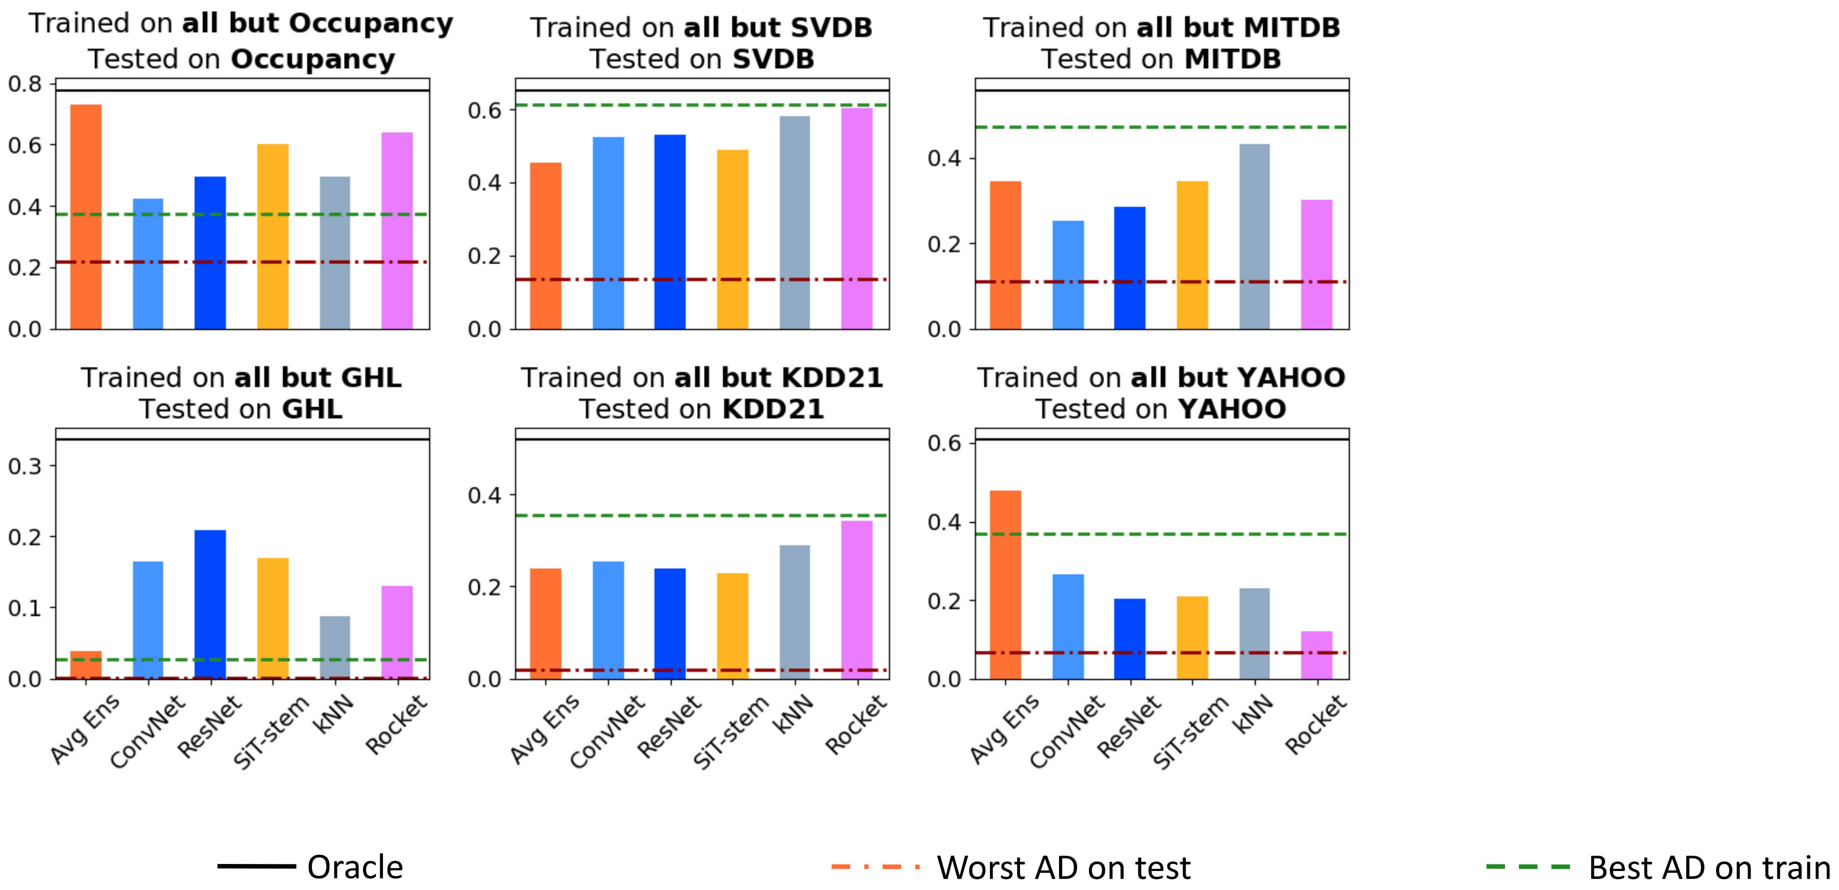

MSAD: *Experimental Evaluation*

Out-of-distribution testing: How well a model handles unfamiliar data?

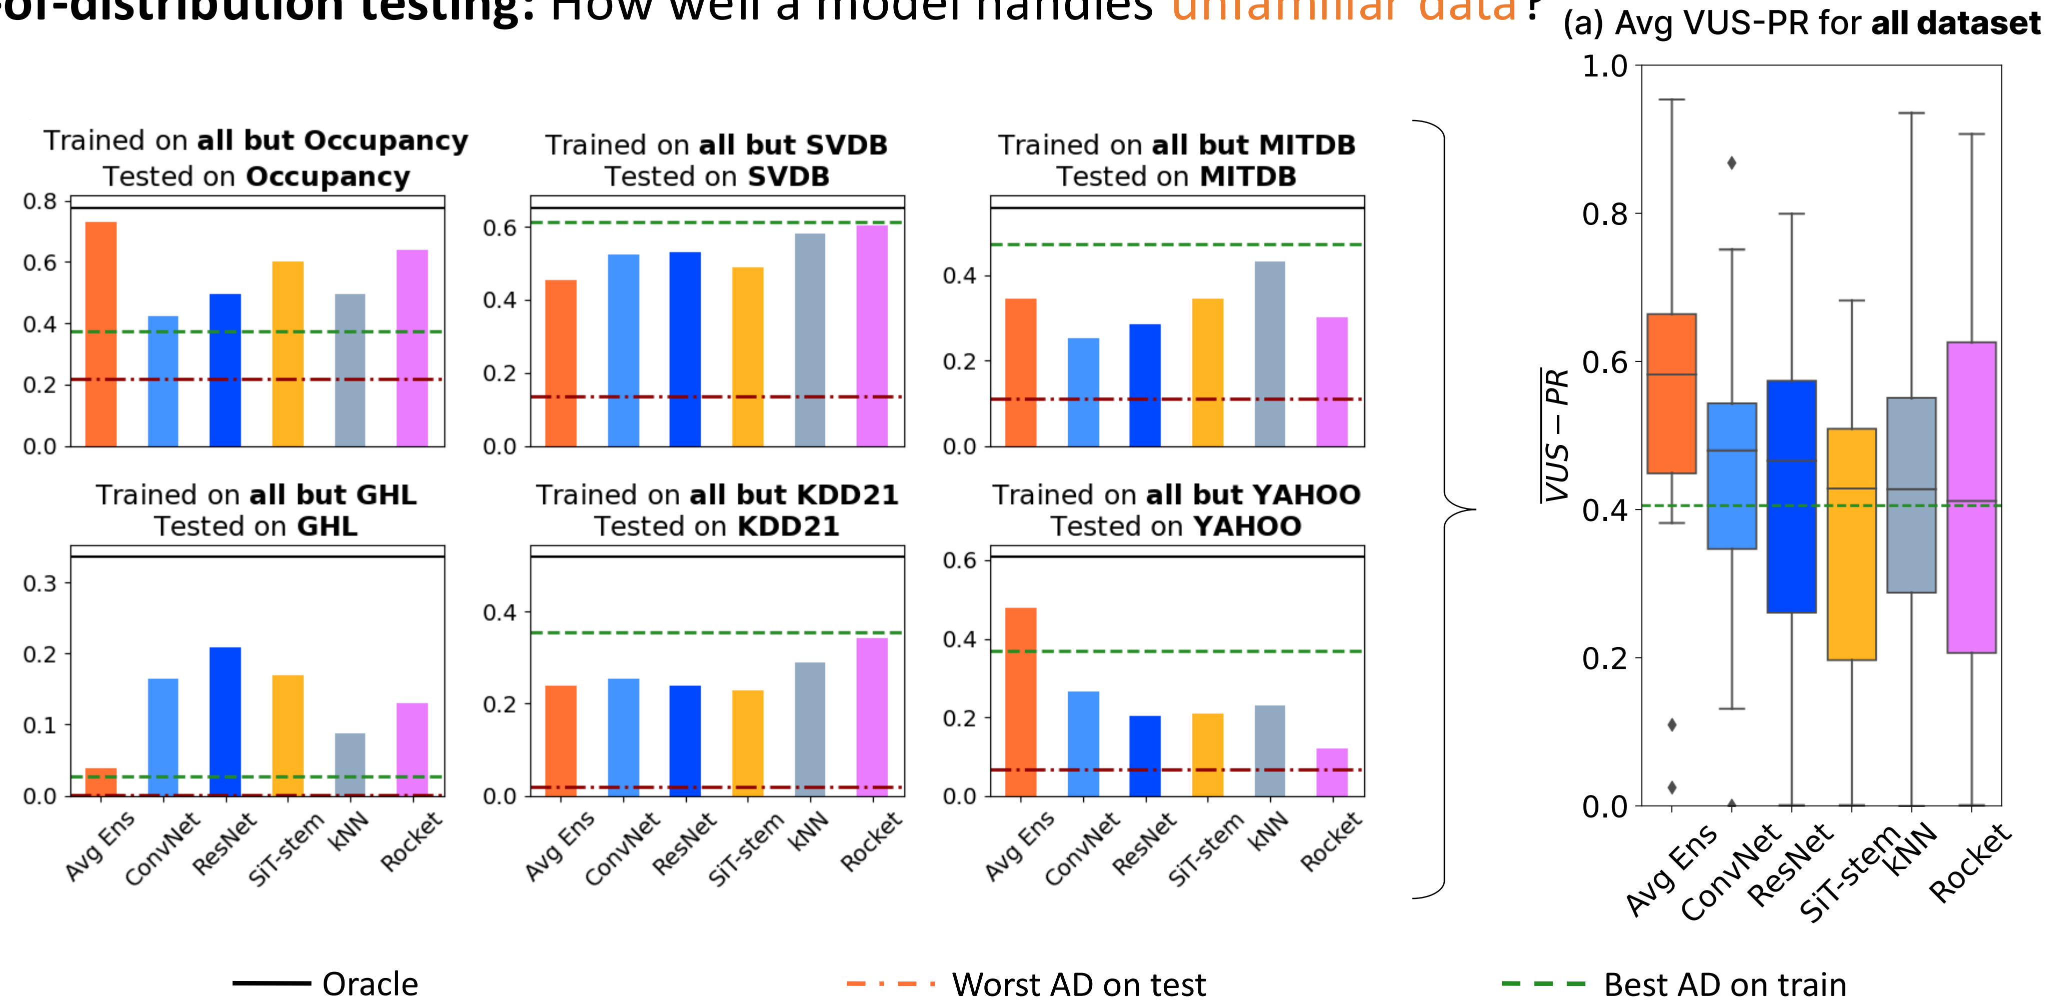

MSAD: *Experimental Evaluation*

Out-of-distribution testing: How well a model handles unfamiliar data?

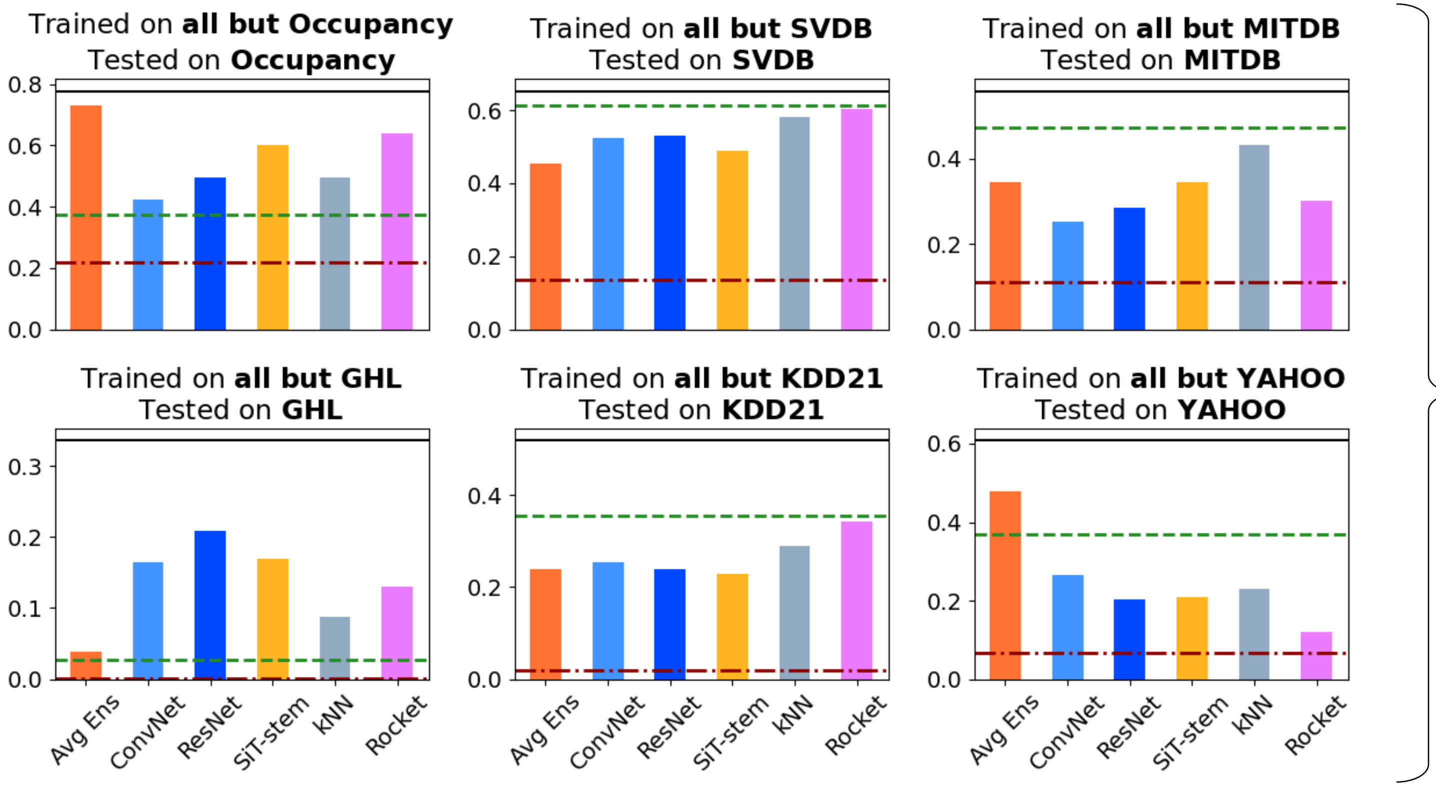

(a) Avg VUS-PR for **all dataset**

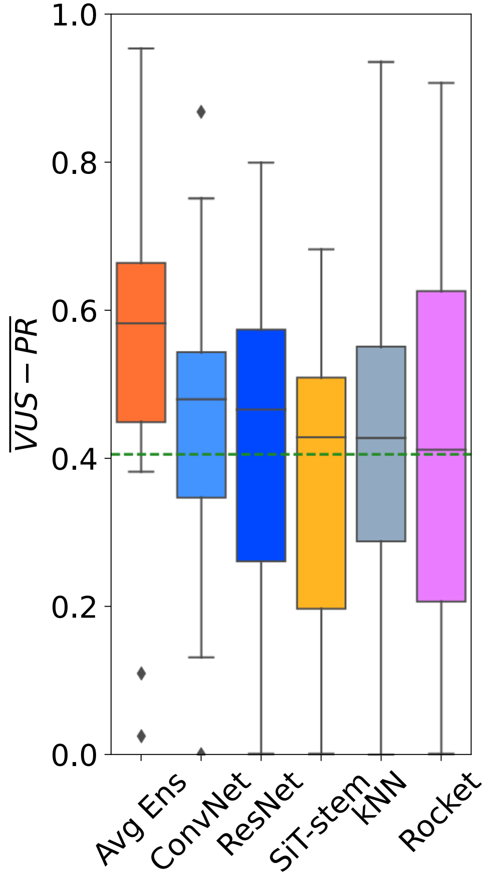

- **Avg Ens** is generally safer in terms of accuracy for new datasets

MSAD: *Experimental Evaluation*

Out-of-

Choose Wisely:

An Extensive Evaluation of Model Selection for Anomaly Detection in Time Series.

Emmanouil Sylligardos, Paul Boniol, John Paparrizos, Panos Trahanias, and Themis Palpanas.

VLDB2023 vancouver

Paper
(VLDB 2023)

https://helios2.mi.parisdescartes.fr/~themisp/publications/pvldb23-msad.pdf

ICDE Utrecht 2024

Demo
(ICDE 2024)

https://adecimots.streamlit.app/

GitHub Repo

boniolp/MSAD

➢ Avg

datasets

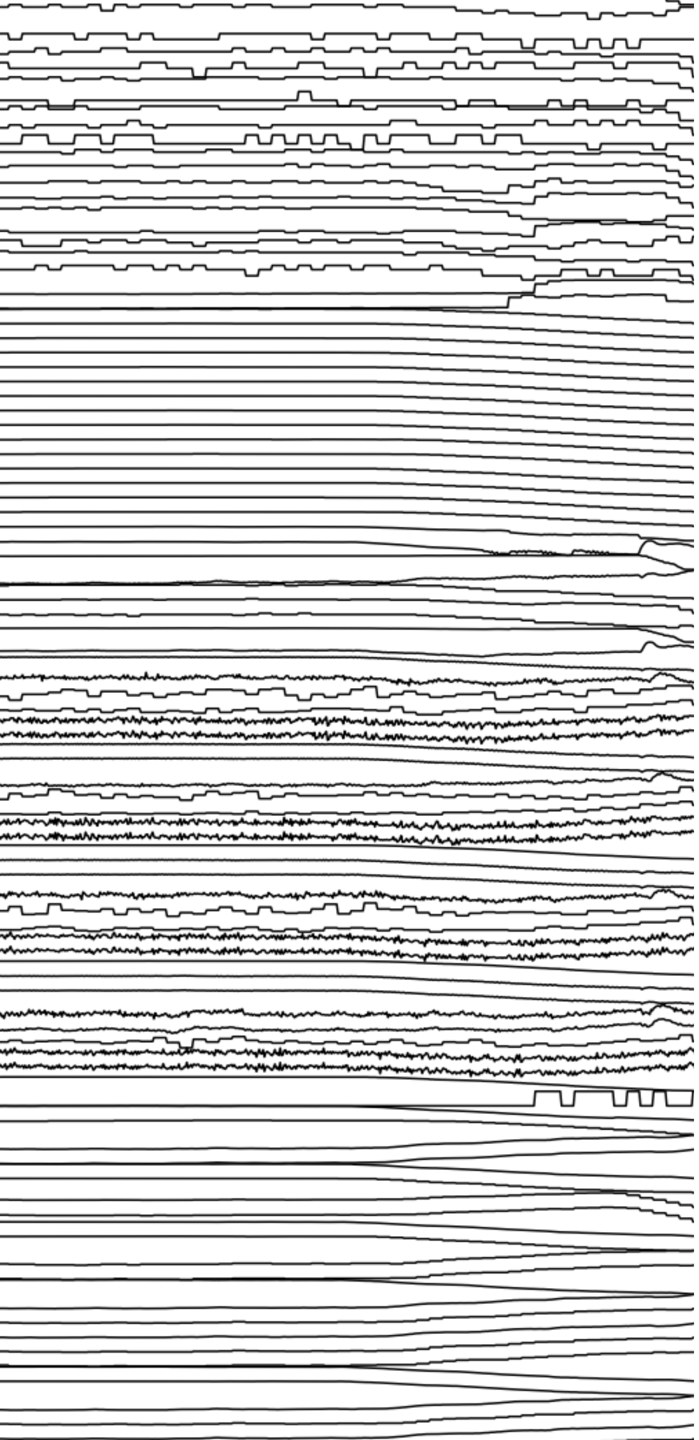

V. Conclusion

Research Directions

Conclusion: *Research Directions*

- Ensembling is still better for out-of-distribution cases

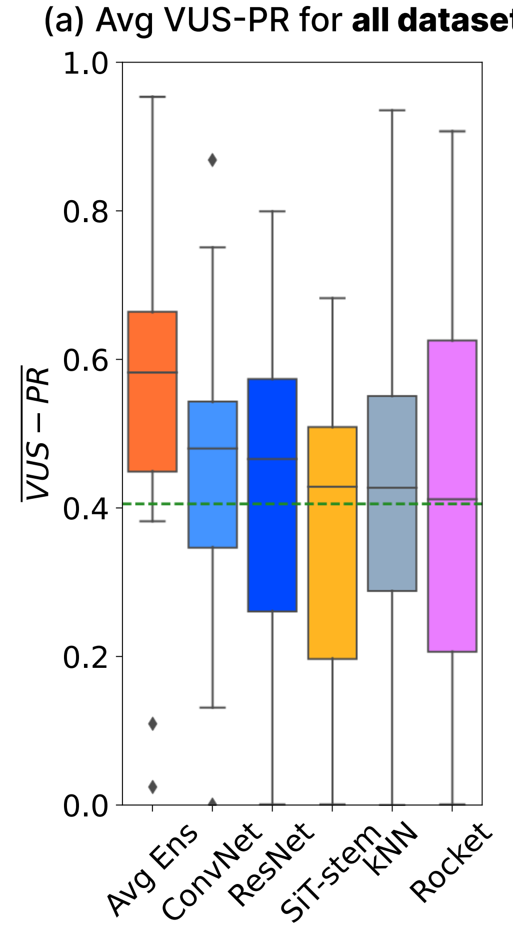

Conclusion: *Research Directions*

- Ensembling is still better for out-of-distribution cases
  - Combining Model Selection and Ensembling

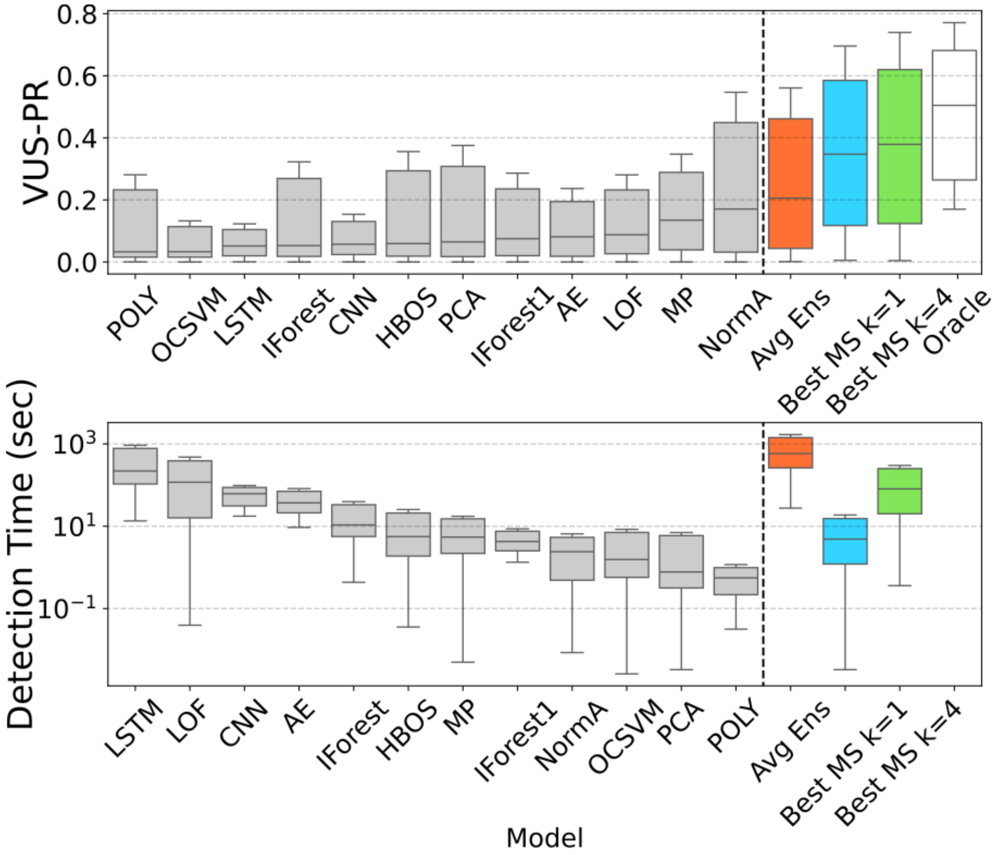

Conclusion: *Research Directions*

- Ensembling is still better for out-of-distribution cases
  - Combining Model Selection and Ensembling

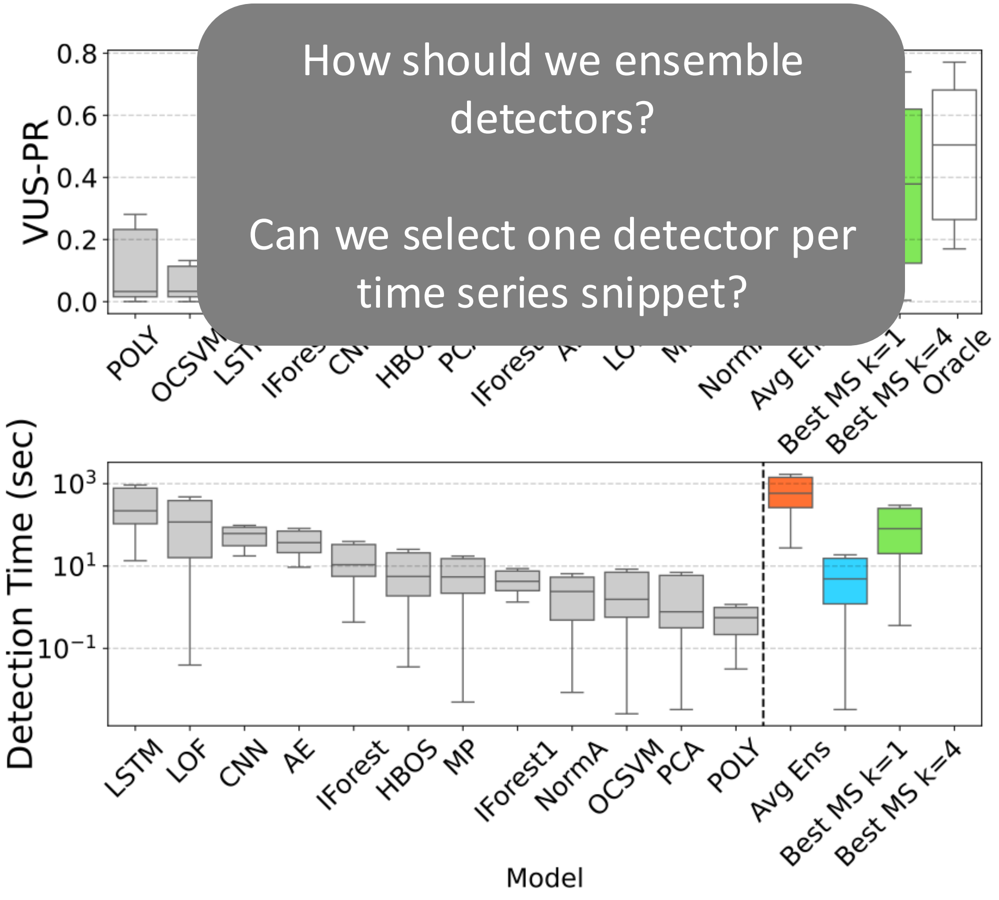

Conclusion: *Research Directions*

- Ensembling is still better for out-of-distribution cases
  - Combining Model Selection and Ensembling

- Ensembling has a strong impact on execution time

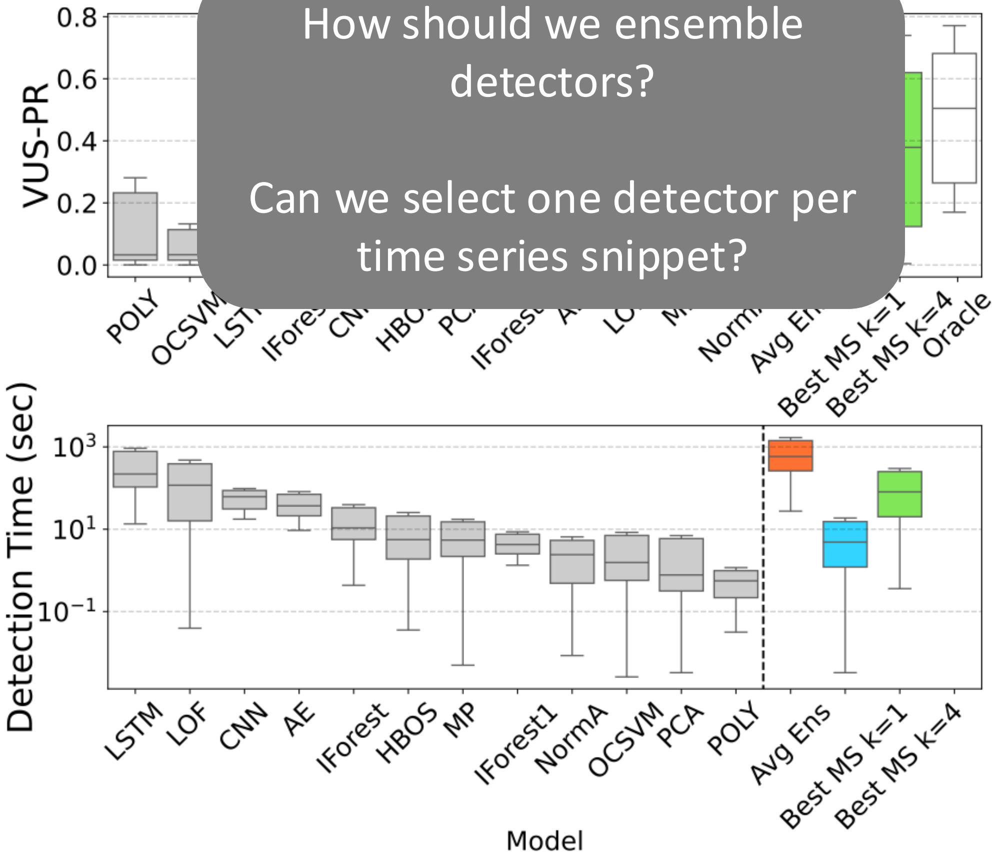

How should we ensemble detectors?

Can we select one detector per time series snippet?

Conclusion: *Research Directions*

- Ensembling is still better for out-of-distribution cases
  - Combining Model Selection and Ensembling

- Ensembling has a strong impact on execution time
  - Trade-off between execution time and accuracy in the selection process

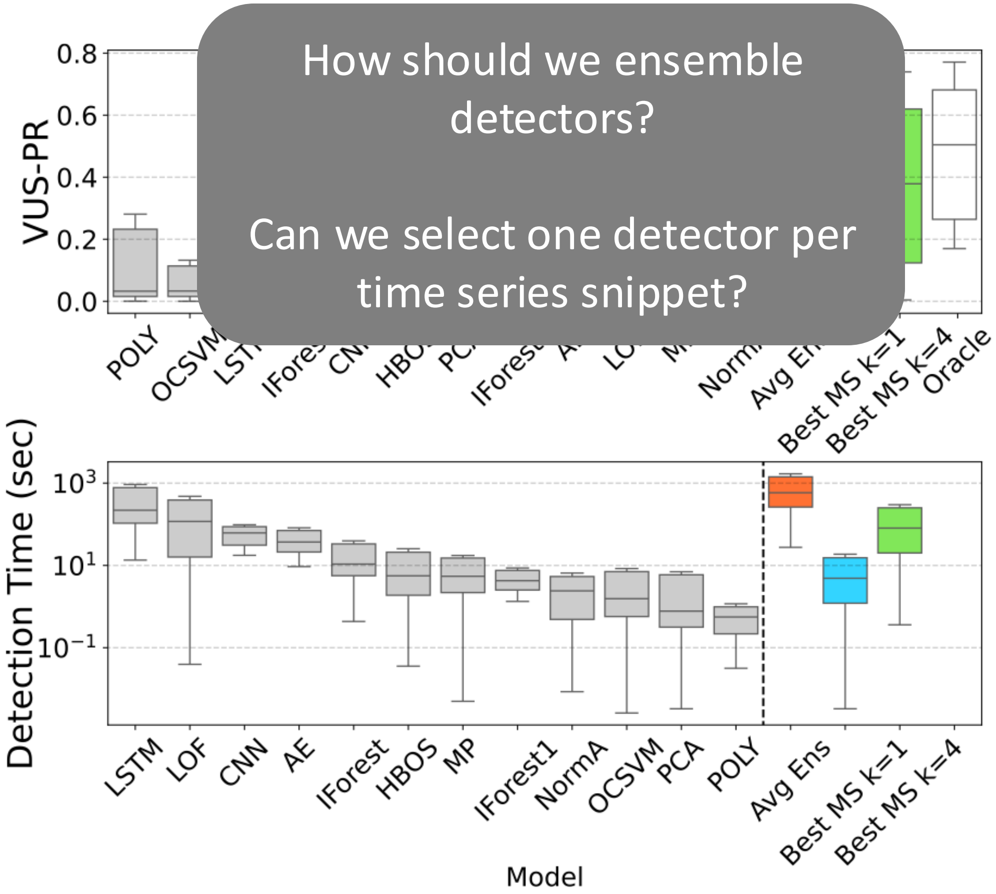

Conclusion: *Research Directions*

- Ensembling is still better for out-of-distribution cases
  - Combining Model Selection and Ensembling
- Ensembling has a strong impact on execution time
  - Trade-off between execution time and accuracy in the selection process
- Adding a new detector require training from scratch the pipeline

Classification-based

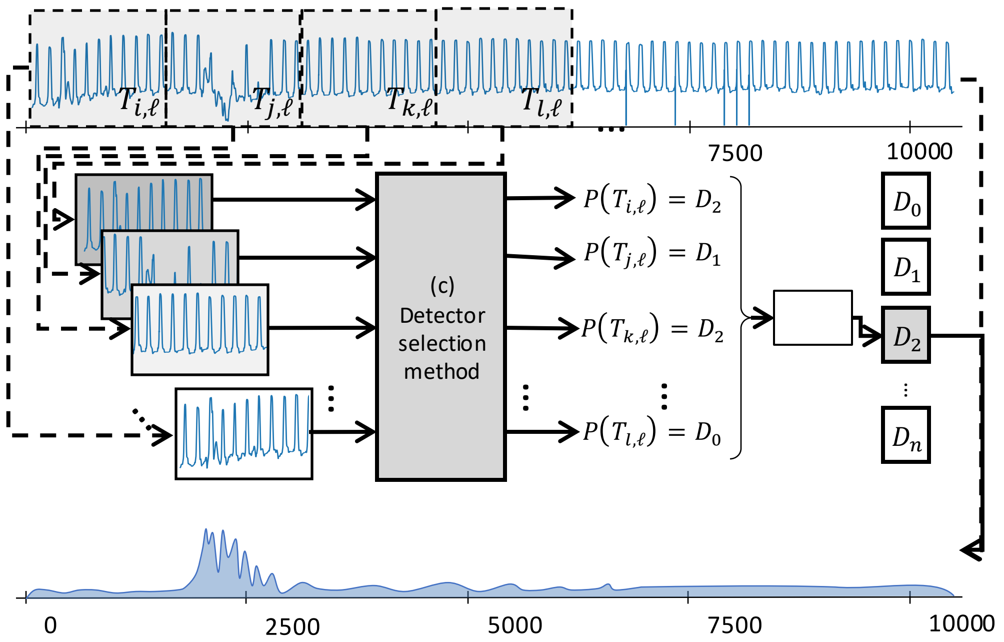

Conclusion: *Research Directions*

- Ensembling is still better for out-of-distribution cases
  - Combining Model Selection and Ensembling
- Ensembling has a strong impact on execution time
  - Trade-off between execution time and accuracy in the selection process
- Adding a new detector require training from scratch the pipeline
  - Improving modularity (regression-based model selection)

Regression-based

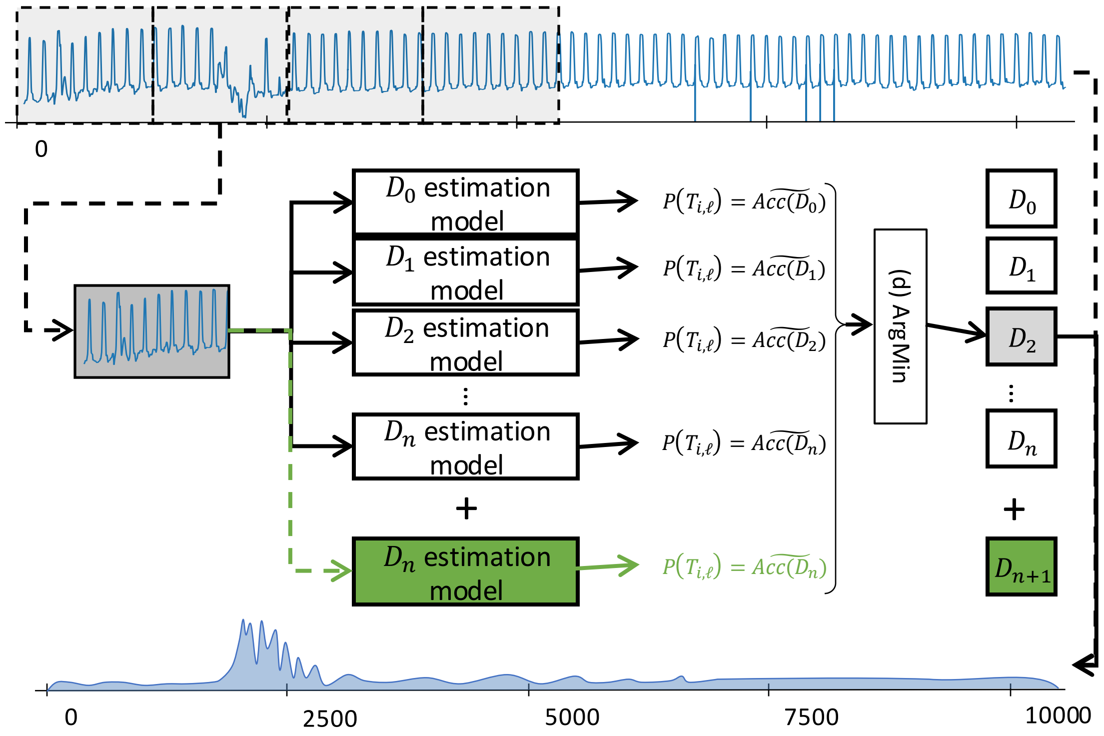

References

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. L. Lacasa et al., **From time series to complex networks: The visibility graph**, PNAS, 2008
7. M. Small et al. **Transforming Time seires into Complex Networks**, Complex Sciences (2009)
8. M. Small et al. **Transforming Time seires into Complex Networks**, Complex Sciences (2009)
9. Paul Boniol and Themis Palpanas. 2020. **Series2Graph: graph-based subsequence anomaly detection for time series.** Proc. VLDB Endow. 13, 12 (August 2020), 1821–1834.
10. 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
11. Schneider, J., Wenig, P. & Papenbrock, T. **Distributed detection of sequential anomalies in univariate time series**. The VLDB Journal **30**, 579–602 (2021).
12. Wenig, P., Papenbrock, T. (2024). **Series2Graph++: Distributed Detection of Correlation Anomalies in Multivariate Time Series**. In: Wrembel, R., Chiusano, S., Kotsis, G., Tjoa, A.M., Khalil, I. (eds) Big Data Analytics and Knowledge Discovery. DaWaK 2024
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. John Paparrizos, et al. 2022. **TSB-UAD: An End-to-End Benchmark Suite for Univariate Time-Series Anomaly Detection**. Proc. VLDB Endow. 15, 8
15. Mononito Goswami, Cristian Challu, Laurent Callot, Lenon Minorics, Andrey Kan. 2023. **Unsupervised Model Selection for Time-series Anomaly Detection**. In Proceedings of the International Conference on Learning Representations.
16. Lei Cao, Yizhou Yan, Yu Wang, Samuel Madden, Elke A Rundensteiner. 2023. **Autood: Automatic outlier detection**. Proceedings of the ACM on Management of Data, 1(1): 1-27. ACM, New York, NY, USA.
17. Charu C. Aggarwal and Saket Sathe. 2015. **Theoretical Foundations and Algorithms for Outlier Ensembles**. SIGKDD Explor. Newsl. 17, 1

People that contributed to this presentation

Prof. Themis Palpanas
Université Paris Cité; IUF

Mohammed Meftah
EDF R&D

Prof. Panos Trahanias
ICS-FORTH

Prof. John Paprrizos
Ohio State University

Prof. Michael Franklin
University of Chicago

Emmanouil Sylligardos
Ecole Normale Supérieure

Qinghua Liu
Ohio State University

And many others…

Thank you for attending!