Doctoral Candidate 2

Anomaly Detection for Time Series

Leonardo Schiavo

Departamento de Sistemas Informáticos,
A.I.D.A. - Applied Intelligence and Data Analysis research group,
Universidad Politécnica de Madrid, Spain

July 22, 2026

AI+DA

POLITÉCNICA

This work was supported by the **MSCA Doctoral Networks project T.U.A.I. - Towards an Understanding of Artificial Intelligence via a transparent, open and explainable perspective (TUAI) project**, *N°101168344*.

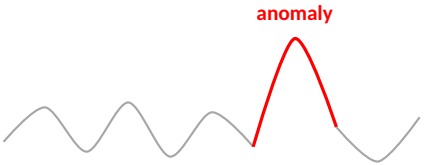
Table of Contents

- Introduction
- Methodology
- Benchmarks
- Federation
- Conclusions and Future Directions

What is time-series anomaly detection?

Goal: identify observations or temporal patterns that deviate from expected behaviour.

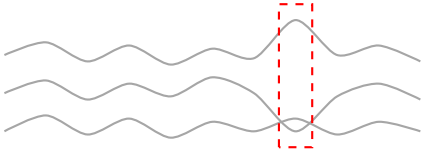
Univariate setting



One variable observed over time.

An anomaly is identified from its **temporal behaviour**.

Multivariate setting



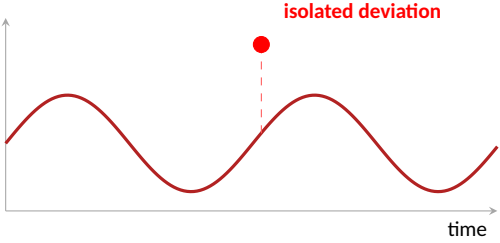
Multiple variables observed simultaneously.

Anomalies may involve individual channels or **dependencies between channels**.

Types of Anomalies [1]

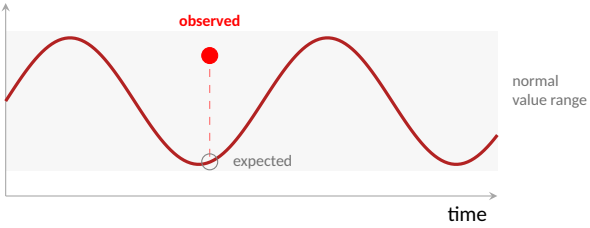
Point anomalies

A single timestamp strongly deviates from the normal behaviour of the signal. In the multivariate case, the anomaly may affect one or multiple channels.



Contextual anomalies

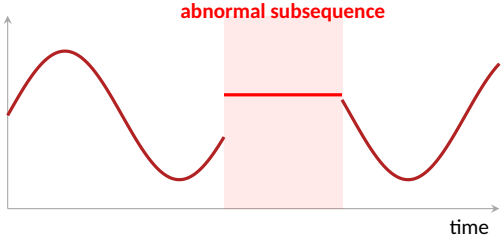
A value is globally plausible but anomalous in its specific temporal context, such as the current phase, season, or operating regime.



Types of Anomalies [2]

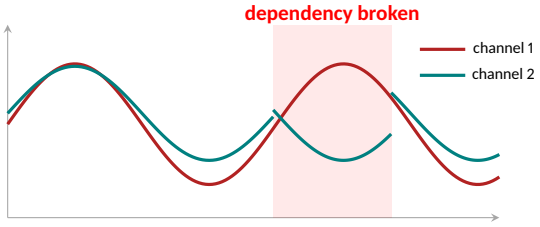
Sequential anomalies
(collective or subsequence anomalies)

An entire interval exhibits abnormal temporal dynamics. Individual values may appear plausible, while the sequence as a whole is not.



Correlation anomalies

Individual channels remain within their normal ranges, but their usual cross-channel relationship is violated. This anomaly may therefore be invisible when channels are analysed independently.



Families of TS Anomaly Detectors

- **Statistical and decomposition-based detectors** model normality through distributional assumptions, trend, seasonality, and residuals (*EWMA*, *Holt-Winters*, *STL-ESD*).
- **Distance- and density-based detectors** identify anomalies as observations or subsequences that are far from normal patterns or lie in low-density regions (*k-NN*, *LOF*, *Matrix Profile*).
- **Change-point and state-space detectors** detect anomalies as regime changes or unlikely observations under a latent dynamical model (*CUSUM*, *BOCPD*, *Kalman Filter*).
- **Forecasting-based detectors** learn to predict normal temporal evolution and use large prediction errors as anomaly scores (*ARIMA*, *Prophet*, *LSTM-AD*).
- **Reconstruction-based detectors** learn a compact representation of normal windows and flag samples with high reconstruction error (*LSTM Encoder-Decoder*, *USAD*).
- **Generative and probabilistic detectors** estimate the likelihood of a window under a learned normal data distribution (*Donut*, *OmniAnomaly*, *InterFusion*).
- **Graph-based detectors** model cross-channel dependencies and detect violations of normal inter-variable relationships (*GDN*, *MTAD-GAT*).

Future Directions in TSAD&XAI

- **Adaptive and online detection** to handle evolving normality, concept drift, streaming data, and adaptive thresholding.
- **Truly multivariate anomaly detection** to capture genuine cross-channel dependencies and evaluate models on realistic multivariate anomalies.
- **Actionable explanations** that combine temporal localization, responsible variables, normal references, and plausible counterfactuals.
- **Reliable evaluation of detection and explanations** through realistic benchmarks, robust metrics, principled thresholding, and systematic validation of explanation quality.
- **Causal diagnosis and root-cause analysis** to move beyond feature attribution toward identifying anomaly sources and their propagation across variables.
- **Scalable and data-efficient learning** through lightweight models, uncertainty-aware methods, few-shot and foundation models, and privacy-preserving learning.

Table of Contents

- Introduction
- Methodology
- Benchmarks
- Federation
- Conclusions and Future Directions

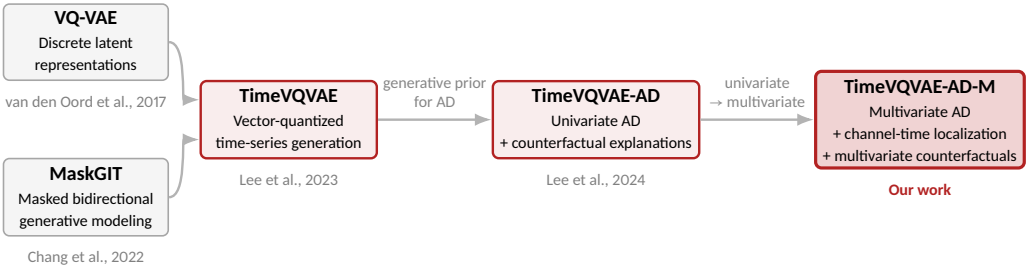
TimeVQVAE-AD-M

Time Vector Quantized Anomaly Detection Multivariate

- **Multivariate time-series**
- **Generative and probabilistic**
- **Discrete representation learning**
- **Counterfactual explainability**

Based on TimVQVAE-AD Univariate by Daesoo Lee, Sara Malacarne, Erlend Aune **??**

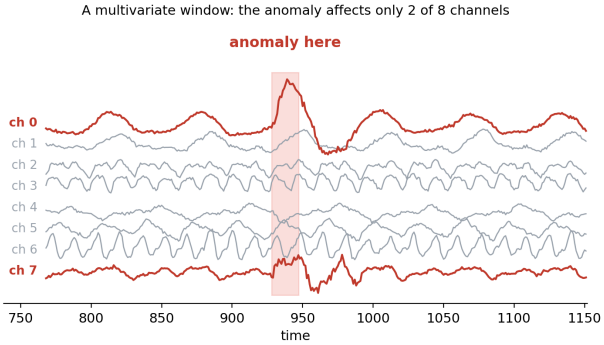
From Foundations to TimeVQVAE-AD-M



Wouldn't it be nice?

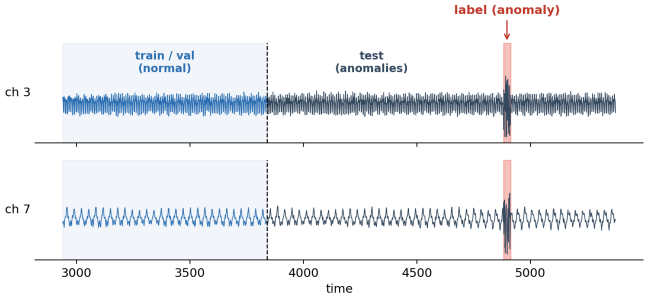
A single framework that tells you:

- **Where?** → which channels
- **When?** → which time interval
- **What if?** → a counterfactual
  (the signal without the anomaly)



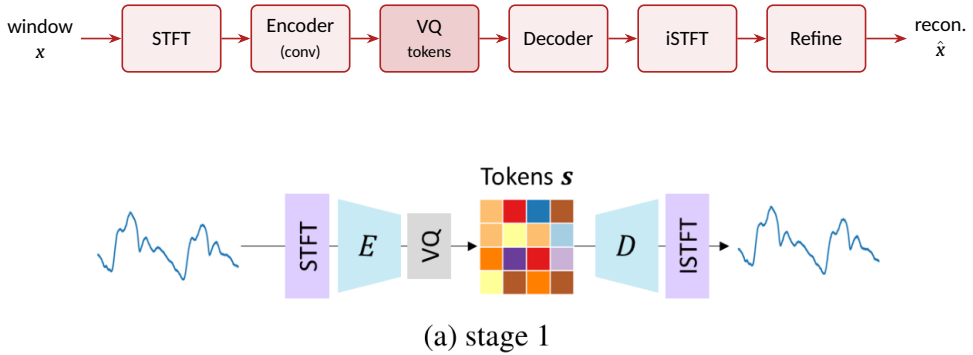
Premises: an unsupervised model

- Training is **unsupervised**: the model only learns normal behaviour.
- Labels exist **only in test**: used to evaluate, never to train.



Stage 1: The tokenizer

Two stages. Stage 1 is a tokenizer: a VQ-VAE in the time frequency domain.



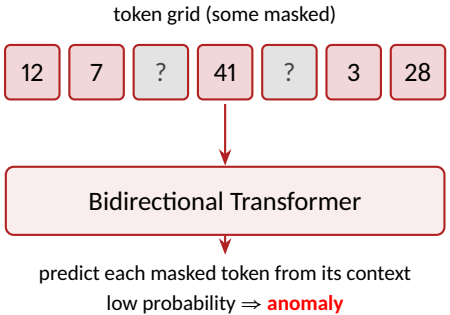
Stage 1 of the original TimeVQVAE-AD (Lee et al., 2024).

Stage 1: Structured tokens

- Convolutions act **only along time** → no mixing across channels and frequencies.
- One token for every (channel, frequency, time) coordinate.
- **Shared codebook** across channels → the same token id is comparable between variables (it captures values that normally occur together).

Stage 2: The prior over tokens

- Stage 1 is **frozen**: a window is now a grid of tokens.
- A **bidirectional transformer** (MaskGIT) learns the normal token distribution.
- Trained by **masking**: predict hidden tokens from the visible context.



Detection and counterfactual

Detection

- how likely is each token, given its context?
- low probability → **anomaly**, kept **per channel**.

Counterfactual

- regenerate only the most surprising tokens;
- get an **anomaly free** version of the signal.

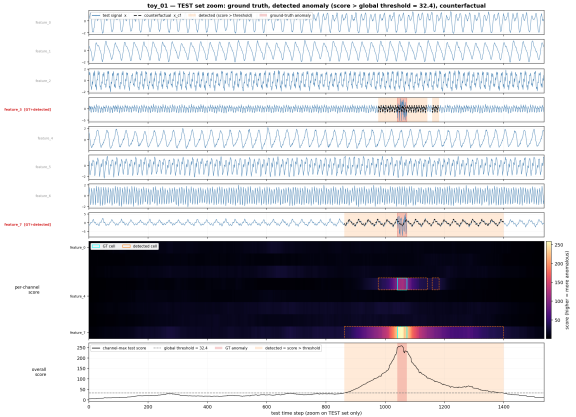


Table of Contents

- Introduction
- Methodology
- Benchmarks
- Federation
- Conclusions and Future Directions

Why do we need a new benchmark?

When? **Which channels?** **What should be there instead?**

- Many existing multivariate benchmarks provide mainly **timestamp-level anomaly labels**.
- This is enough to evaluate **temporal detection**, but not precise **channel attribution**.
- Real-world datasets rarely provide exact information about **which variables were truly affected** by each anomaly.
- Without this ground truth, localization and explanations cannot be evaluated quantitatively.

What Does the Benchmark Test?

8 main anomaly families · 26 anomaly types

Per-channel anomalies

- Point
- Amplitude
- Level
- Pattern
- Morphological
- Categorical

Stress the ability to model **normal temporal behaviour within individual variables**.

Cross-channel anomalies

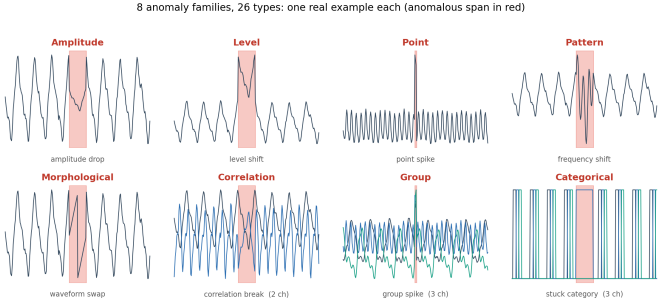
- Correlation
- Group

Stress the ability to model **dependencies between variables**.

These anomalies may be invisible when each channel is analysed independently.

Synthetic benchmark

- **9 anomaly families, 26 anomaly types.**
- per channel *shape* anomalies vs cross channel *dependency* anomalies.



Metric Evaluation

Time-series anomaly detectors usually output an **anomaly score** for each timestamp:

$$s_t = \text{degree of abnormality at time } t.$$

To evaluate the model, scores are compared with ground-truth labels.

- **Point-wise evaluation**
  Each timestamp is treated independently:

  AUROC, AUPRC, Precision, Recall, F1.

- **Event / range evaluation**
  Anomalies are evaluated as temporal intervals, not isolated points.

- **Multivariate evaluation**
  Besides detecting *when* the anomaly happens, we may also evaluate *which variables/channels* are responsible.

The Problem of Point Adjustment

Point Adjustment: if the detector fires at least once inside a ground-truth anomaly window, the whole window is counted as detected.

- It rewards **existence** of a hit, not precise localization.
- A detector can fire only once in a long anomaly and receive full credit.
- Late or temporally imprecise detections are weakly penalized.
- Scores can become artificially high and model rankings misleading.

Better practice:

- report raw point-wise AUROC/AUPRC;
- add range-aware or affiliation-based metrics;
- measure detection delay and false alarm rate;
- avoid relying on PA-F1 as the main result.

Counterfactual Evaluation

A counterfactual asks:

What minimal change would make this anomalous window normal?

Counterfactuals are evaluated along four dimensions:

- **Validity** — does it flip the decision (anomalous → normal)?
- **Proximity & sparsity** — is the change small, touching few timestamps or channels?
- **Plausibility** — does it still look like a realistic series (temporal & cross-channel dynamics)?
- **Usefulness** — does it point to interpretable channel-time regions?

A good counterfactual is valid, minimal, plausible, and informative.

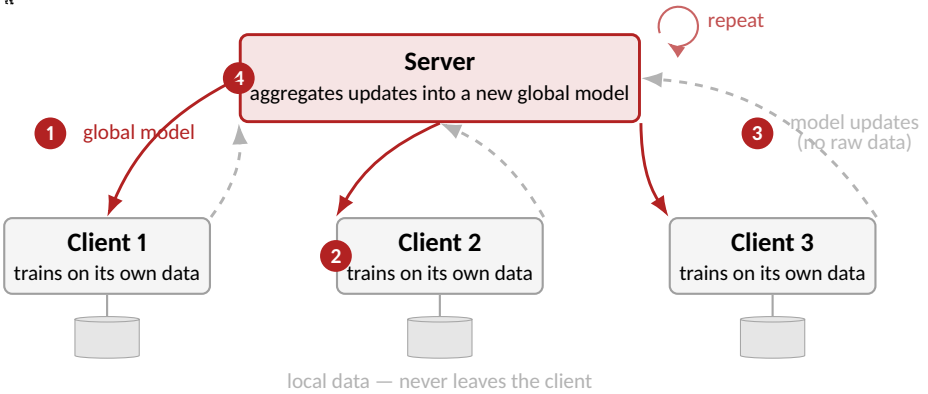


Table of Contents



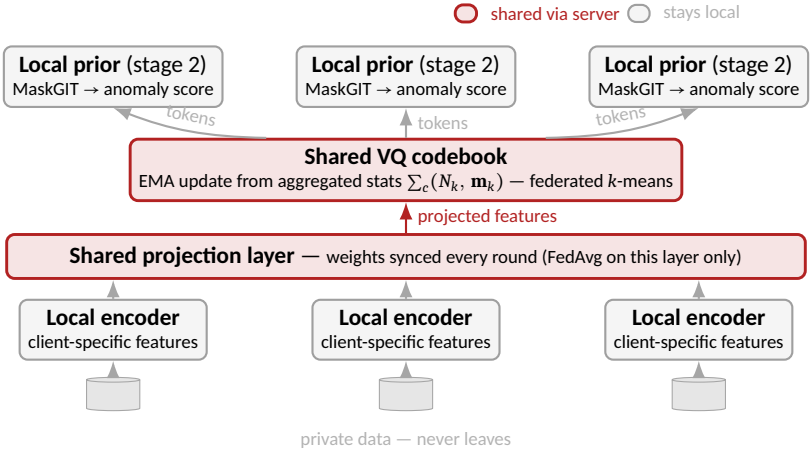
- Introduction
- Methodology
- Benchmarks
- Federation
- Conclusions and Future Directions

What is Federated Learning?



One round: (1) server sends the current model → **(2)** each client trains it locally → **(3)** clients send back only their updates → (4) server averages them into a new global model. Repeat.

Federating TimeVQVAE-AD



Shared = the token space (projection + codebook anchor all clients to one vocabulary). **Local = features and "normal"** (each ent



Table of Contents

AI+DA

- Introduction
- Methodology
- Benchmarks
- Federation
- Conclusions and Future Directions

Future work

- **Takeaway:** one framework that detects, localizes per channel, and explains via counterfactuals, all from the same prior.

Current limitations

- evaluated only on the controlled synthetic benchmark;
- very short *point* anomalies are smoothed and not well detected;
- the fixed per channel threshold is suboptimal.

Next steps

- extend to **real** multivariate benchmarks (SMD, MSL, SWaT, ESA-ADB);
- reduce the number of parameter and move to lightweight
- calibrate / learn the detection threshold;
- adopt online learning approaches

AI+DA | POLITÉCNICA

Anomaly Detection for Time Series

Thank you for listening!

 leonardo.schiavo@upm.es

 0009-0007-5023-0946