Detection of Anomalies and Identification of their Precursors in Large Data Series Collections

Presented by Paul Boniol

Supervised by:

Prof. Themis Palpanas, Mohammed Meftah, Emmanuel Remy

1. Introduction

1. Introduction: *data series*

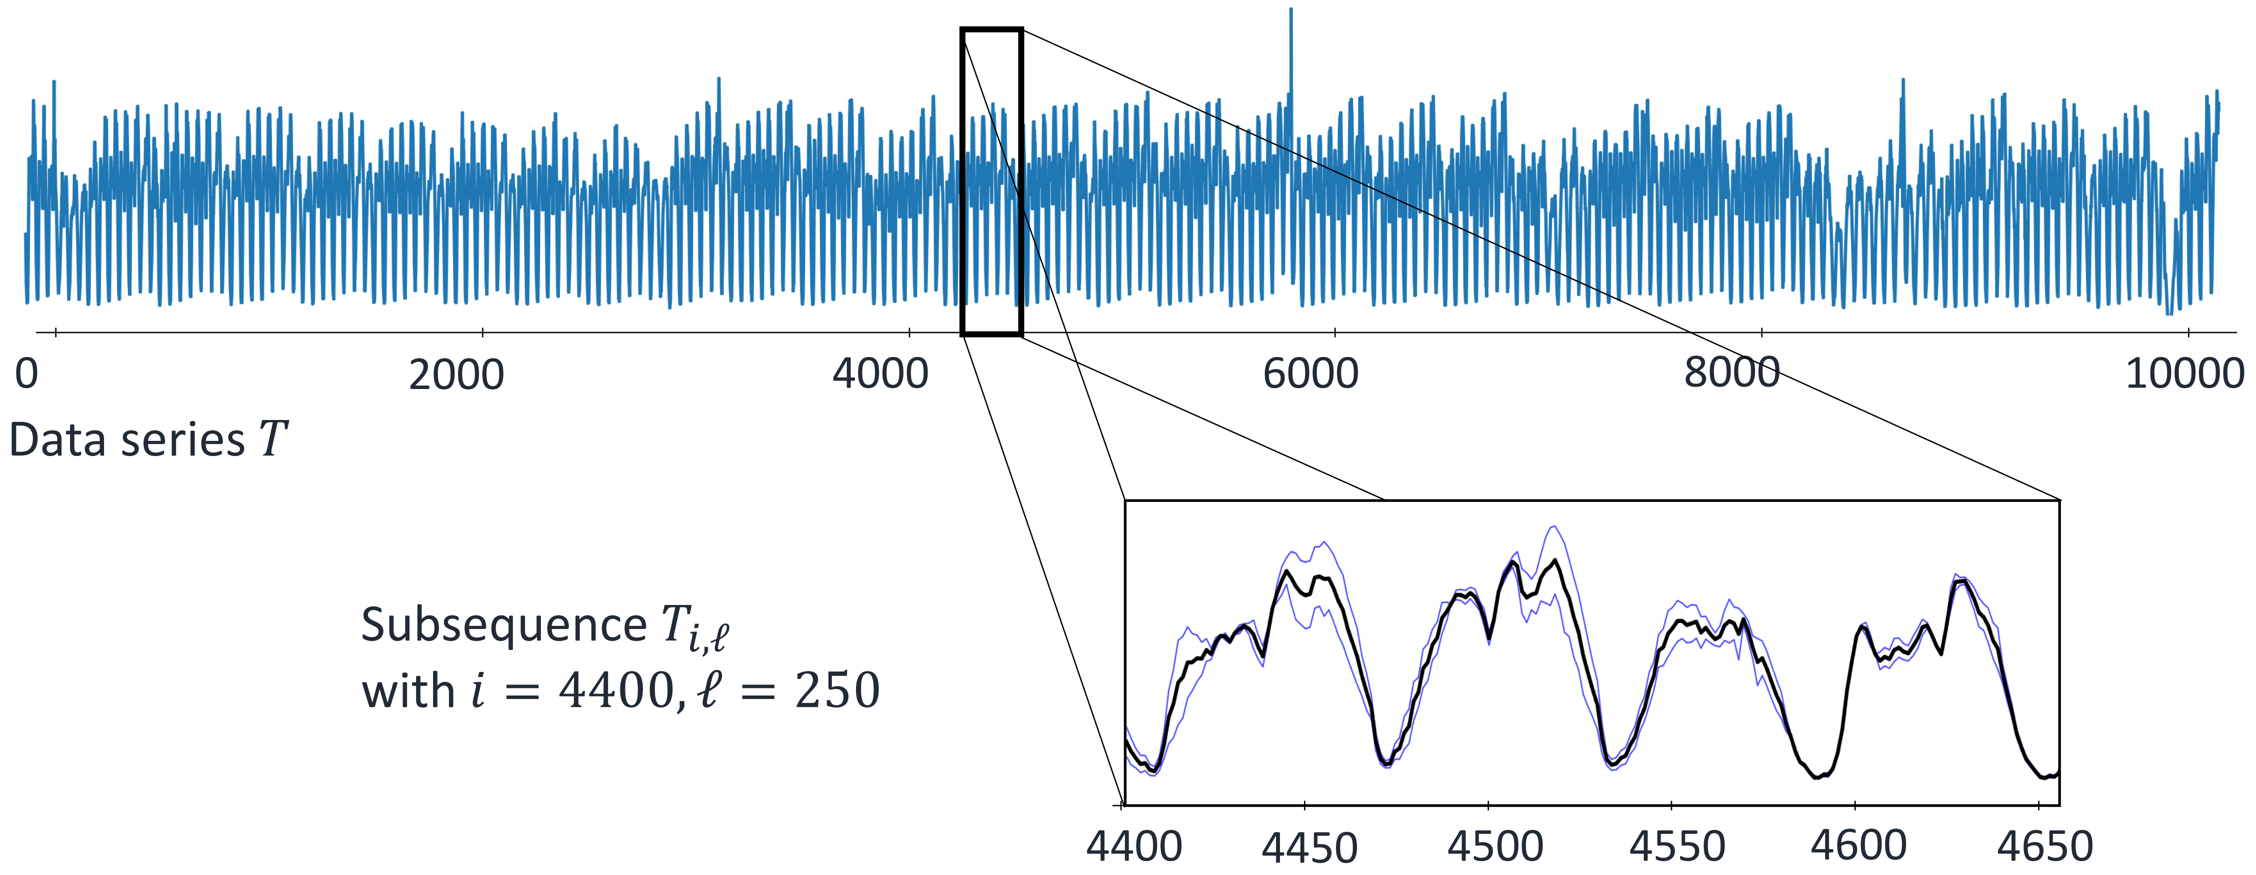

1. Introduction: *anomaly detection*

Anomaly: a rare subsequence (and potentially not desired) of a given length ℓ
(for instance, 10 points, or one hour)

ECG (Electrocardiogram) [1]:
Premature ventricular contraction

Dutch energy consumption [1]:
Unusual days of the year

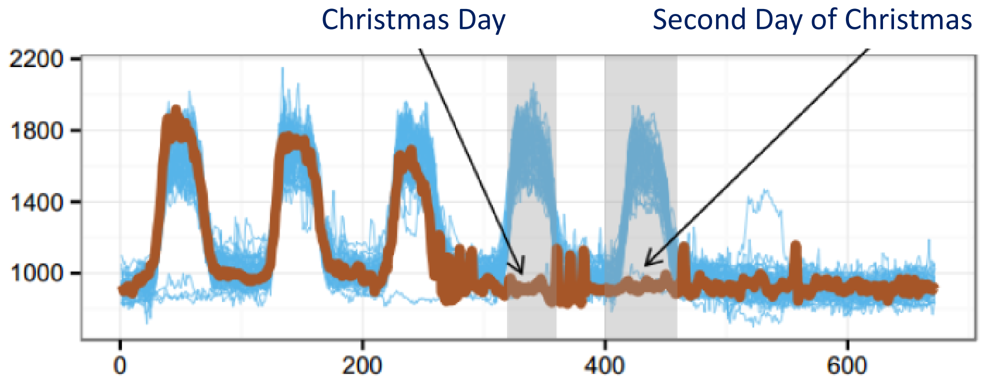

[1] E. Keogh et al., HOT SAX: Finding the Most Unusual Time Series Subsequence: Algorithm and Applications, ICDM, 2005

1. Introduction: *research directions*

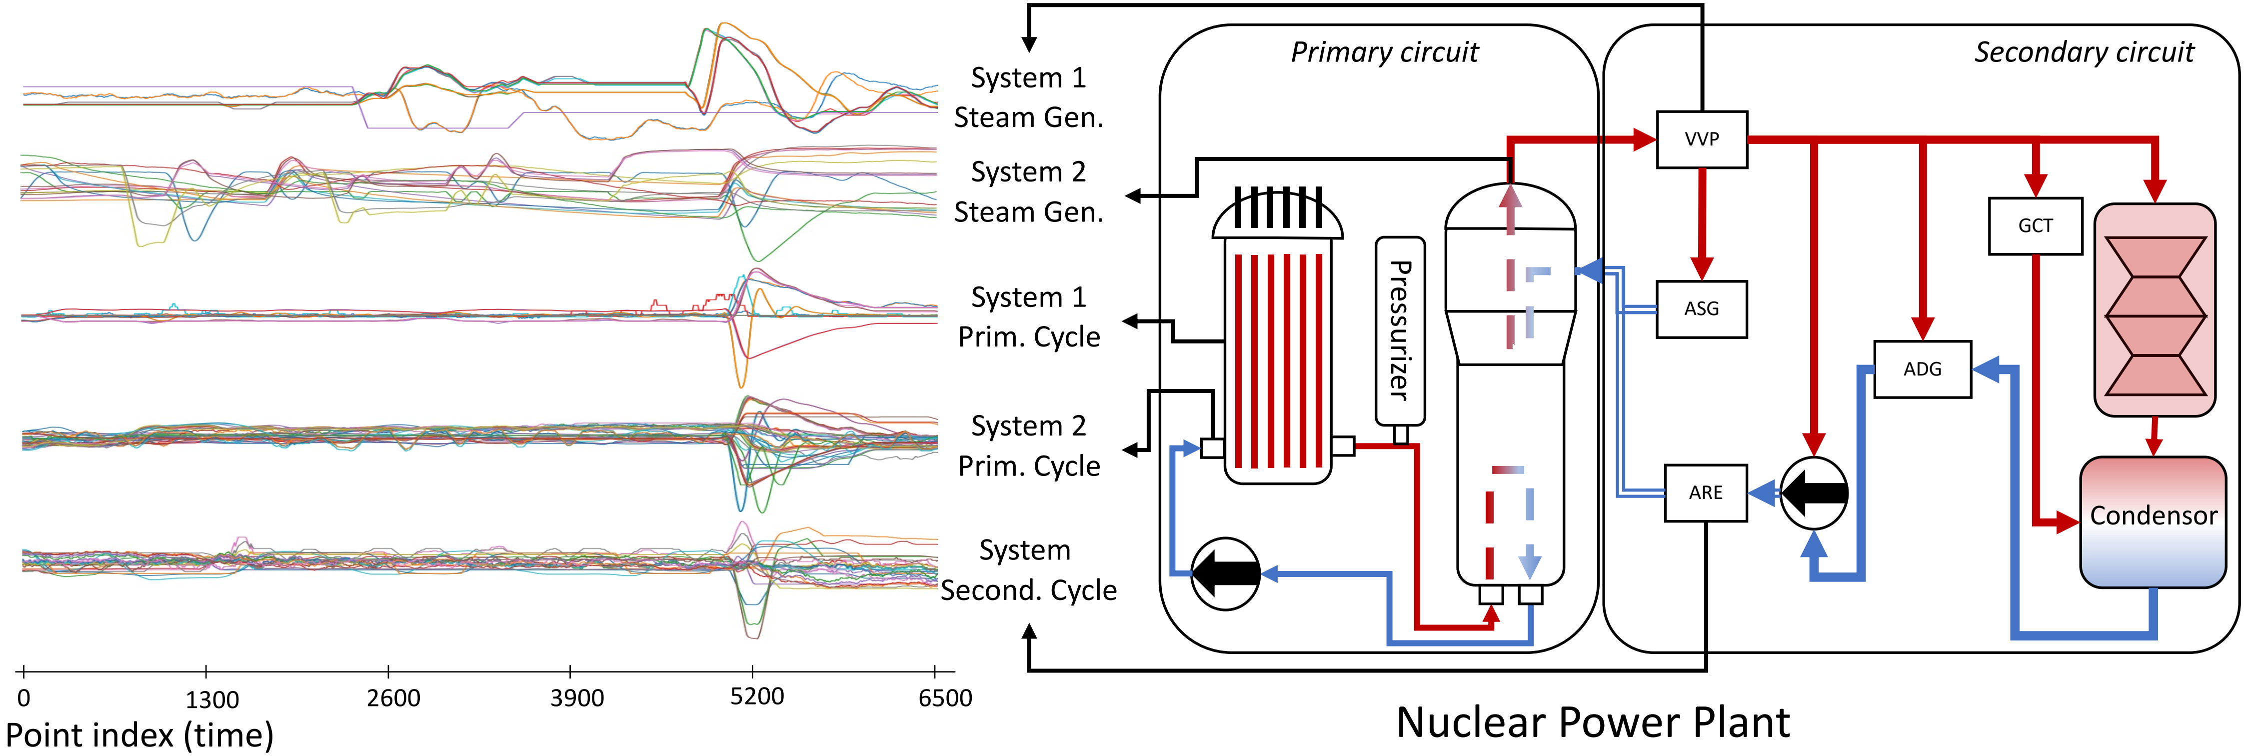

1. Introduction: *research directions*

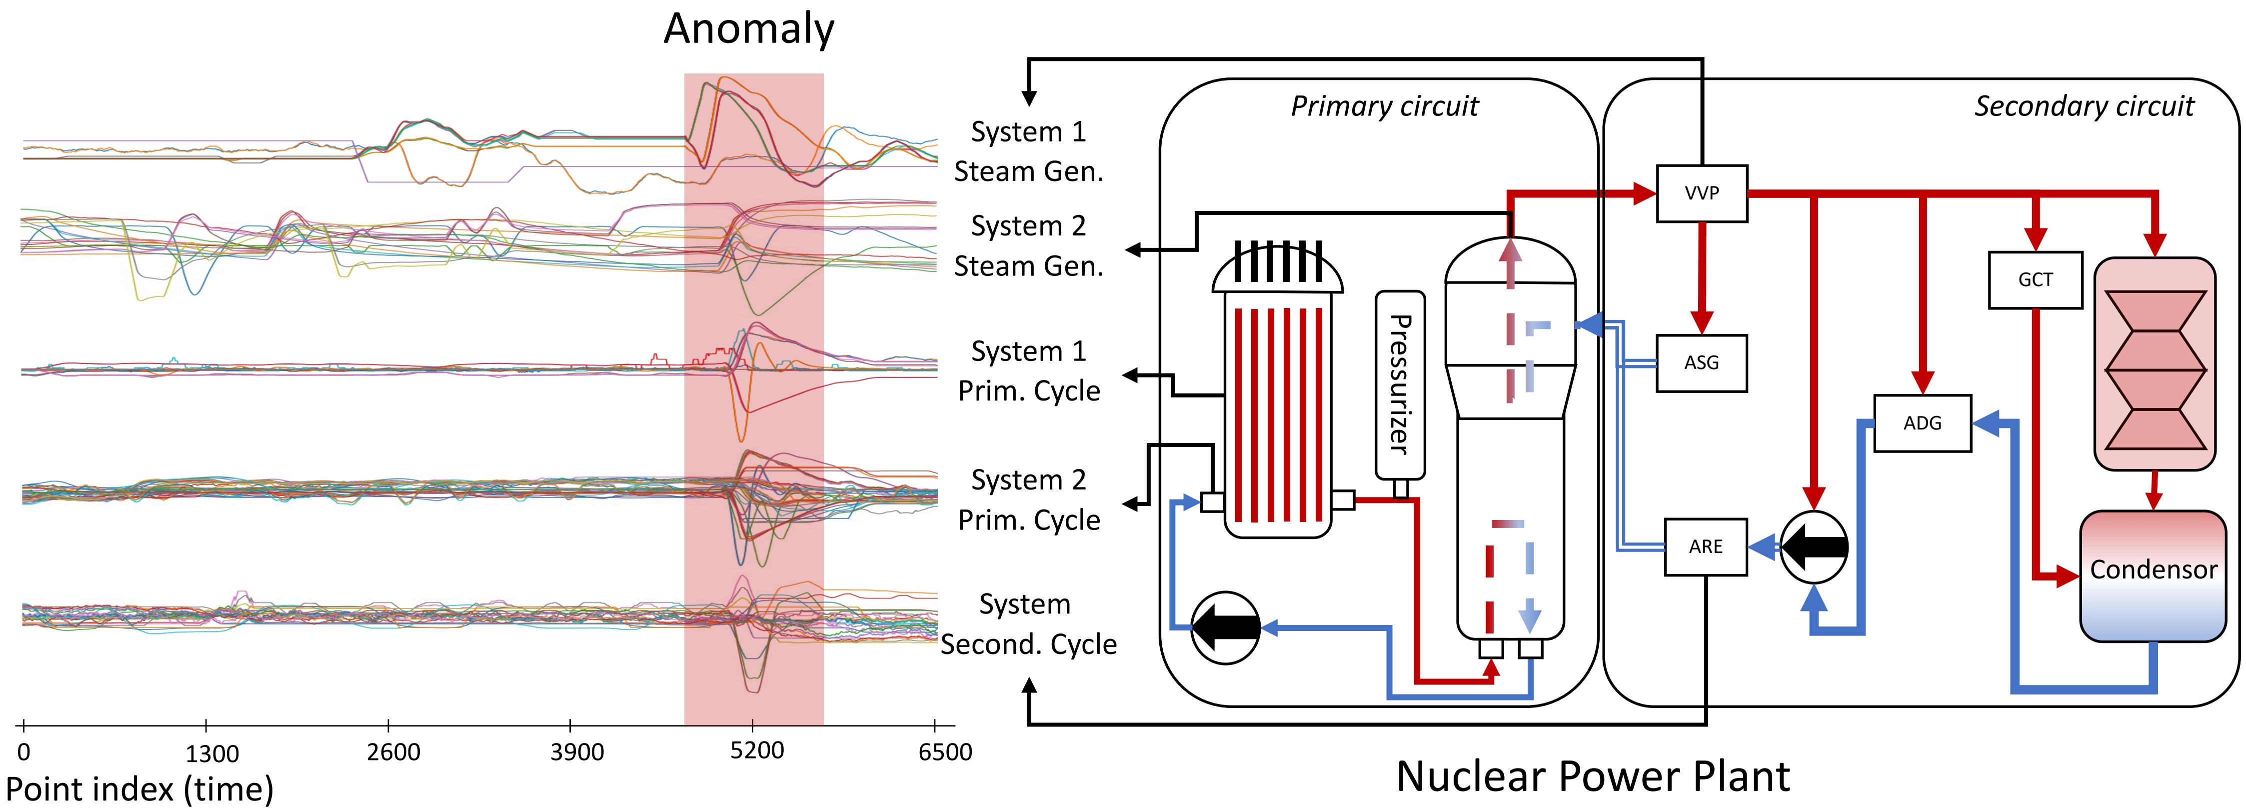

Nuclear Power Plant

Unsupervised subsequence anomalies detection

- Detect unknown anomalies in power plants sensors
- Raise real time alarms to assist expert decisions

1. Introduction: *research directions*

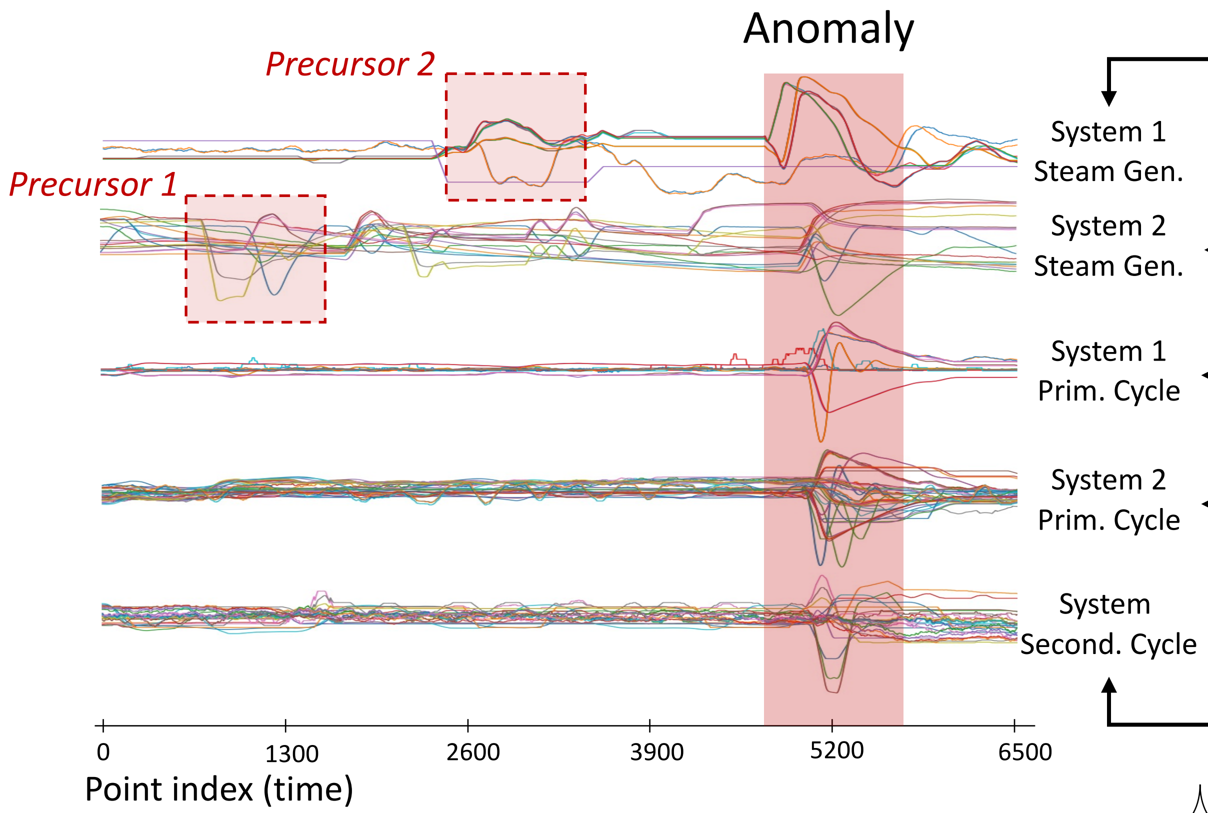

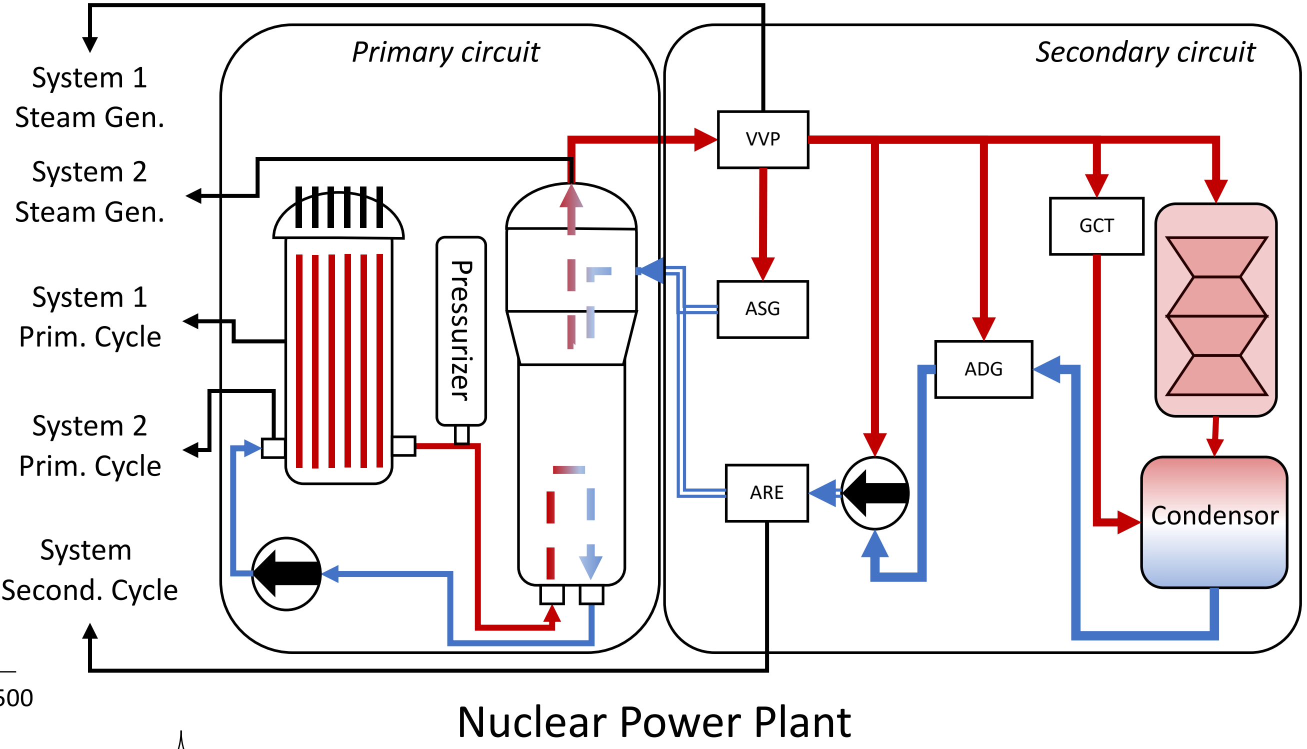

Unsupervised subsequence anomalies detection

- Detect unknown anomalies in power plants sensors
- Raise real time alarms to assist expert decisions

Supervised detection of anomaly precursors

- Classify known anomalies
- Detect precursors (discriminant subsequences) of known anomalies in the power plant

1. Introduction: *research directions*

"Detection of Anomalies and Identification of their Precursors in Large Data Series Collections"

Unsupervised subsequence anomalies detection | *Supervised identification of anomaly precursors*

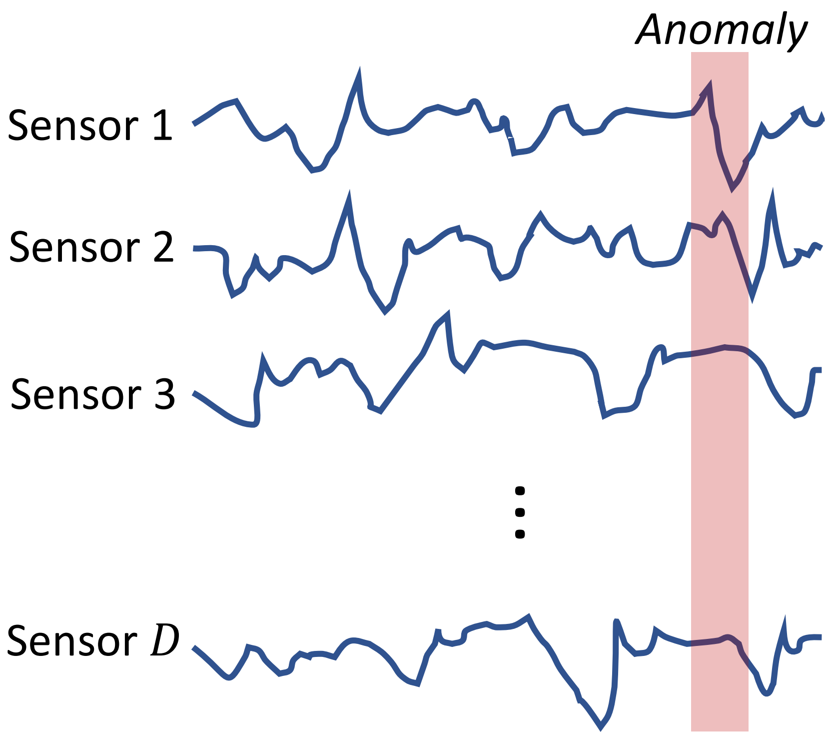

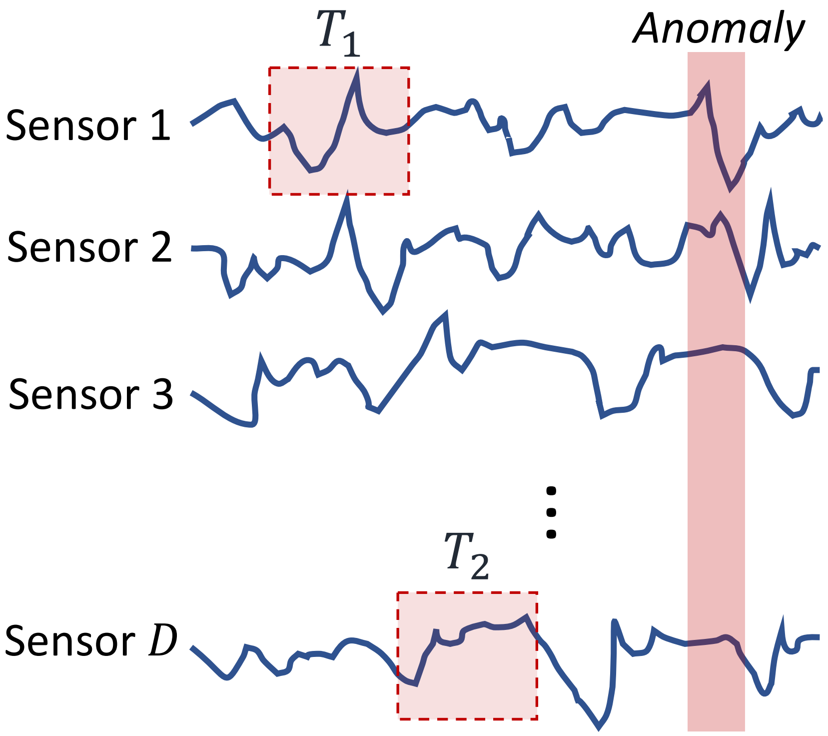

2. Unsupervised detection of abnormal subsequences

2.1. Series2Graph: *Background*

Graph G_{ℓ_G} [2]:

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

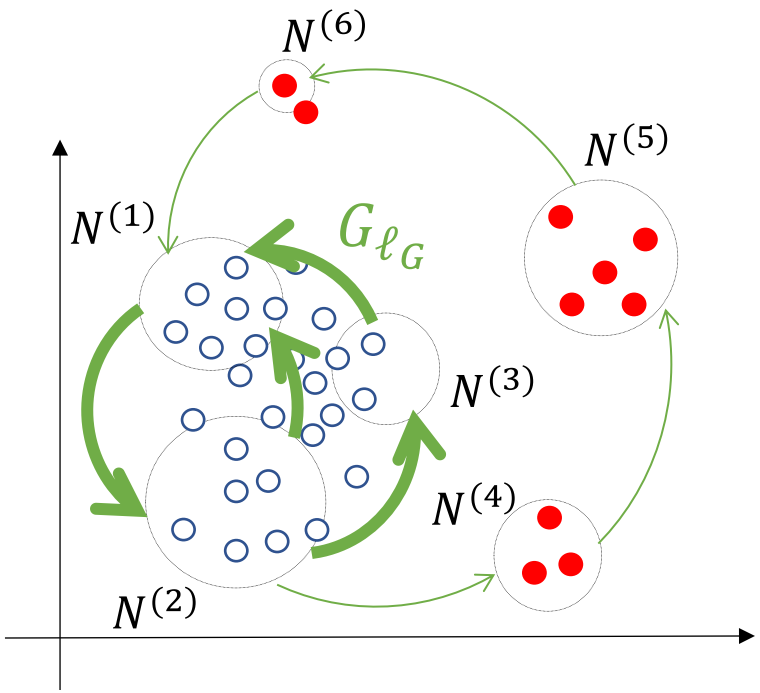

$\mathcal{N}$ is a set of nodes. Each node is an ensemble of similar subsequences.

$\mathcal{E}$ is a set of edges. Each edge (between two nodes i and j) is associated to a weight w that corresponds to the number of times a subsequence of node j follows a subsequence of node i.

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

2.1. Series2Graph: *Background*

Graph G_{ℓ_G} [2]:
Given a data series T, and an input length ℓ_G, we build a graph $G_{\ell_G}(\mathcal{N}, \mathcal{E})$ for which:

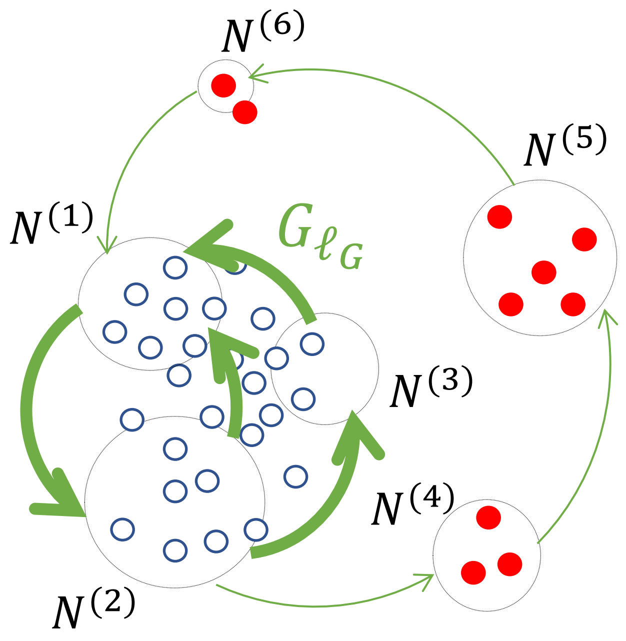

For a given subsequence $T_{i,\ell}$ and its corresponding path
$P_{th} =< N^{(i)}, N^{(i+1)}, \dots, N^{(i+\ell)} >$,
we define the normality score as follows:

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

2.1. Series2Graph: *Background*

Graph G_{ℓ_G} [2]:
Given a data series T, and an input length ℓ_G, we build a graph $G_{\ell_G}(\mathcal{N},\mathcal{E})$ for which:

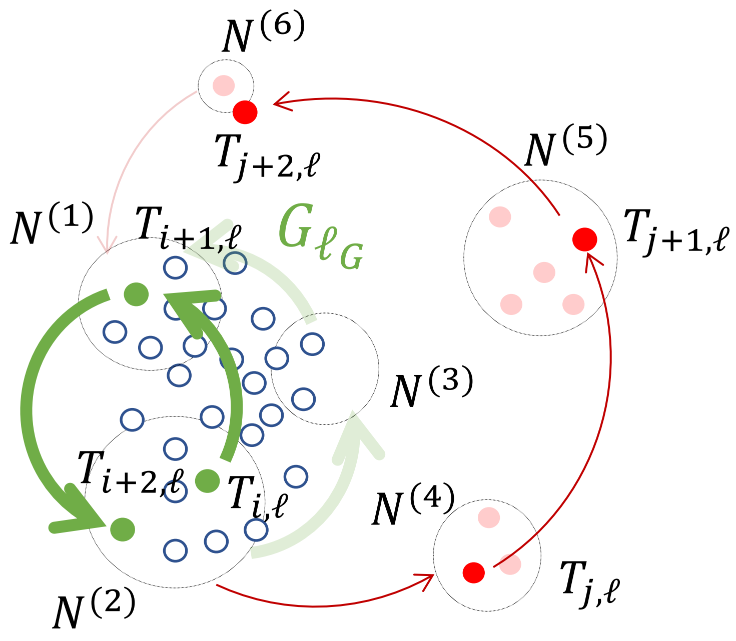

The subsequence $T_{i,\ell+2}$ is a path in the graph
$P_{th}(T_{i,\ell+2}) = < T_{i,\ell}, T_{i+1,\ell}, T_{i+2,\ell} > = < N^{(4)}, N^{(5)}, N^{(6)} >$ in G_{ℓ_G}.

The subsequence $T_{j,\ell+2}$ is a path in the graph
$P_{th}(T_{j,\ell+2}) = < T_{j,\ell}, T_{j+1,\ell}, T_{j+2,\ell} > = < N^{(2)}, N^{(1)}, N^{(2)} >$ in G_{ℓ_G}.

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

2.2. Series2Graph: *Computation steps*

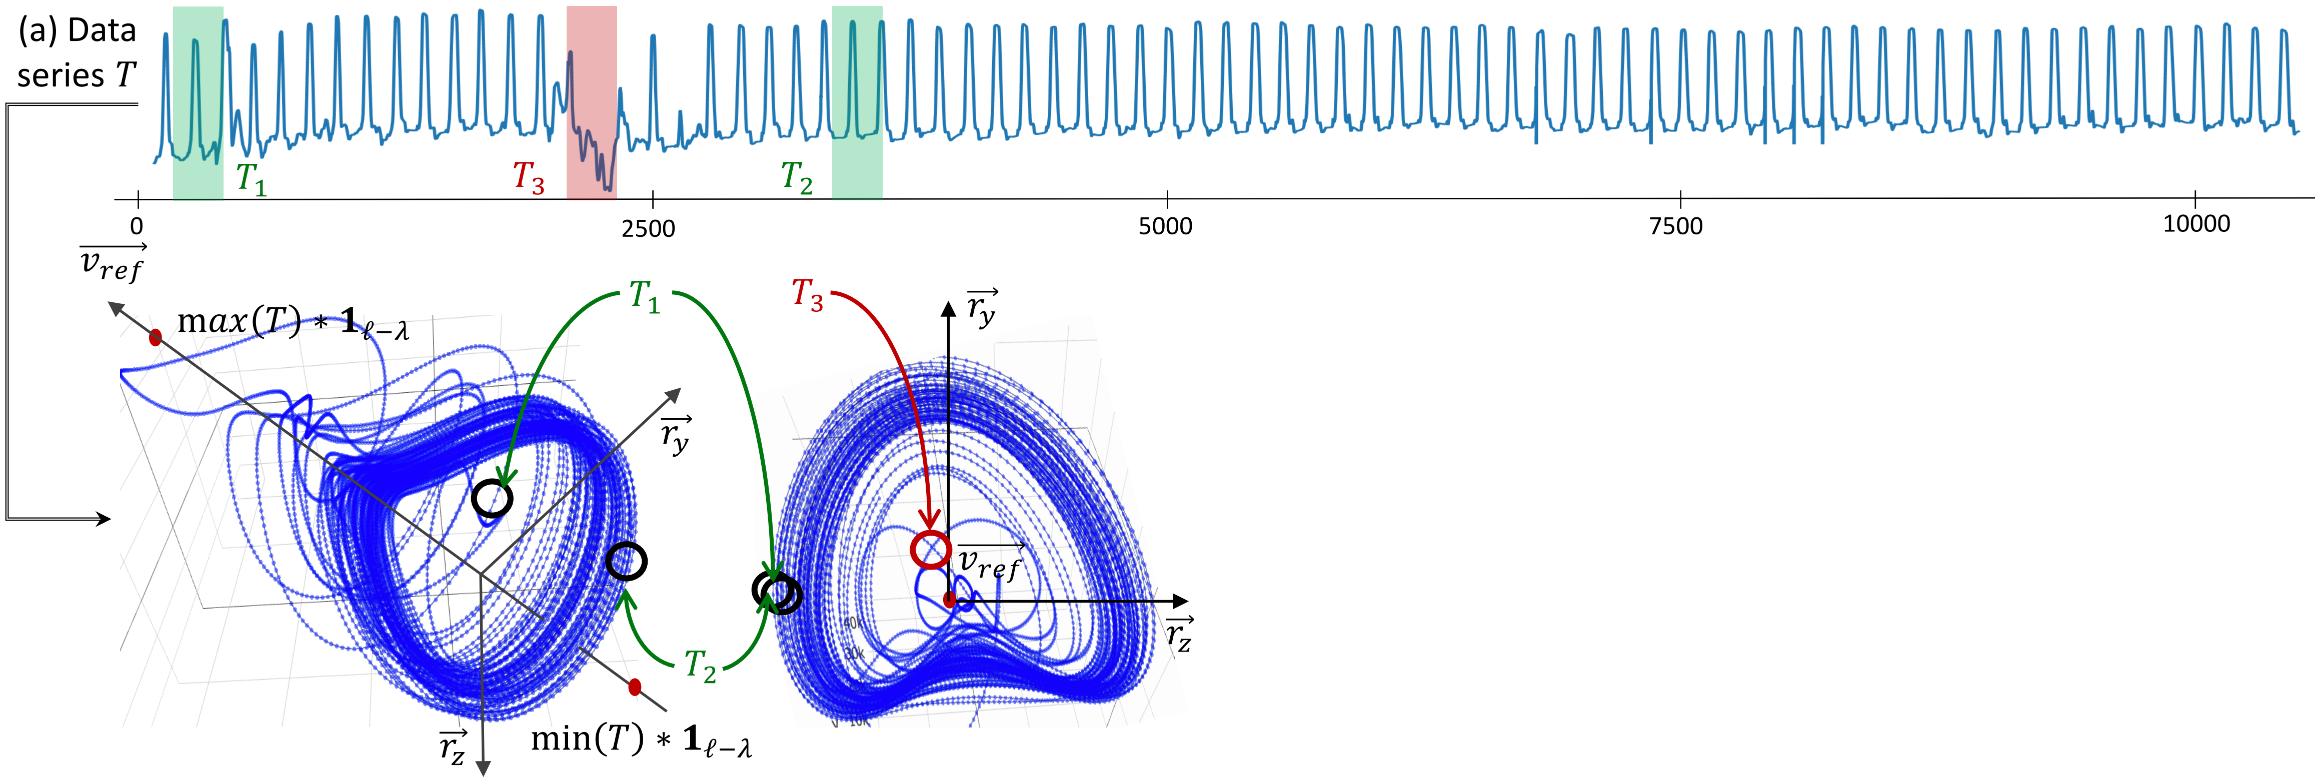

(b) First step: subsequence embedding

Method used:
3 components of the *Principal Component Analysis* applied on all subsequences of T

2.2. Series2Graph: *Computation steps*

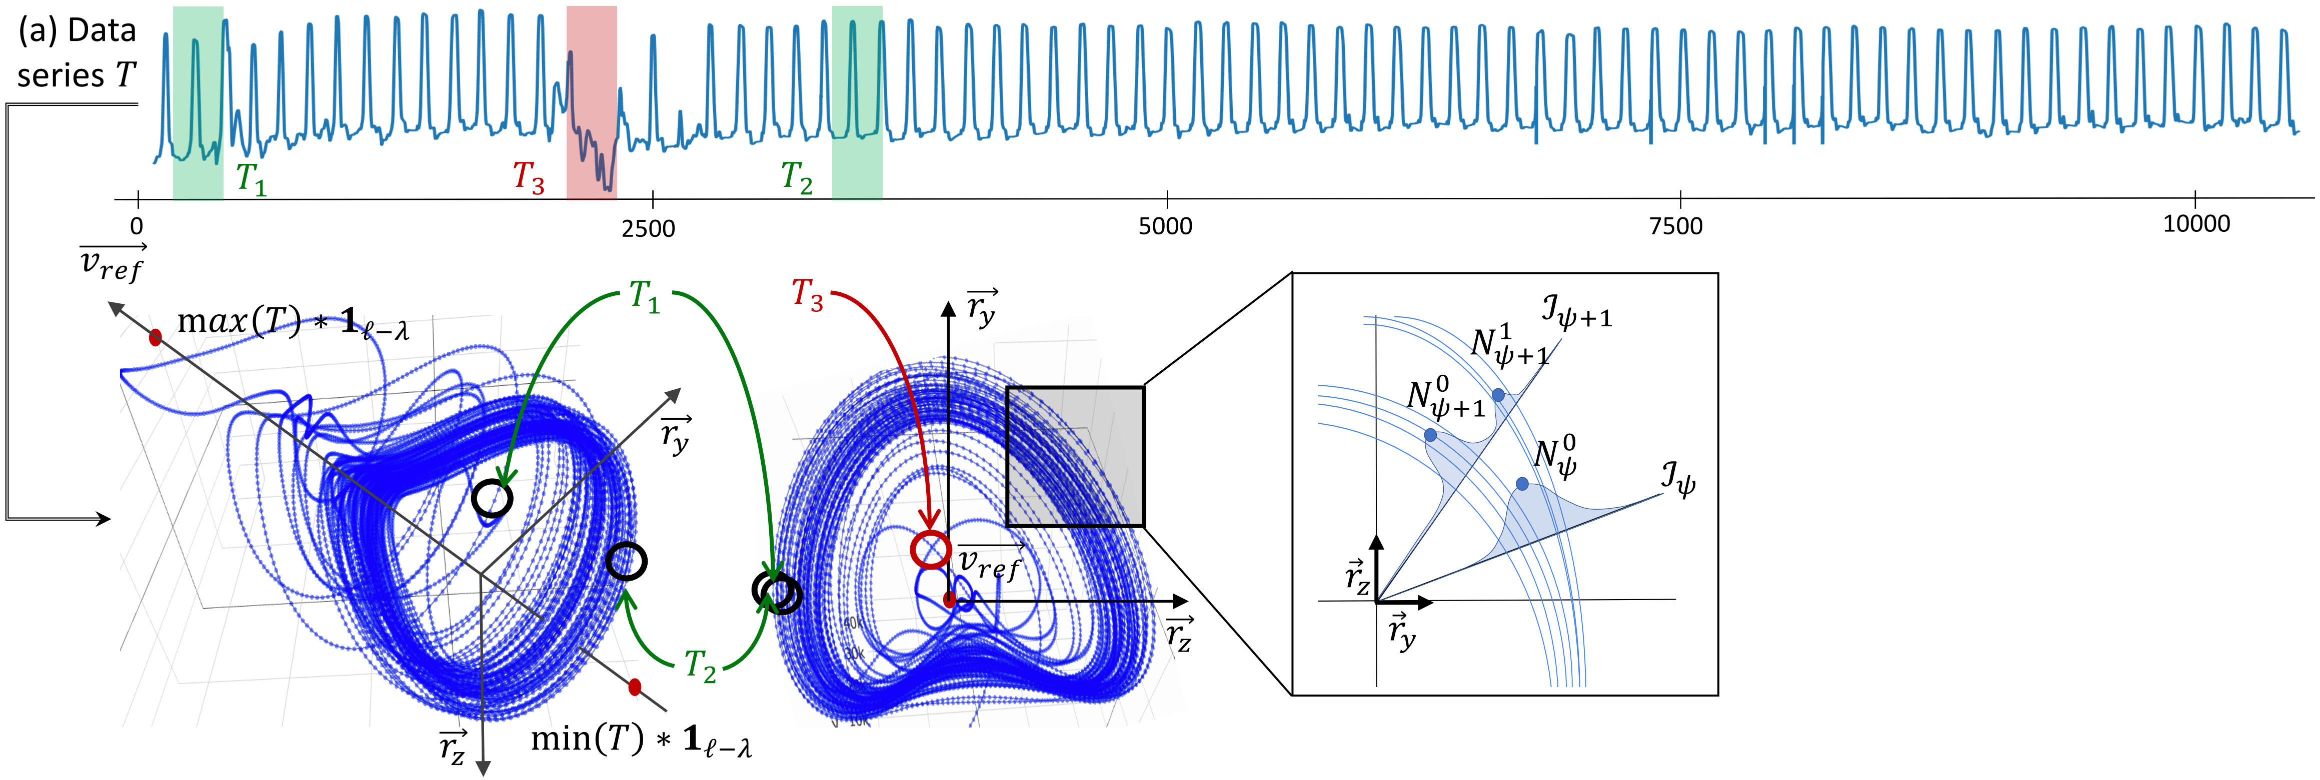

(b) First step: subsequence embedding

(c) Second step: node creation

Method used:

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

2.2. Series2Graph: *Computation steps*

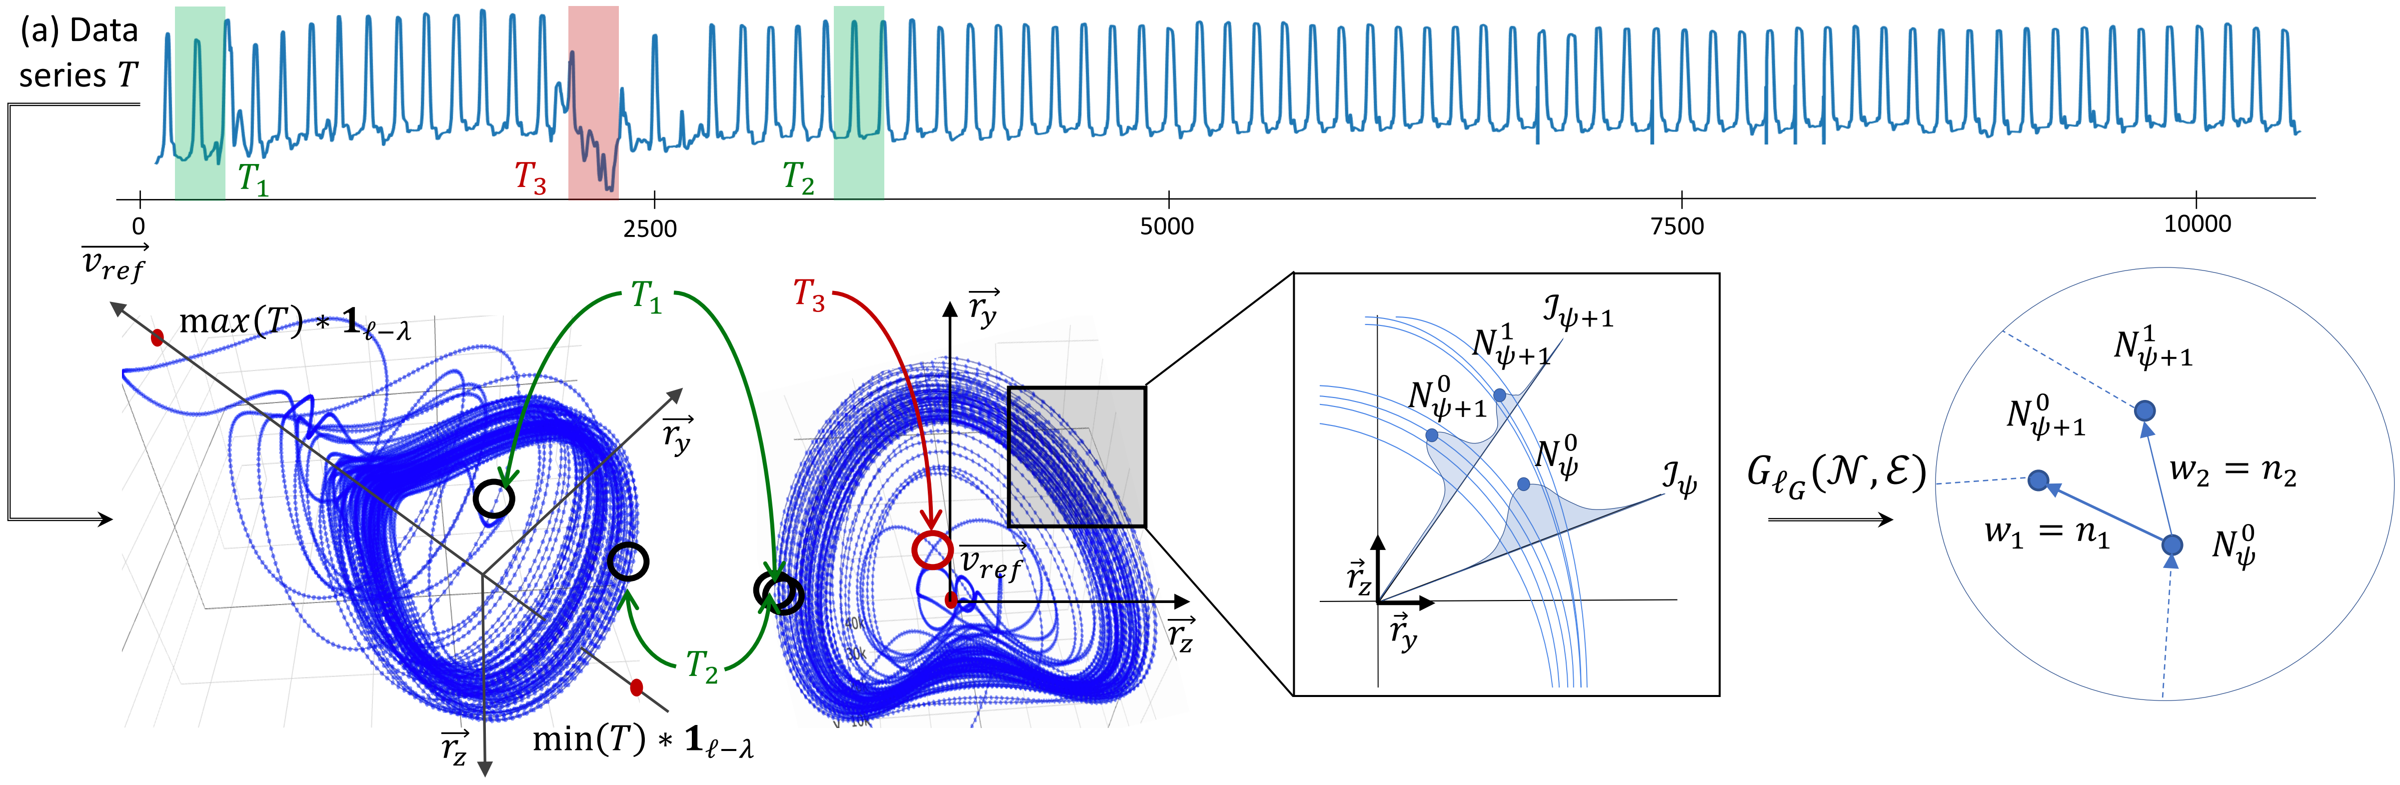

(b) First step: subsequence embedding

(c) Second step: node creation

(d) Third step: edge creation

Method used:

Assign each subsequence to a node and set an edge for each transition between two nodes

3. Supervised identification of anomaly precursors

3.1. Background and definitions

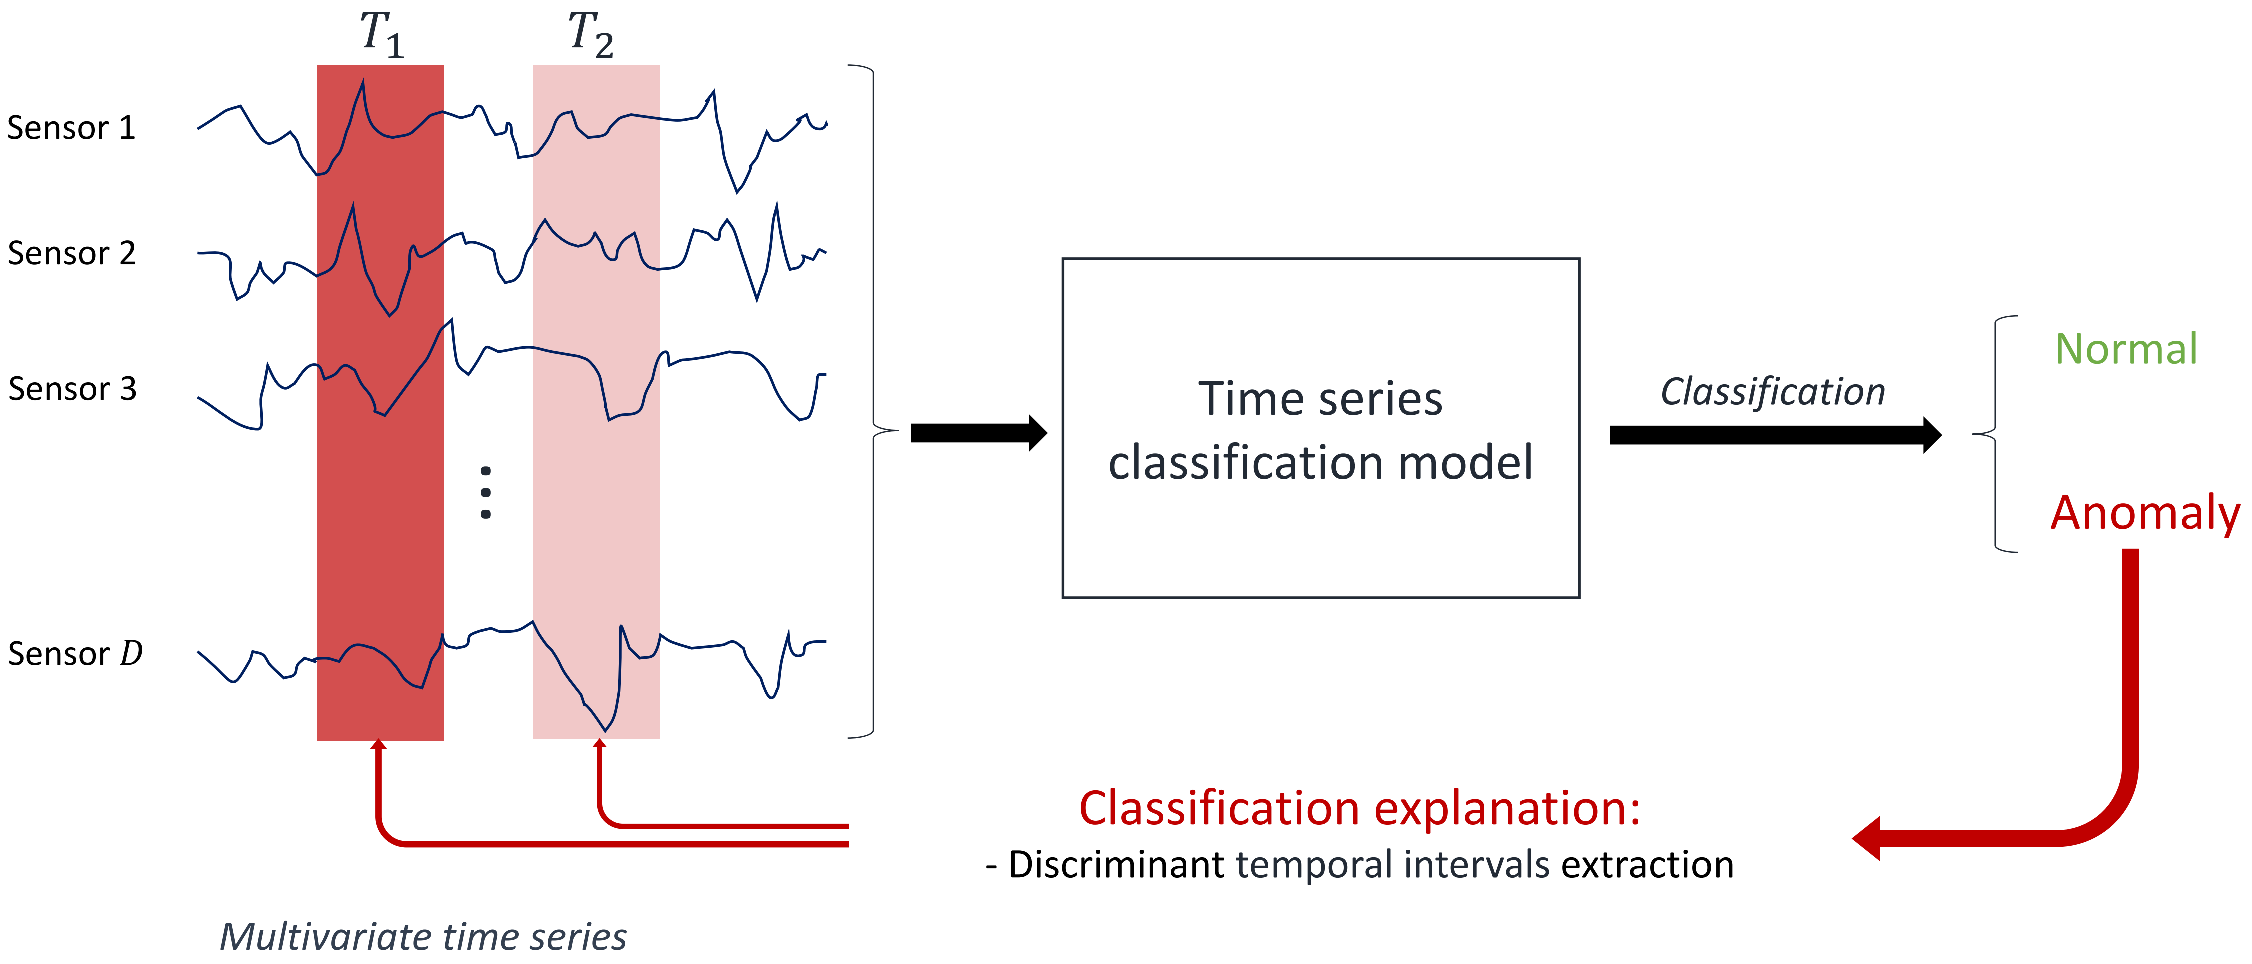

3.1. Background and definitions

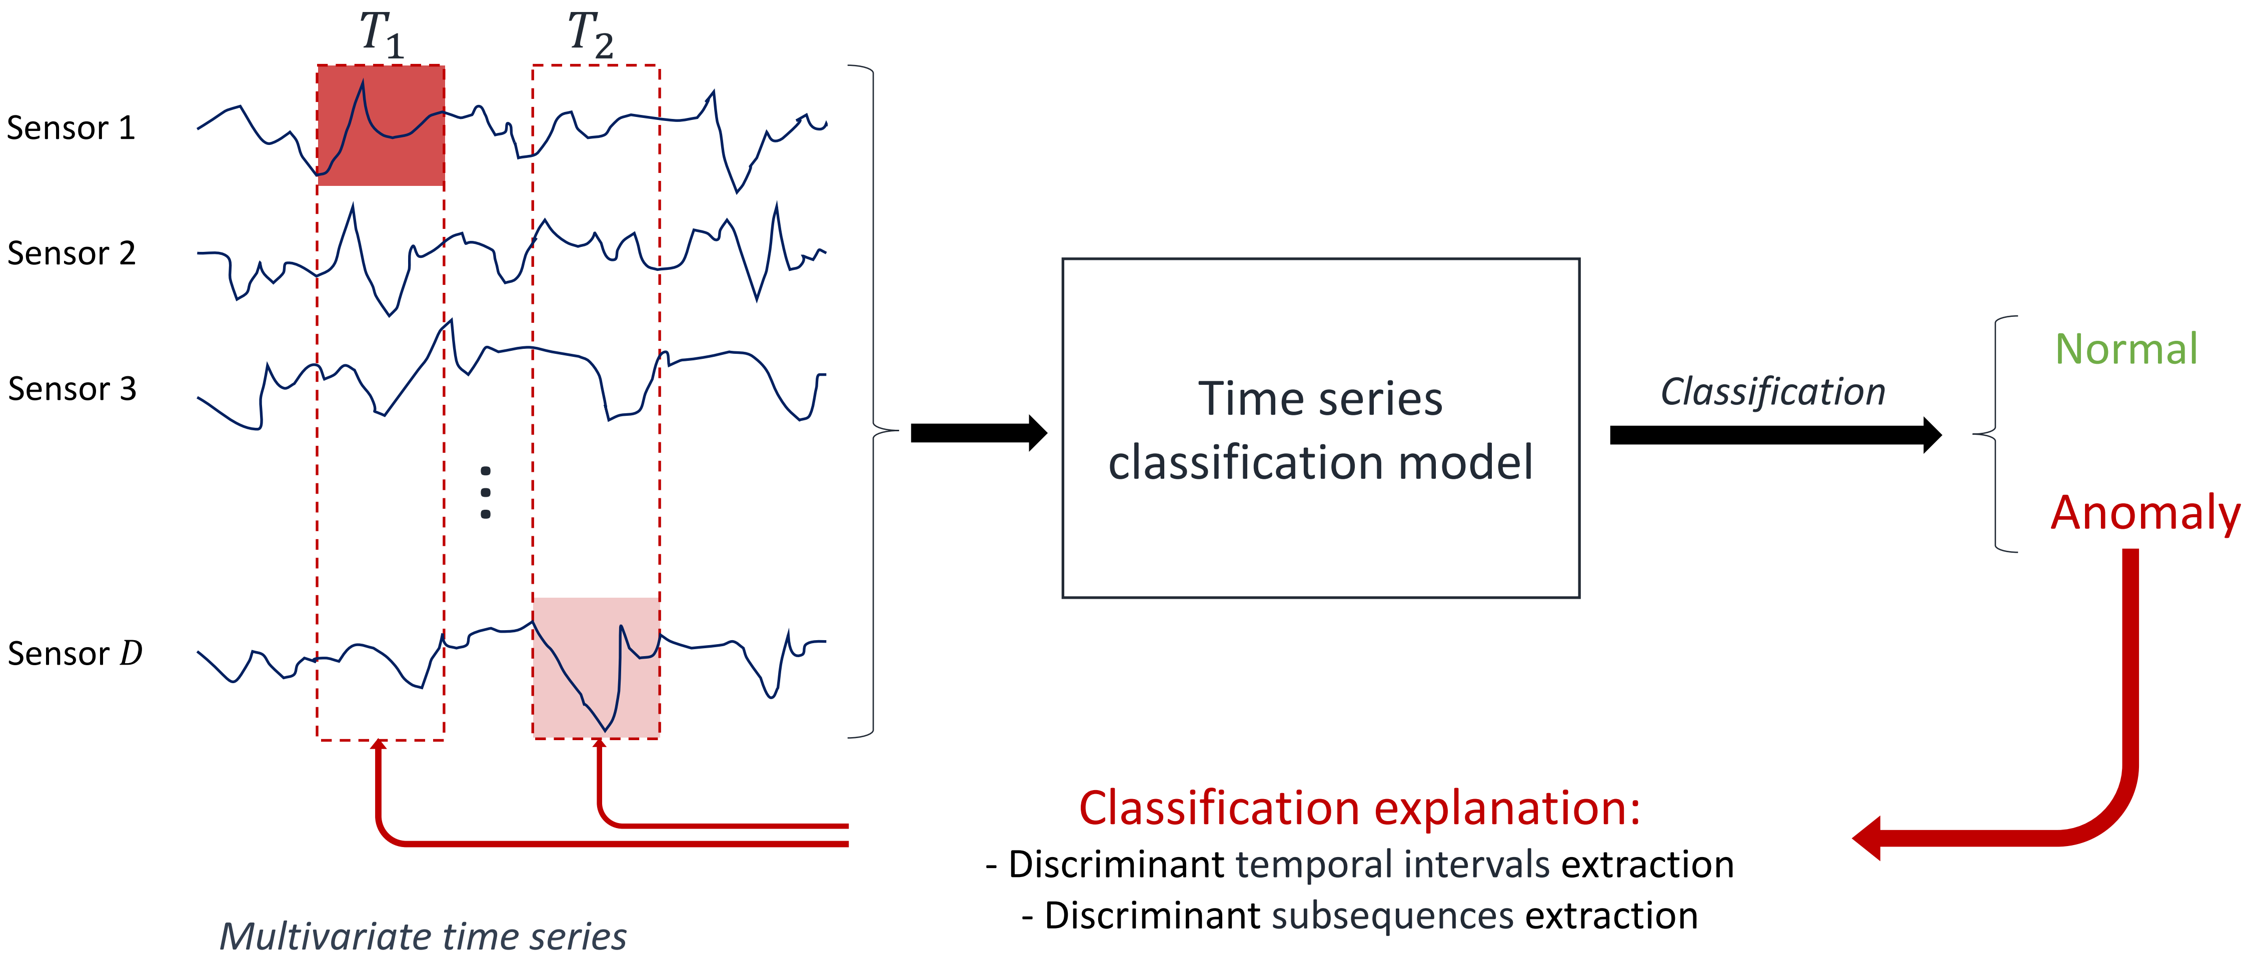

3.2. Related works

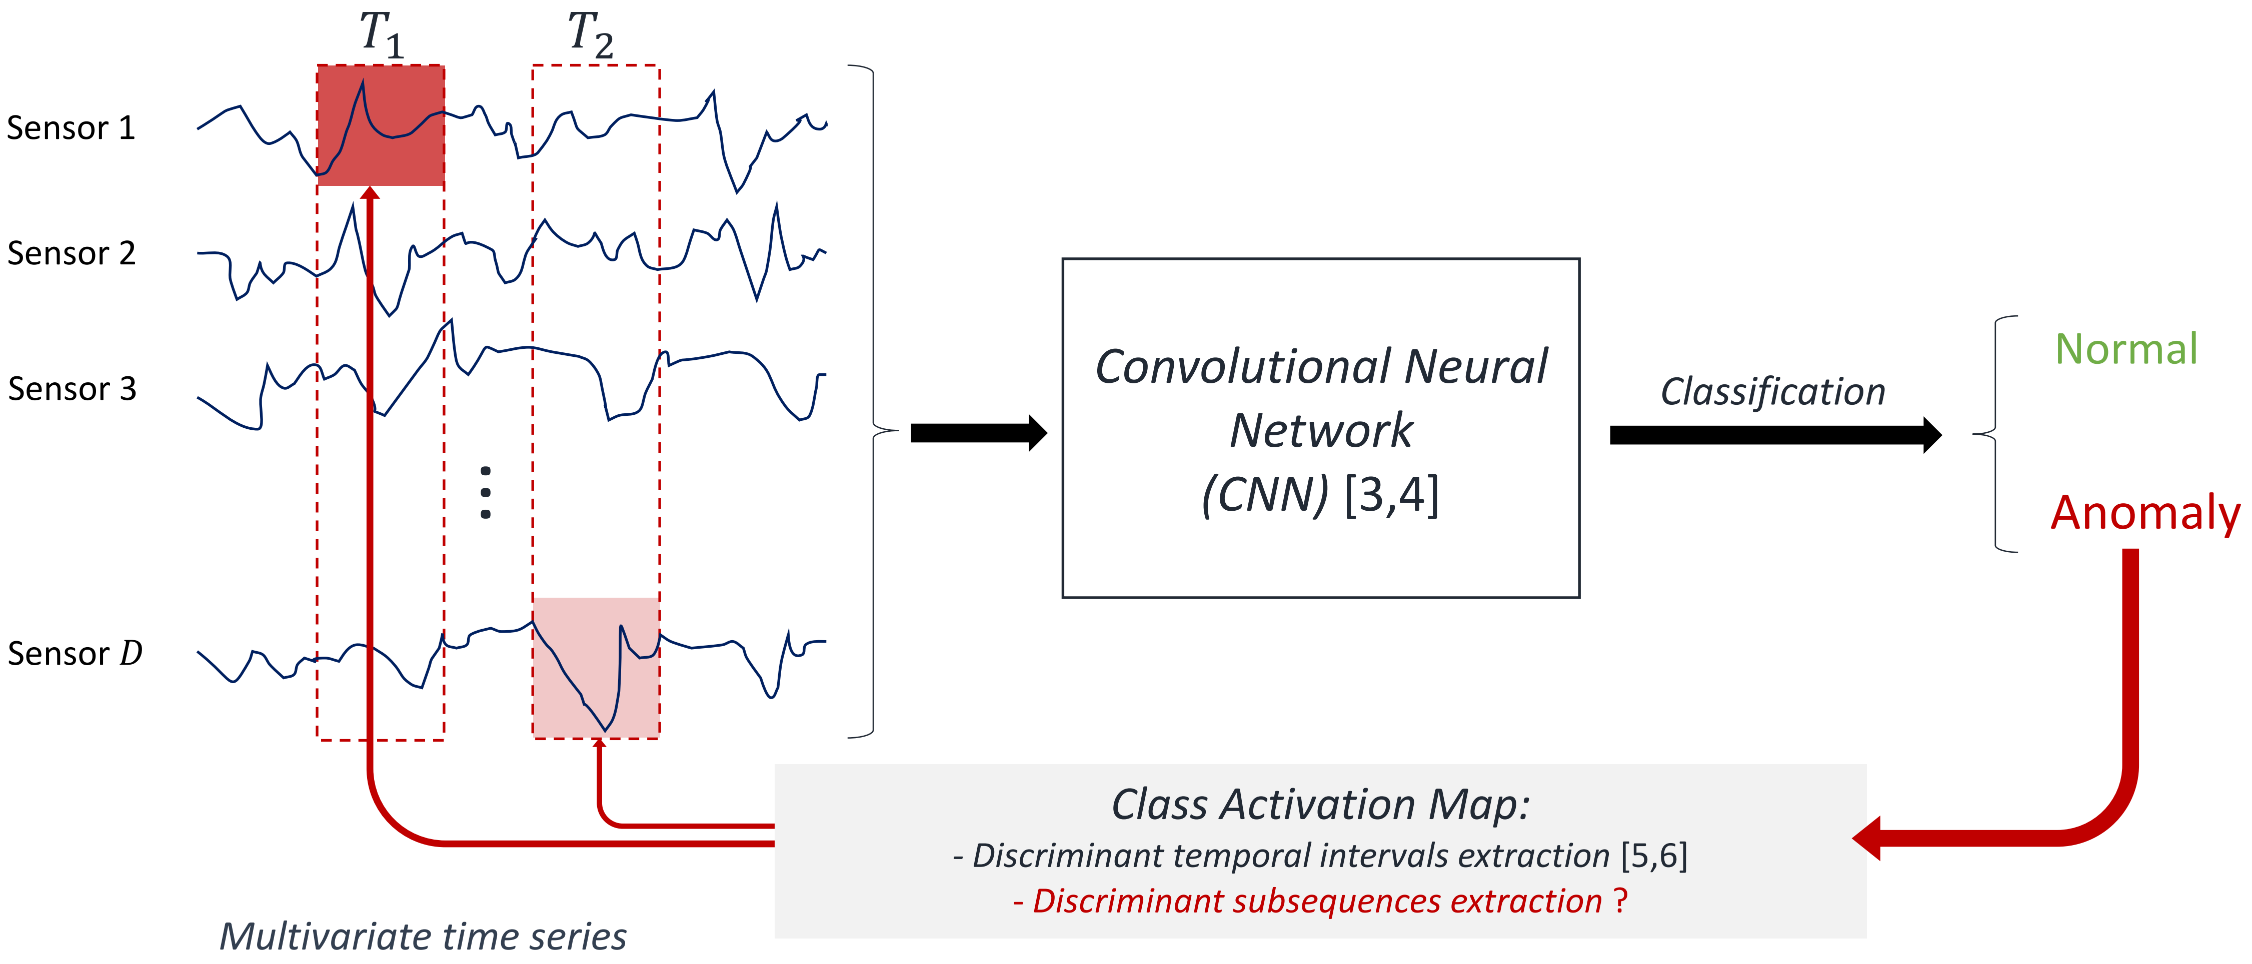

[3] H. I. Fawaz et al., Deep Learning for Time Series Classification: A Review, Data Min. Knowl. Discov., 33, 4, 2019

[4] Z. Wang et al., Time series classification from scratch with deep neural networks: A strong baseline, IJCNN, 2017

[5] B. Zhou et al., Learning Deep Features for Discriminative Localization, CVPR, 2016

[6] H. I. Fawaz et al., Evaluating Surgical Skills from Kinematic Data Using Convolutional Neural Networks, MICCAI, 2018

3.2. Related works: *Class Activation Map example*

Originally proposed for the identification and localization of objects in a picture that would explain the classification outcome by a convolutional neural network [5].
It can also be applied on data series [6,7].

Electrocardiogram classified as “arythmia” [7]

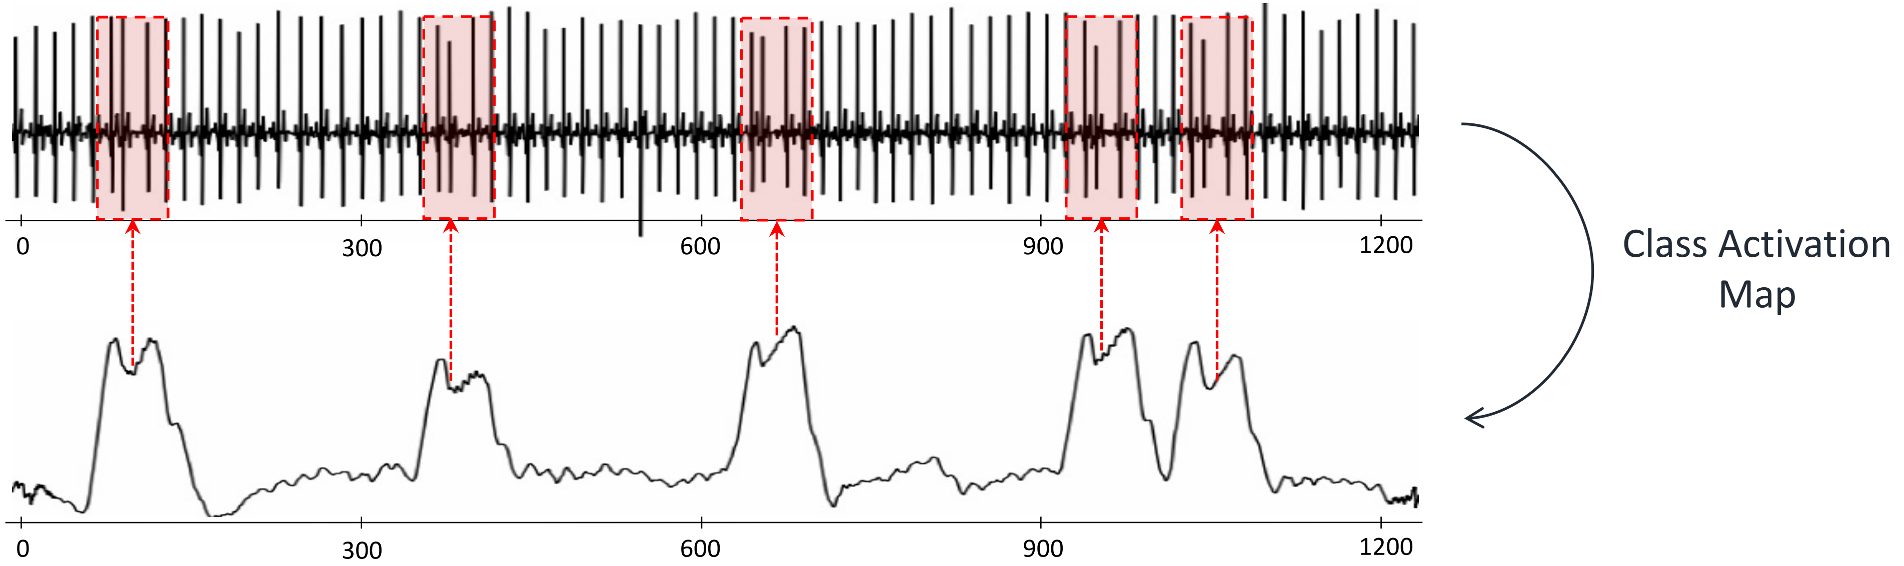

Subsequences of the time series used by the model to classify as “arythmia”

[5] B. Zhou et al., Learning Deep Features for Discriminative Localization, CVPR, 2016
[6] H. I. Fawaz et al. Evaluating Surgical Skills from Kinematic Data Using Convolutional Neural Networks. In MICCAI (2018)
[7] Sebastian D. Goodfellow et al. Towards Understanding ECG Rhythm Classification Using Convolutional Neural Networks and Attention Mappings. In PMLR (2018)

3.2. Related works: *limitation of CAM*

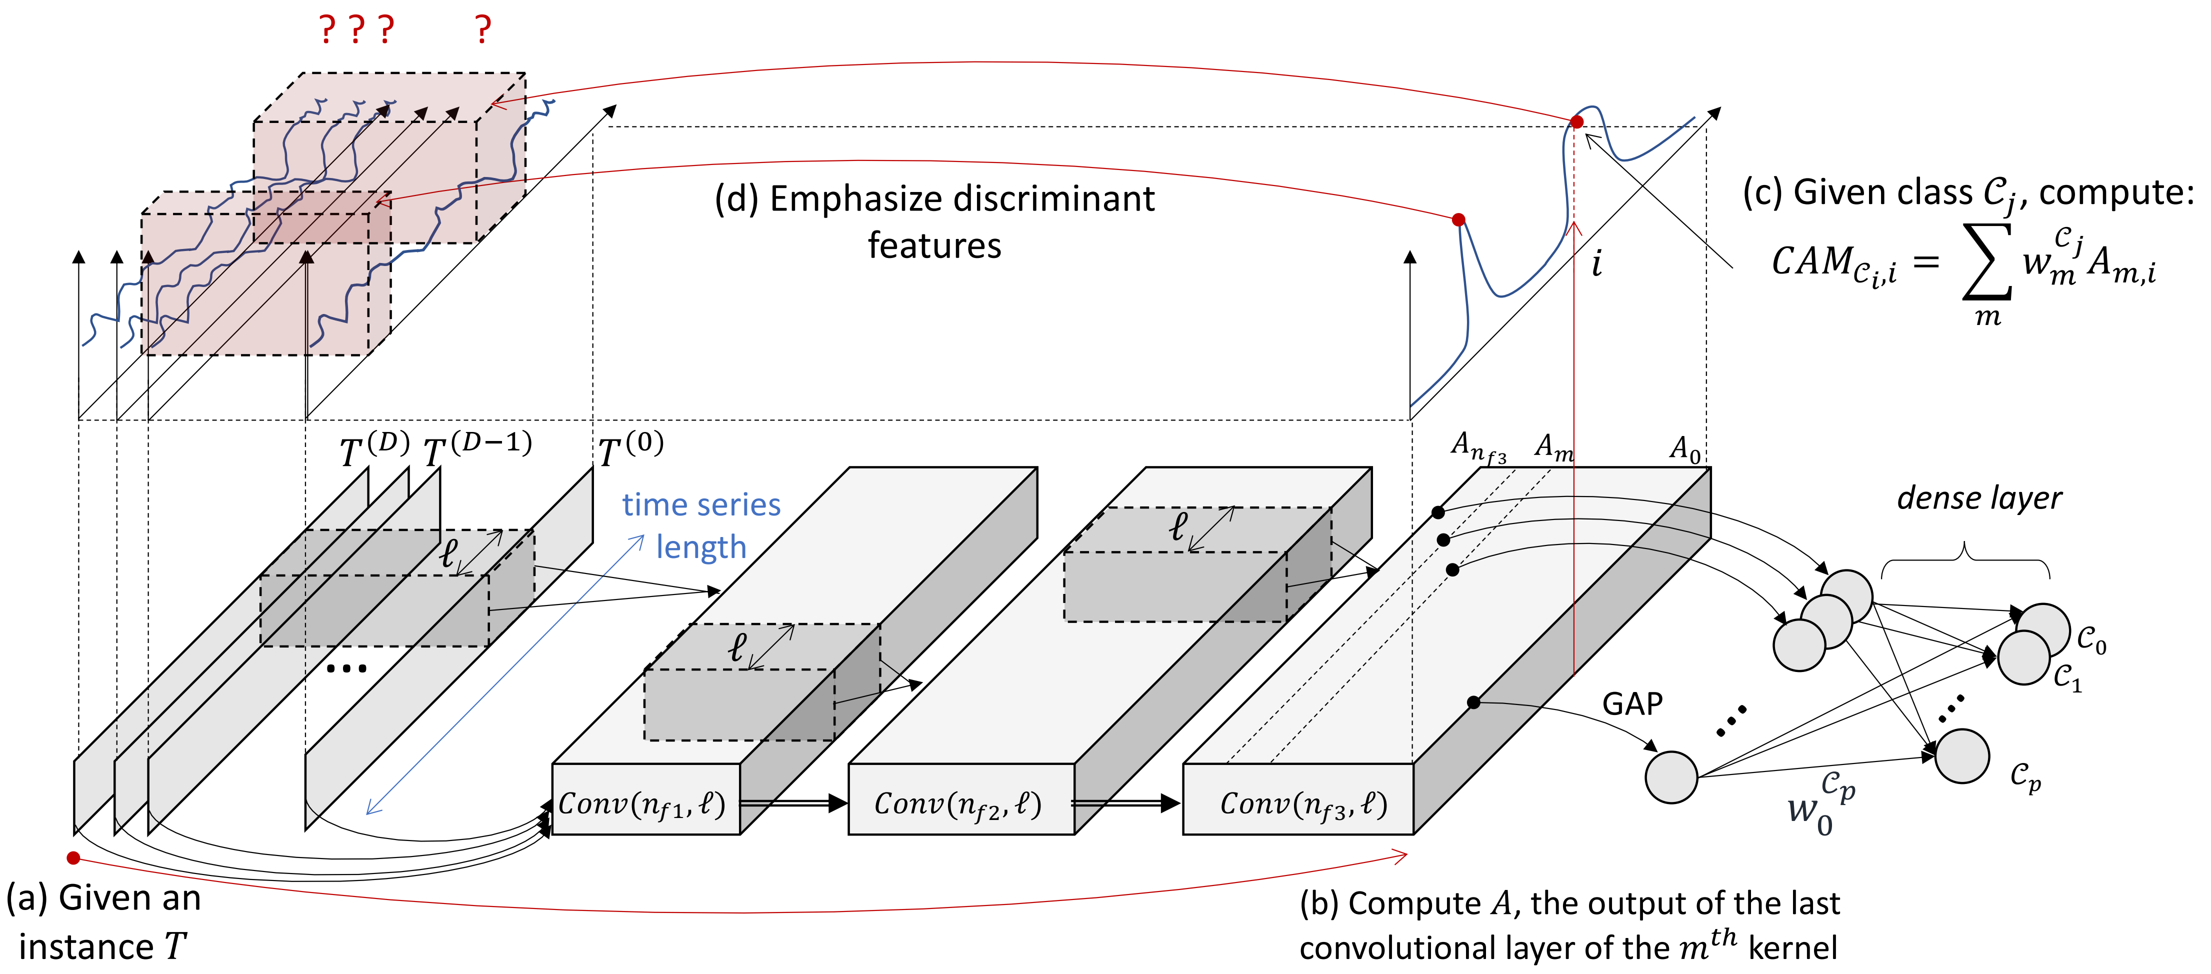

3.3. Proposed approach: *dimension-wise Class Activation Map (dCAM)*

- A generic proposed method [8].
- New input structure to the model (dCNN), without any constraint on the architecture.
- Permutations of the input to identify discriminant subsequences within each dimension.

For a given multivariate time series $T = \{T^{(0)}, \dots, T^{(D)}\}$:

$$C(T) = \begin{pmatrix} T^{(D)} & T^{(0)} & \cdots & T^{(D-2)} & T^{(D-1)} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ T^{(1)} & T^{(2)} & \ddots & T^{(D)} & T^{(0)} \\ T^{(0)} & T^{(1)} & \cdots & T^{(D-1)} & T^{(D)} \end{pmatrix} =$$

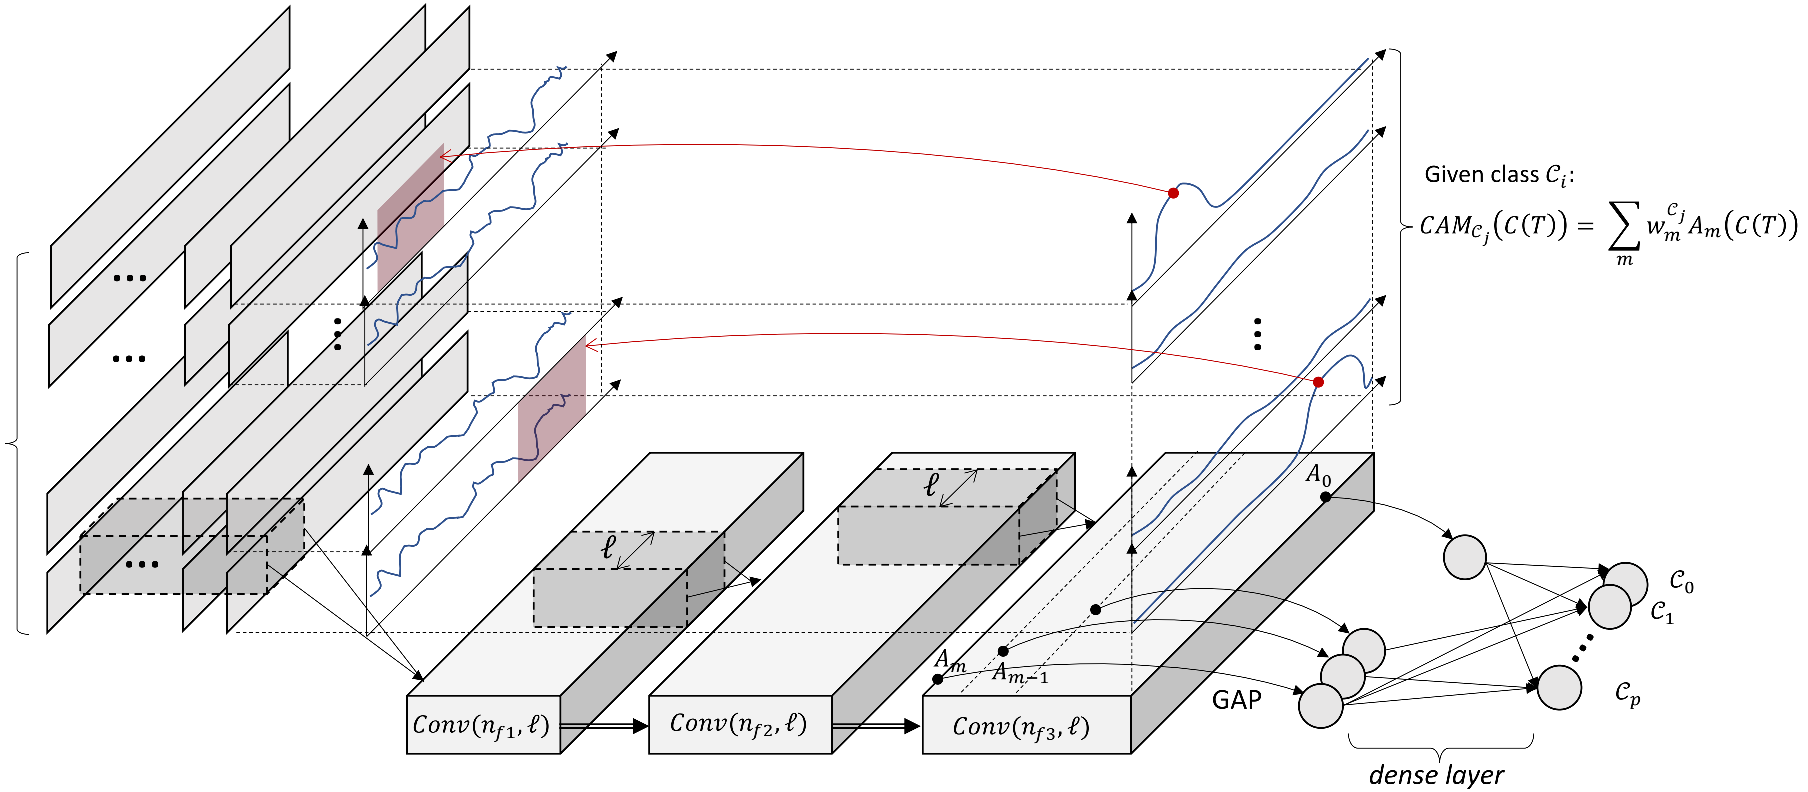

[8] P. Boniol et al., dCAM: Dimension-Wise Class Activation Map for Explaining Multivariate Data series Classification, SIGMOD, 2022

3.3. Proposed approach: *dCAM*

$$C(T) = \begin{pmatrix} T^{(D)} & T^{(0)} & \cdots & T^{(D-1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(1)} & T^{(2)} & \ddots & T^{(0)} \\ T^{(0)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix} \xrightarrow{CAM_{\mathcal{C}_j}(S_T^0)}$$

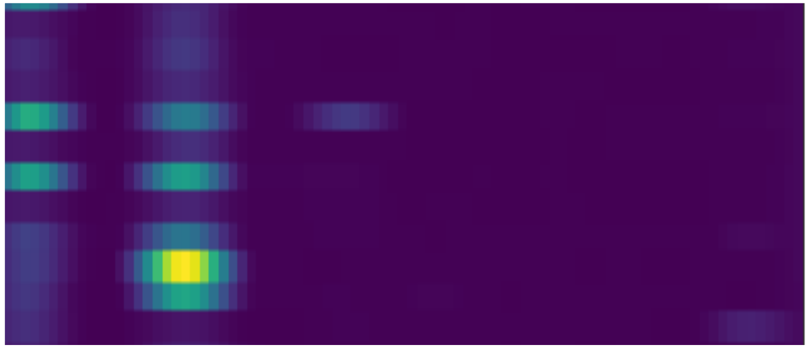

3.3. Proposed approach: *dCAM*

$$C(T) = \begin{pmatrix} T^{(D)} & T^{(0)} & \cdots & T^{(D-1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(1)} & T^{(2)} & \ddots & T^{(0)} \\ T^{(0)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix}$$

$T^{(0)}, T^{(1)}, T^{(2)}, T^{(3)}, T^{(4)}, T^{(5)}$

3.3. Proposed approach: *dCAM*

$$C(T) = \begin{pmatrix} T^{(D)} & T^{(0)} & \cdots & T^{(D-1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(1)} & T^{(2)} & \ddots & T^{(0)} \\ T^{(0)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix}$$

$$T^{(0)}, T^{(1)}, T^{(2)}, T^{(3)}, T^{(4)}, T^{(5)}$$

$$S_T^1 = \begin{pmatrix} T^{(3)} & T^{(D)} & \cdots & T^{(D-2)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(D)} & T^{(4)} & \ddots & T^{(1)} \\ T^{(2)} & T^{(1)} & \cdots & T^{(0)} \end{pmatrix} \xrightarrow{CAM_{\mathcal{C}_j}(S_T^1)}$$

3.3. Proposed approach: *dCAM*

$$C(T) = \begin{pmatrix} T^{(D)} & T^{(0)} & \cdots & T^{(D-1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(1)} & T^{(2)} & \ddots & T^{(0)} \\ T^{(0)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix}$$

$T^{(0)}, T^{(1)}, T^{(2)}, T^{(3)}, T^{(4)}, T^{(5)}$

$$S_T^1 = \begin{pmatrix} T^{(3)} & T^{(D)} & \cdots & T^{(D-2)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(D)} & T^{(4)} & \ddots & T^{(1)} \\ T^{(2)} & T^{(1)} & \cdots & T^{(0)} \end{pmatrix}$$

$T^{(2)}, T^{(1)}, T^{(0)}, T^{(3)}, T^{(4)}, T^{(5)}$

3.3. Proposed approach: *dCAM*

$$C(T) = \begin{pmatrix} T^{(D)} & T^{(0)} & \cdots & T^{(D-1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(1)} & T^{(2)} & \ddots & T^{(0)} \\ T^{(0)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix}$$

$T^{(0)}, T^{(1)}, T^{(2)}, T^{(3)}, T^{(4)}, T^{(5)}$

$$S_T^1 = \begin{pmatrix} T^{(3)} & T^{(D)} & \cdots & T^{(D-2)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(D)} & T^{(4)} & \ddots & T^{(1)} \\ T^{(2)} & T^{(1)} & \cdots & T^{(0)} \end{pmatrix}$$

$T^{(2)}, T^{(1)}, T^{(0)}, T^{(3)}, T^{(4)}, T^{(5)}$

$$S_T^2 = \begin{pmatrix} T^{(2)} & T^{(D)} & \cdots & T^{(1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(D-1)} & T^{(3)} & \ddots & T^{(2)} \\ T^{(3)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix} \xrightarrow{CAM_{\mathcal{C}_j}(S_T^2)}$$

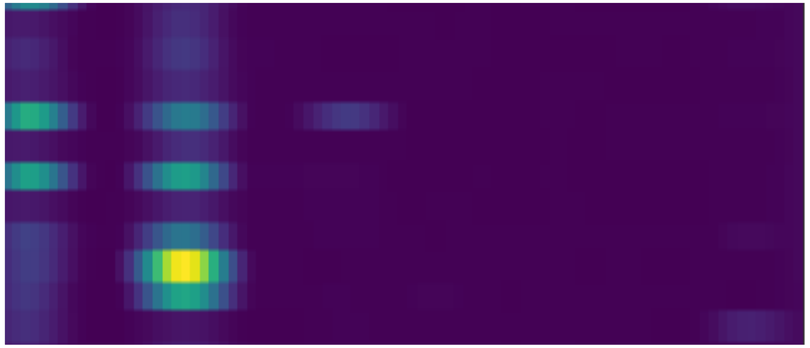

3.3. Proposed approach: *dCAM*

$$C(T) = \begin{pmatrix} T^{(D)} & T^{(0)} & \cdots & T^{(D-1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(1)} & T^{(2)} & \ddots & T^{(0)} \\ T^{(0)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix}$$

$T^{(0)}, T^{(1)}, T^{(2)}, T^{(3)}, T^{(4)}, T^{(5)}$

$$S_T^1 = \begin{pmatrix} T^{(3)} & T^{(D)} & \cdots & T^{(D-2)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(D)} & T^{(4)} & \ddots & T^{(1)} \\ T^{(2)} & T^{(1)} & \cdots & T^{(0)} \end{pmatrix}$$

$T^{(2)}, T^{(1)}, T^{(0)}, T^{(3)}, T^{(4)}, T^{(5)}$

$$S_T^2 = \begin{pmatrix} T^{(2)} & T^{(D)} & \cdots & T^{(1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(D-1)} & T^{(3)} & \ddots & T^{(2)} \\ T^{(3)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix}$$

$T^{(4)}, T^{(5)}, T^{(2)}, T^{(1)}, T^{(3)}, T^{(0)}$

3.3. Proposed approach: *dCAM*

$$C(T) = \begin{pmatrix} T^{(D)} & T^{(0)} & \cdots & T^{(D-1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(1)} & T^{(2)} & \ddots & T^{(0)} \\ T^{(0)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix}$$

$$S_T^1 = \begin{pmatrix} T^{(3)} & T^{(D)} & \cdots & T^{(D-2)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(D)} & T^{(4)} & \ddots & T^{(1)} \\ T^{(2)} & T^{(1)} & \cdots & T^{(0)} \end{pmatrix}$$

$$S_T^2 = \begin{pmatrix} T^{(2)} & T^{(D)} & \cdots & T^{(1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(D-1)} & T^{(3)} & \ddots & T^{(2)} \\ T^{(3)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix}$$

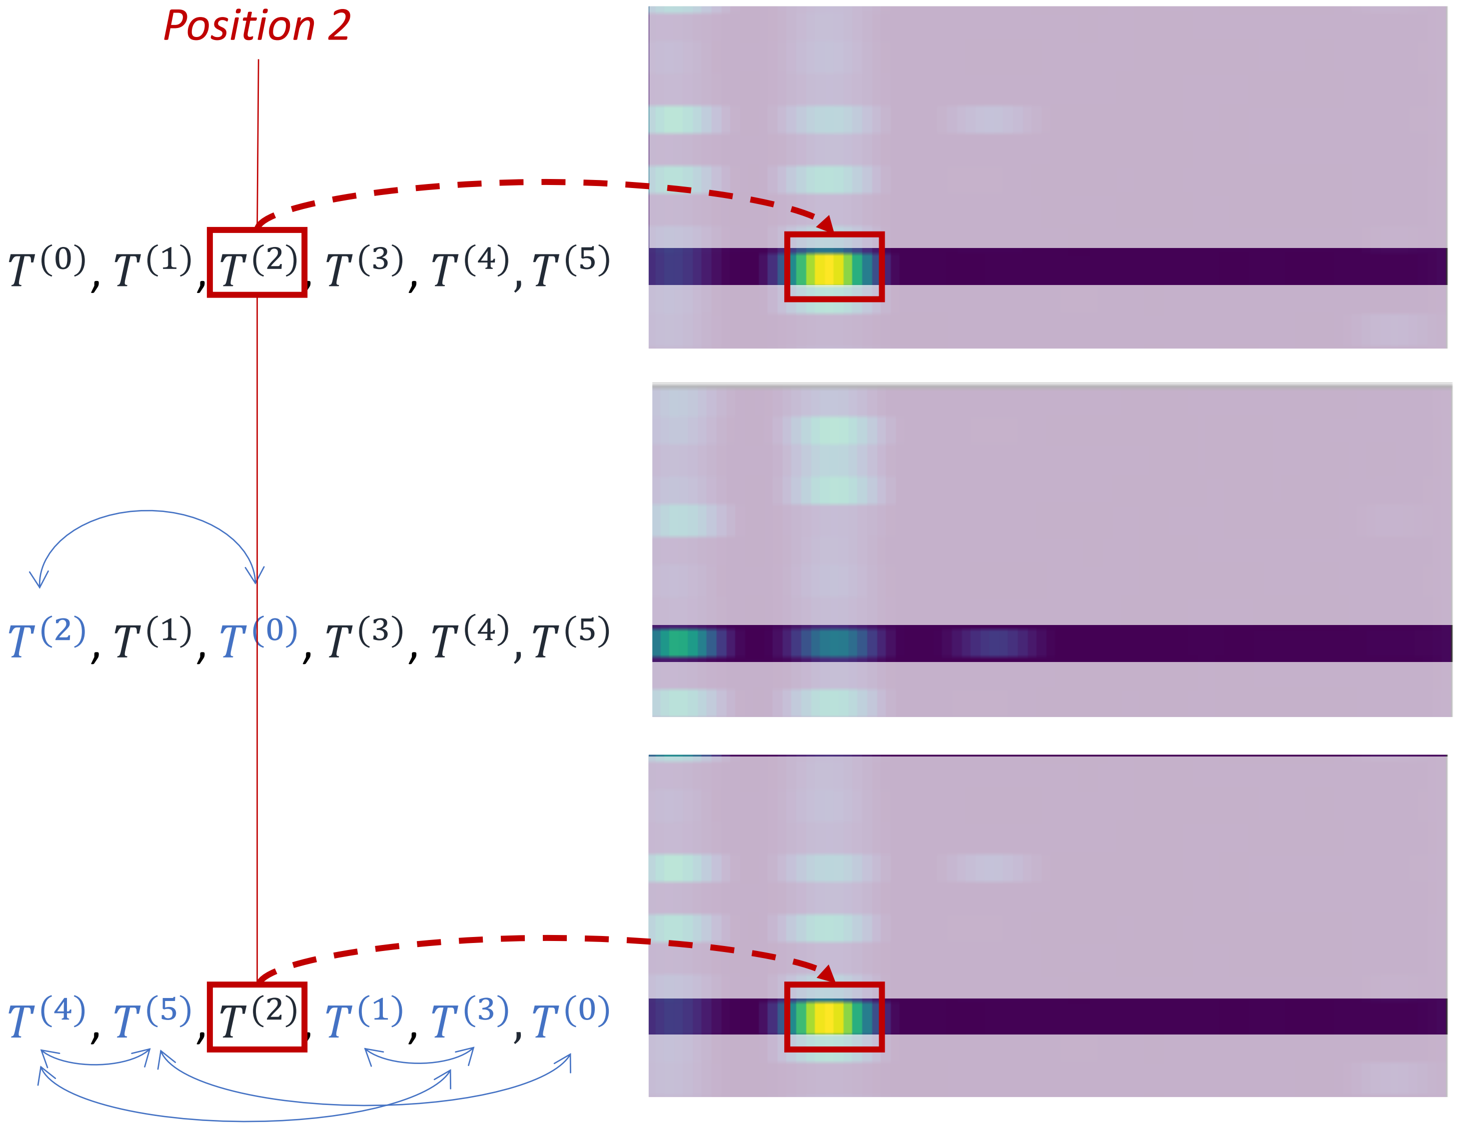

3.3. Proposed approach: *dCAM*

$$C(T) = \begin{pmatrix} T^{(D)} & T^{(0)} & \cdots & T^{(D-1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(1)} & T^{(2)} & \ddots & T^{(0)} \\ T^{(0)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix}$$

$$S_T^1 = \begin{pmatrix} T^{(3)} & T^{(D)} & \cdots & T^{(D-2)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(D)} & T^{(4)} & \ddots & T^{(1)} \\ T^{(2)} & T^{(1)} & \cdots & T^{(0)} \end{pmatrix}$$

$$S_T^2 = \begin{pmatrix} T^{(2)} & T^{(D)} & \cdots & T^{(1)} \\ \vdots & \vdots & \ddots & \vdots \\ T^{(D-1)} & T^{(3)} & \ddots & T^{(2)} \\ T^{(3)} & T^{(1)} & \cdots & T^{(D)} \end{pmatrix}$$

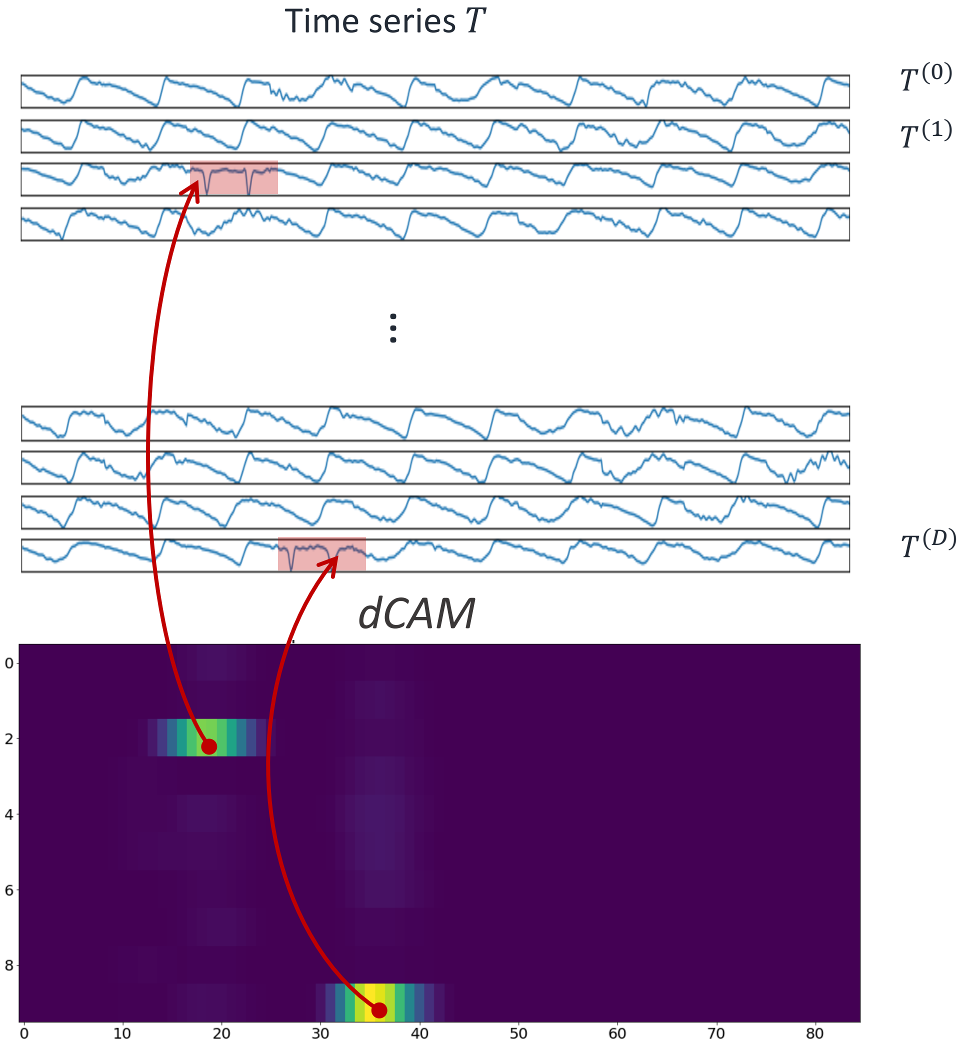

4. Industrial study: pump vibrations case

4.1. Use case: *identification of precursors of vibrations on feed-water pump*

Goal: pattern identification that could explain why vibrations are happening on feed-water pump in French nuclear power plants [9].

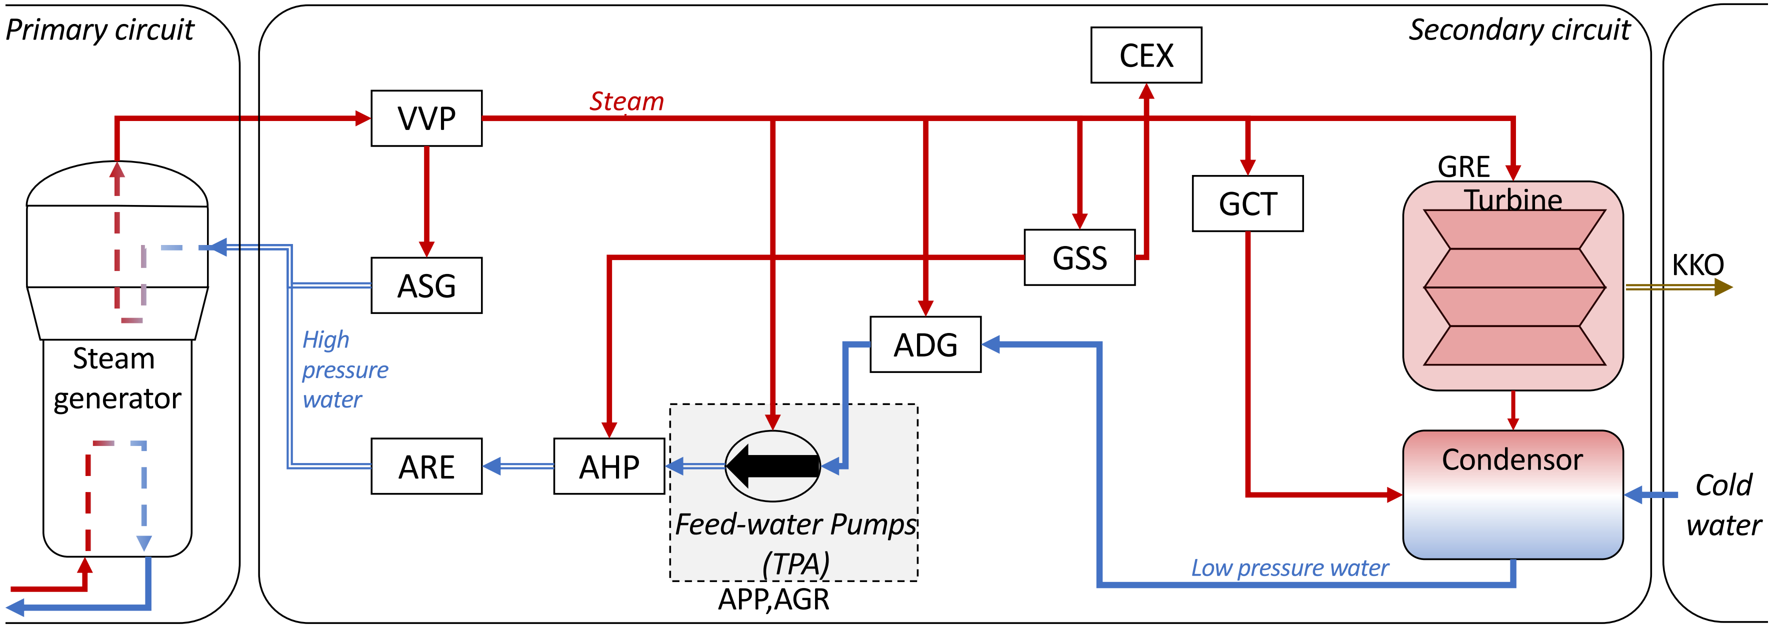

[9] P. Boniol et al., dCNN/dCAM: Anomaly Precursors Discovery in Multivariate Time Series with Deep Convolutional Neural Network, in preparation to MSSP

4.1. Use case: *identification of precursors of vibrations on feed-water pump*

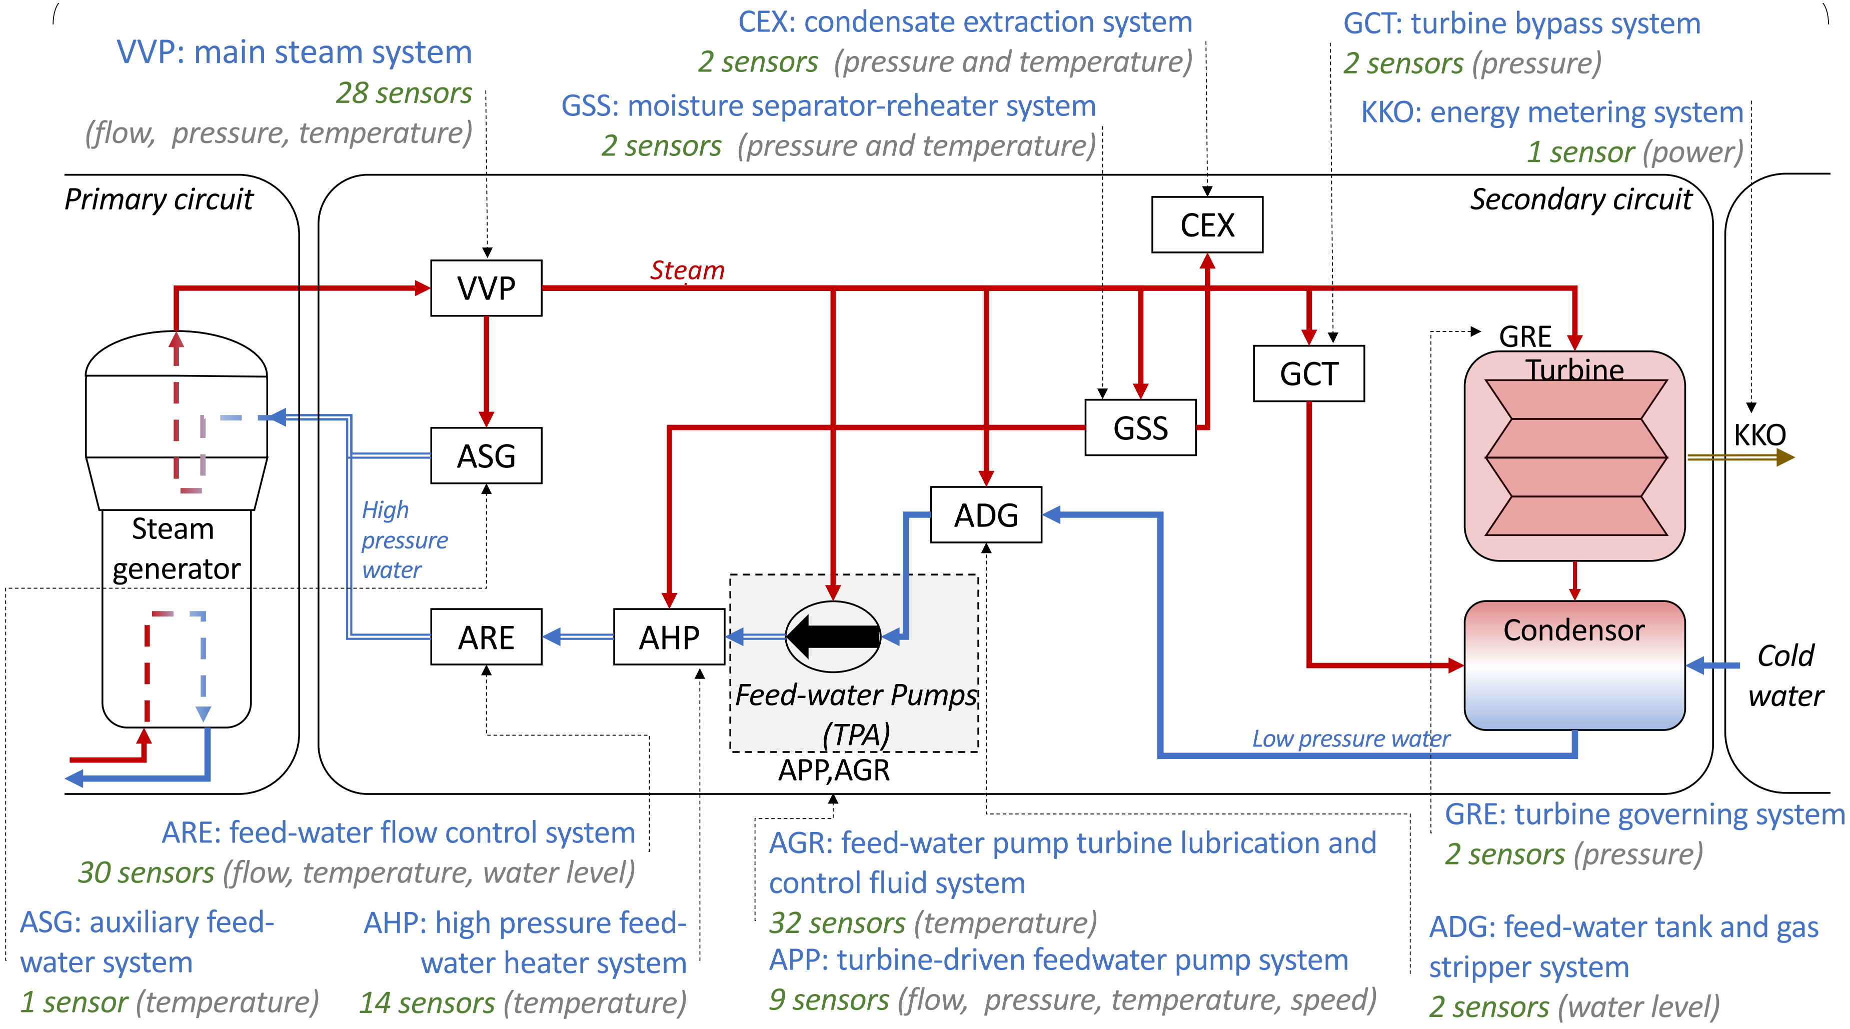

4.1. Use case: *identification of precursors of vibrations on feed-water pump*

Goal: pattern identification that could explain why vibrations are happening on feed-water pump in French nuclear power plants [9].

Class 1:

451 multivariate time series (127 dimensions, 800 points) corresponding to intervals of 80 minutes without any vibration

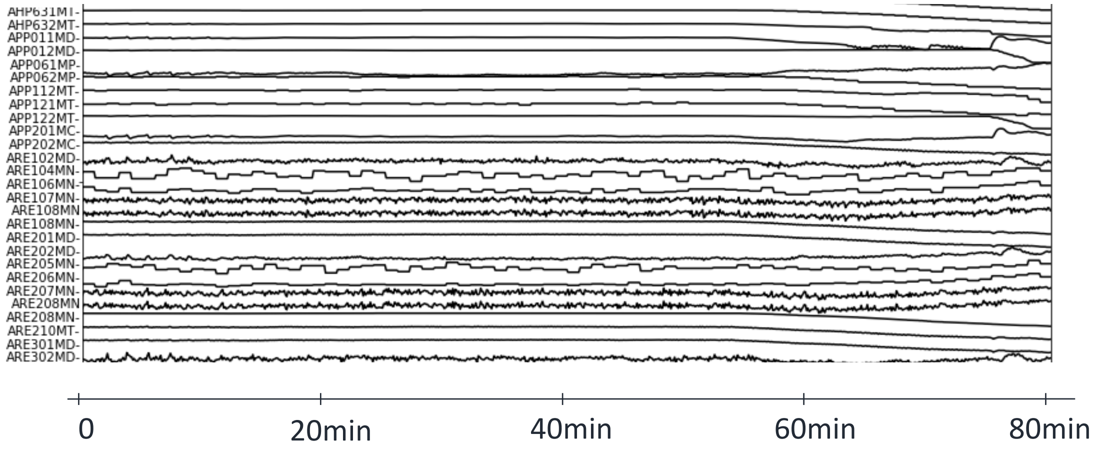

[9] P. Boniol et al., dCNN/dCAM: Anomaly Precursors Discovery in Multivariate Time Series with Deep Convolutional Neural Network, in preparation to MSSP

4.1. Use case: *identification of precursors of vibrations on feed-water pump*

Goal: pattern identification that could explain why vibrations are happening on feed-water pump in French nuclear power plants [9].

Class 1:

451 multivariate time series (127 dimensions, 800 points) corresponding to intervals of 80 minutes without any vibration

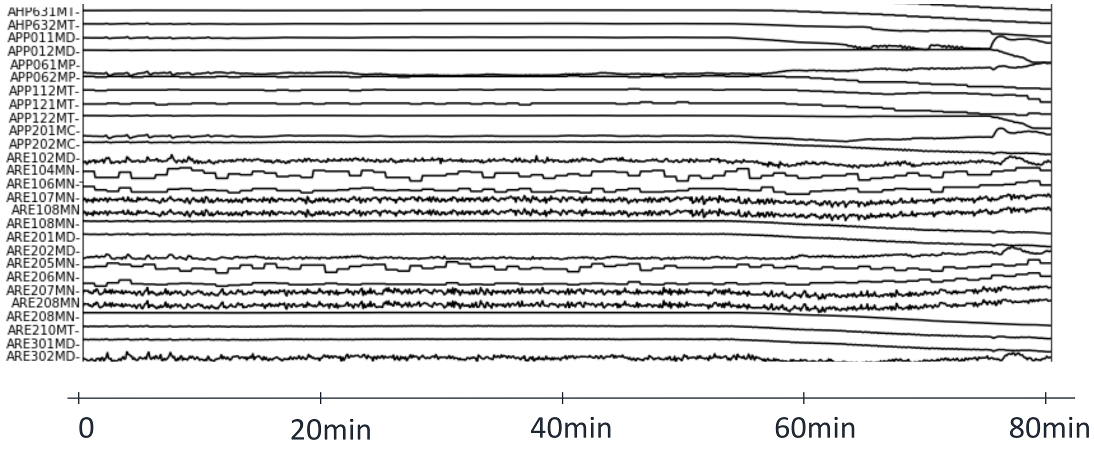

Class 2:

445 multivariate time series (127 dimensions, 800 points) corresponding to intervals of 75 minutes before a vibration and 5 minutes after

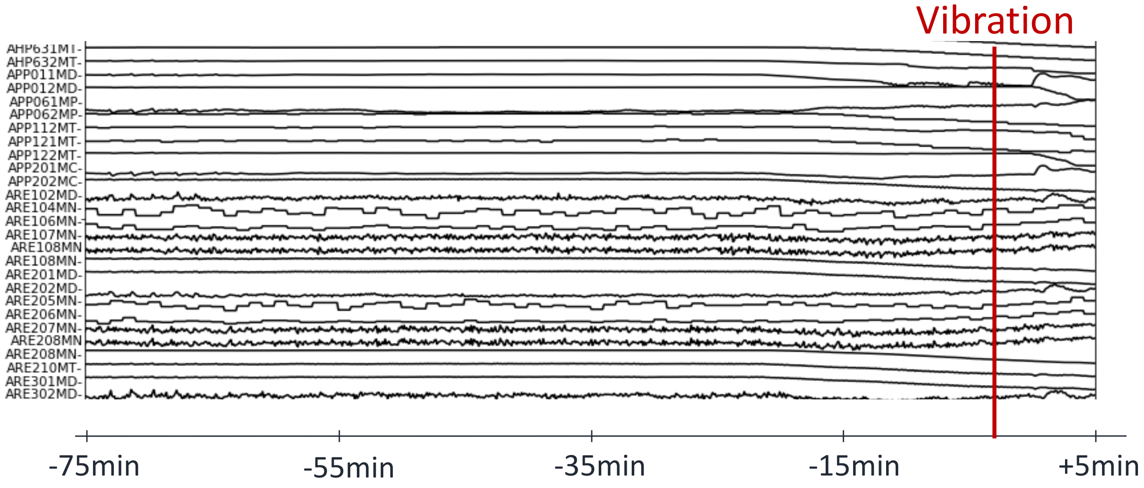

[9] P. Boniol et al., dCNN/dCAM: Anomaly Precursors Discovery in Multivariate Time Series with Deep Convolutional Neural Network, in preparation to MSSP

4.1. Use case: *identification of precursors of vibrations on feed-water pump*

Goal: pattern identification that could explain why vibrations are happening on feed-water pump in French nuclear power plants [9].

Class 1:

451 multivariate time series (127 dimensions, 800 points) corresponding to intervals of 80 minutes without any vibration

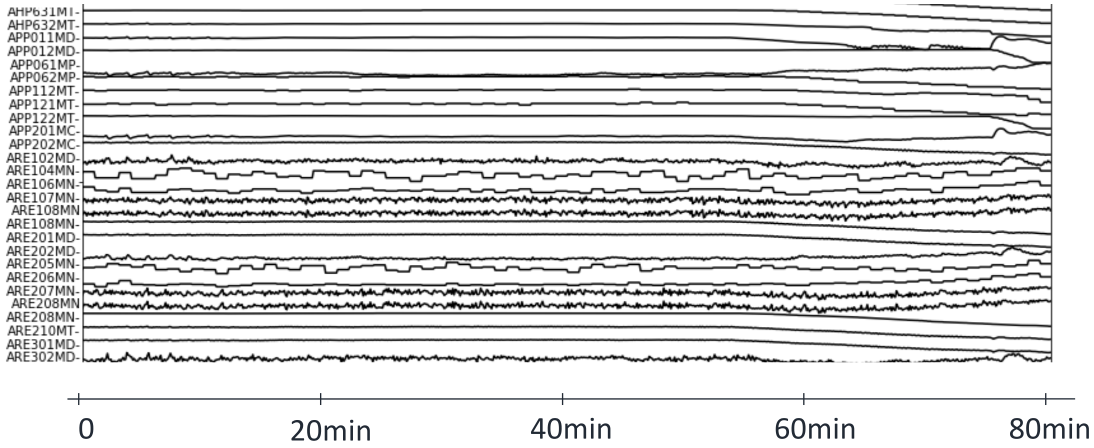

Class 2:

445 multivariate time series (127 dimensions, 800 points) corresponding to intervals of 75 minutes before a vibration and 5 minutes after

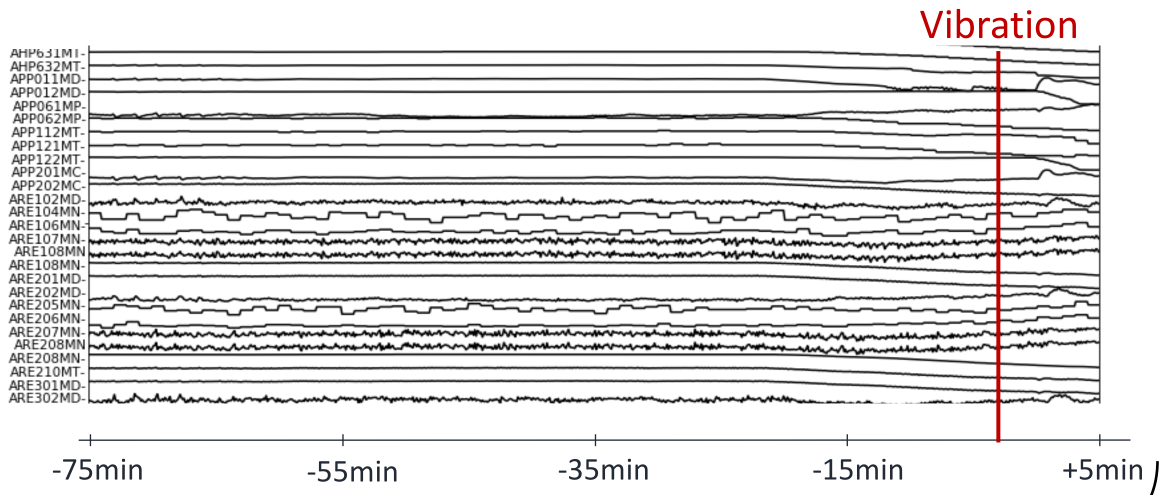

Used to train dCNN model (70/30 split between train/test datasets):

- Classification accuracy: 0.91 (train)/0.89 (test)

[9] P. Boniol et al., dCNN/dCAM: Anomaly Precursors Discovery in Multivariate Time Series with Deep Convolutional Neural Network, in preparation to MSSP

4.2. dCAM example

dCAM:
"Heatmap" that underlines discriminant subsequences for the vibration class.

Multivariate time series: 127 sensors

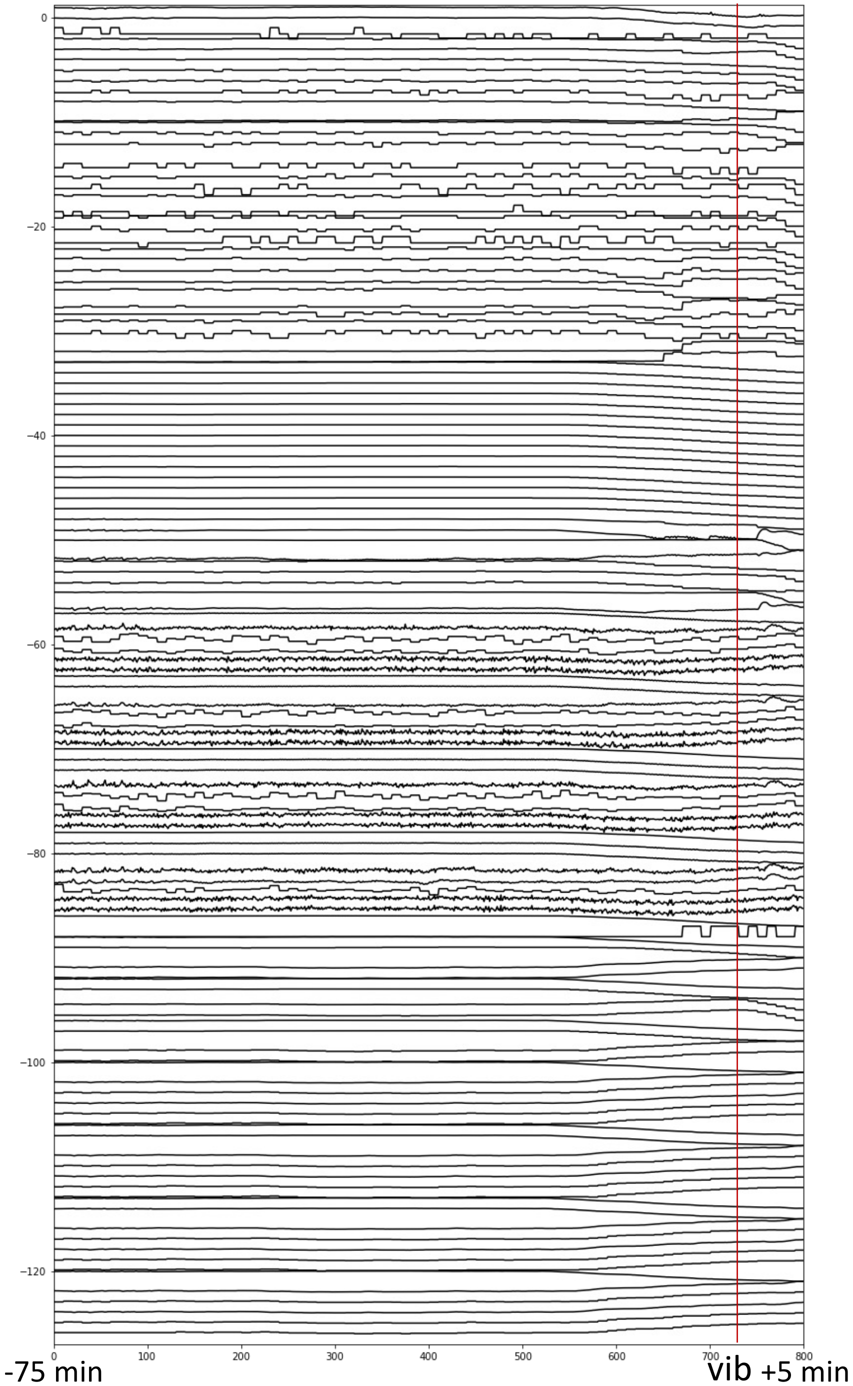

dCAM

4.3. Results analysis: *(1) averaged activation per sensor*

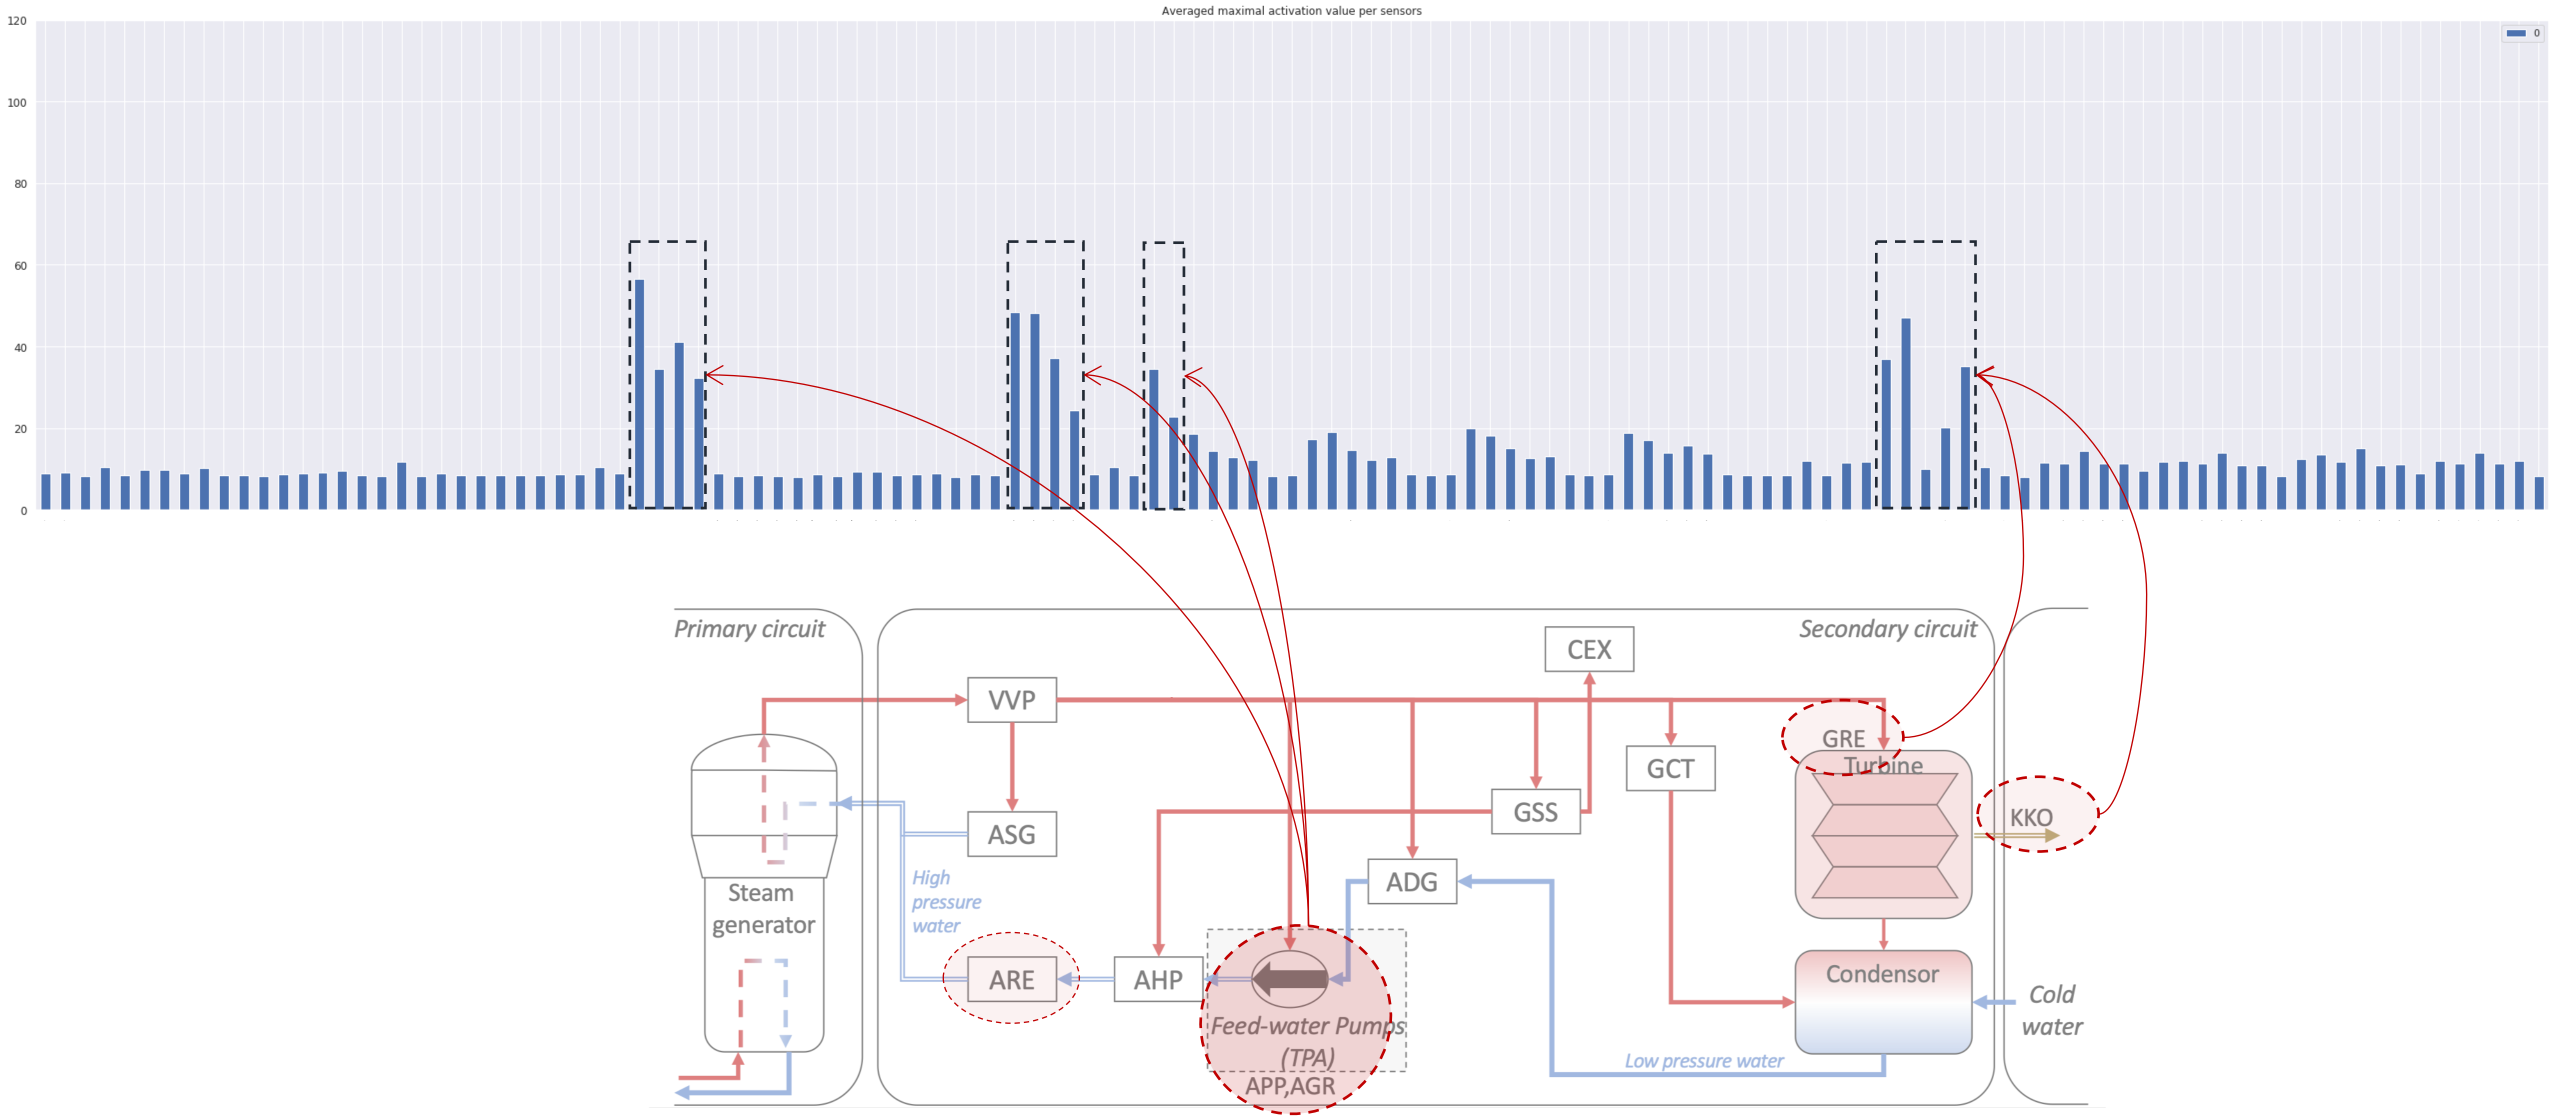

4.3. Results analysis: *(2) averaged activation per temporal index*

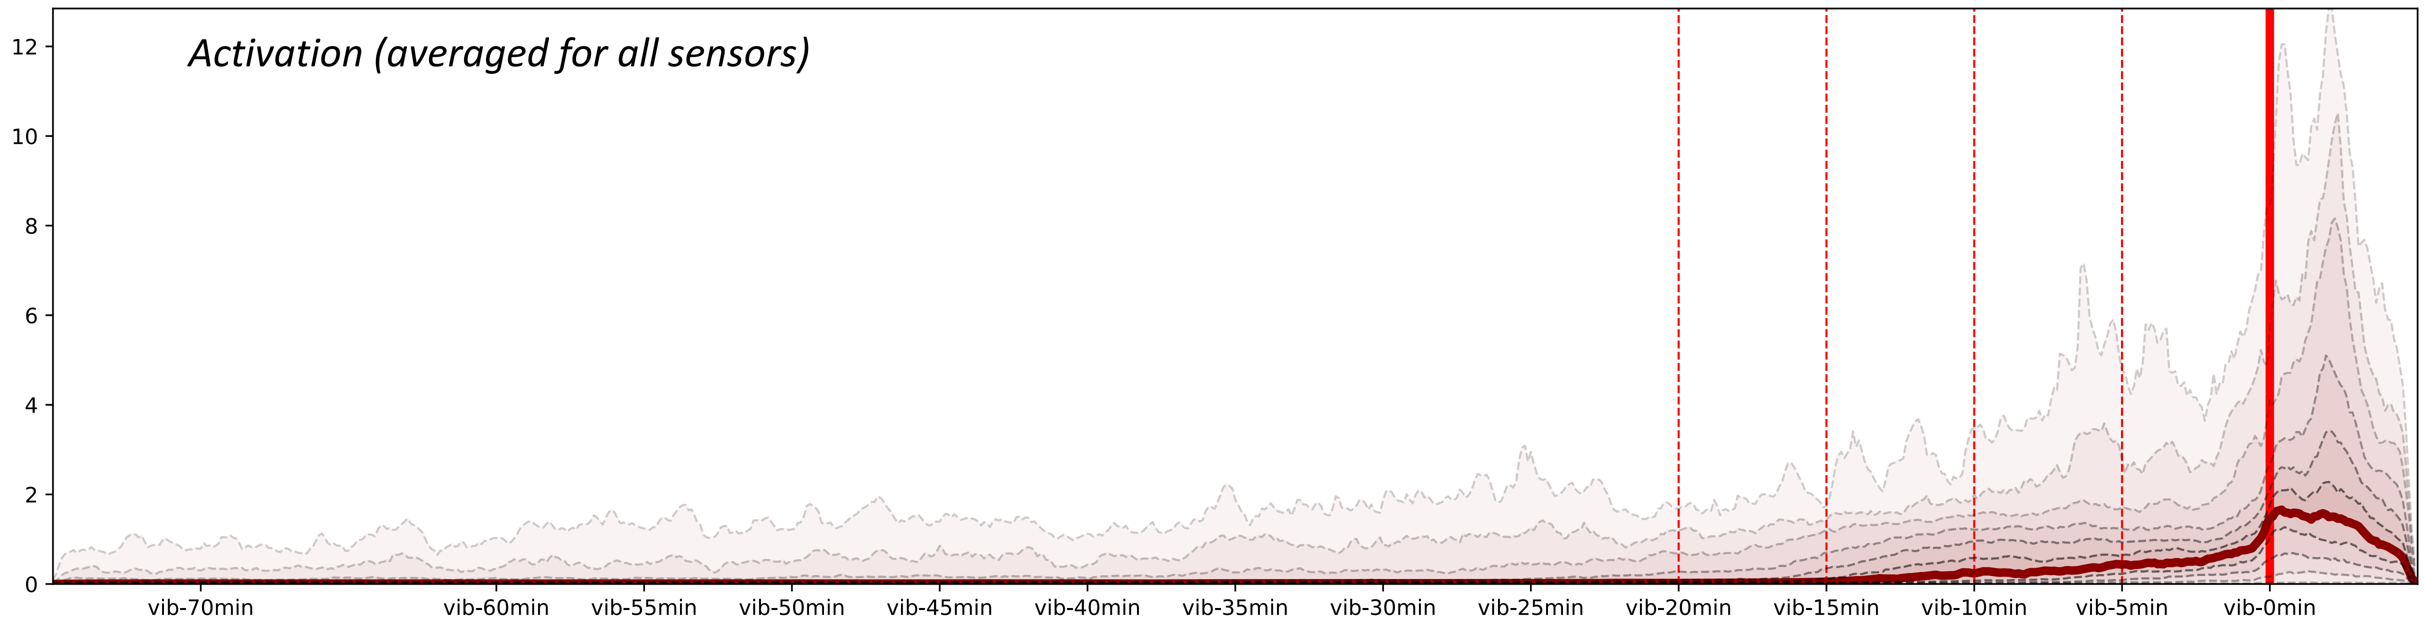

4.3. Results analysis: *(3) types of patterns activated*

For each sensor, clustering of every highly activated subsequences.

sealing temperature TPA1

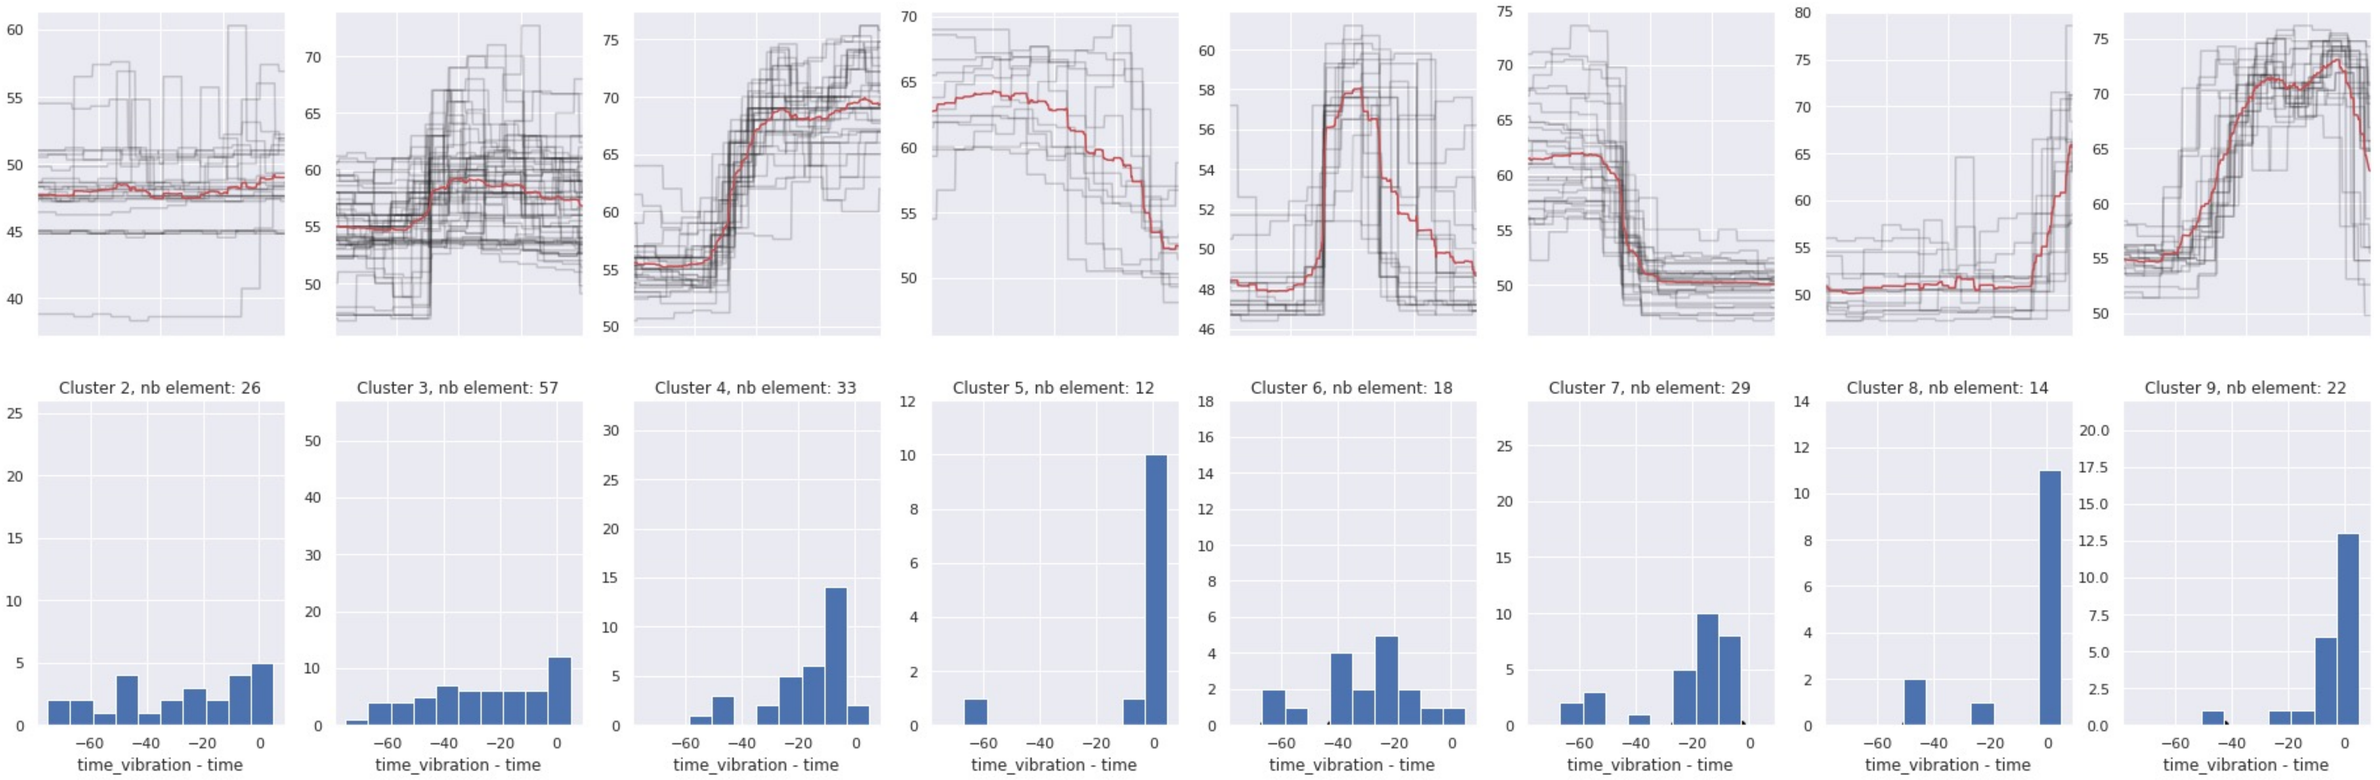

Temporal distribution

4.3. Results analysis: *(3) types of patterns activated*

For each sensor, clustering of every highly activated subsequences.

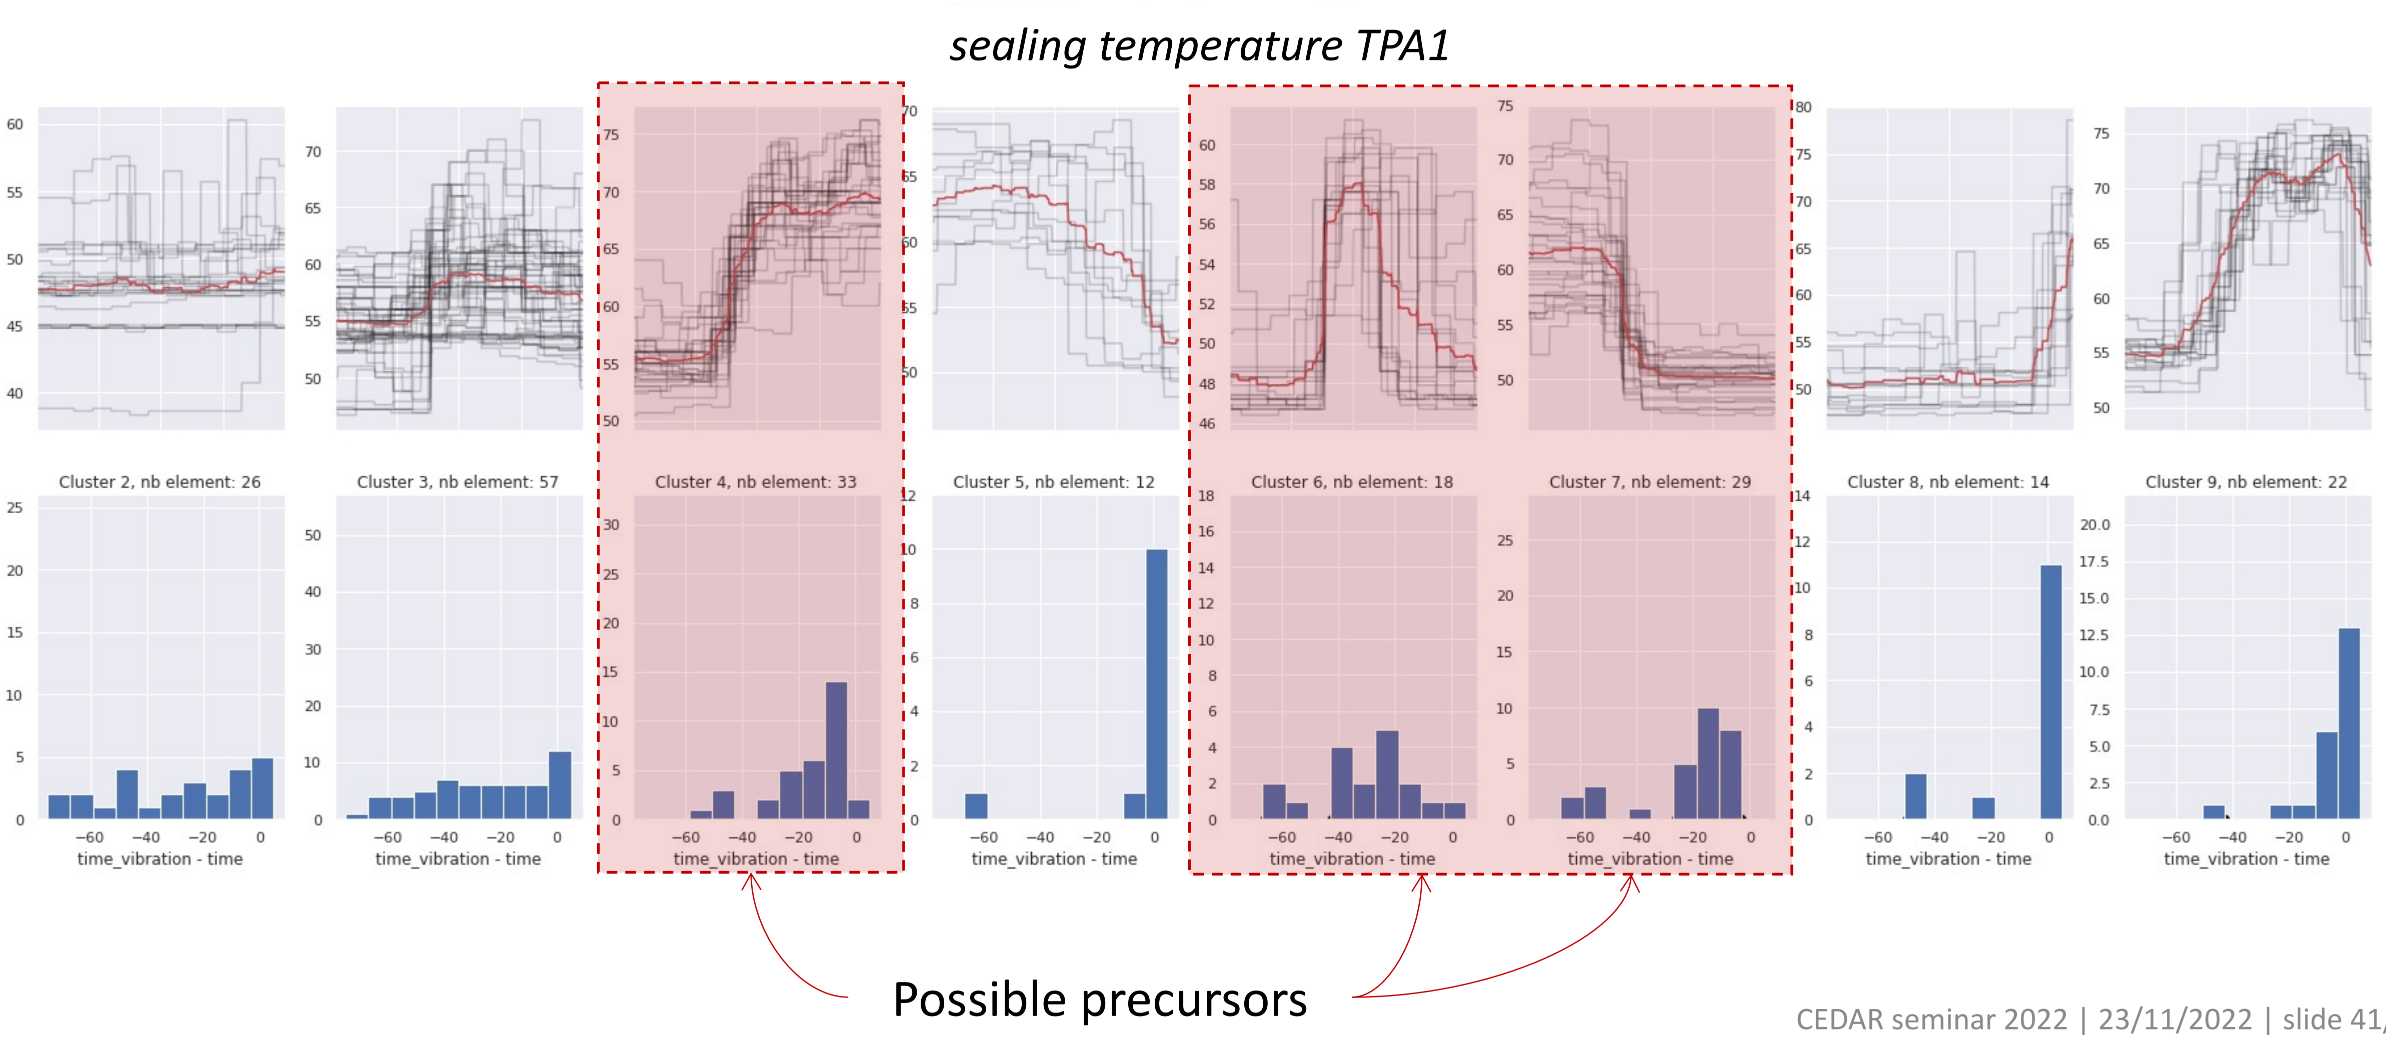

5. Conclusions and open research directions

5.1. Conclusions: *summary of contributions*

We proposed three scalable methods to solve unsupervised subsequence anomaly detection in data series and data streams.

We proposed a new method that adapts the Class Activation Map for multivariate data series.

We studied an industrial use case in which we validated our proposed approaches.

References

[1] E. Keogh et al., HOT SAX: Finding the Most Unusual Time Series Subsequence: Algorithm and Applications, IEEE ICDM, 2005
[2] P. Boniol et al., Series2Graph: Graph-based Subsequence Anomaly Detection in Time Series, PVLDB, 2020
[3] H. I. Fawaz et al., Deep Learning for Time Series Classification: A Review, Data Min. Knowl. Discov., 33, 4, 2019
[4] Z. Wang et al., Time series classification from scratch with deep neural networks: A strong baseline, IJCNN, 2017
[5] B. Zhou et al., Learning Deep Features for Discriminative Localization, CVPR, 2016
[6] H. I. Fawaz et al., Evaluating Surgical Skills from Kinematic Data Using Convolutional Neural Networks, MICCAI, 2018
[7] Sebastian D. Goodfellow et al. Towards Understanding ECG Rhythm Classification Using Convolutional Neural Networks and Attention Mappings. In PMLR (2018)
[8] P. Boniol et al., dCAM: Dimension-Wise Class Activation Map for Explaining Multivariate Data series Classification, Submitted to SIGMOD, 2022
[9] P. Boniol et al., dCNN/dCAM: Anomaly Precursors Discovery in Multivariate Time Series with Deep Convolutional Neural Network, in preparation to MSSP
[10] Sebastian Schmidl, Phillip Wenig, Thorsten Papenbrock: Anomaly Detection in Time Series: A Comprehensive Evaluation. PVLDB, 2022
[11] John Paparrizos, Yuhao Kang, Paul Boniol, Ruey S. Tsay, Themis Palpanas, Michael J. Franklin: TSB-UAD: An End-to-End Benchmark Suite for Univariate Time-Series
Detection. PVLDB, 2022

Thank you!
Any questions?