Modeling Multivariate Sensor Dynamics with STGNNs: An Anomaly Detection Study at the INFN CNAF Tier-1 Facility

Giovanni Zurlo, Elisabetta Ronchieri
L. Scarponi, D. Cesini, A. Costantini *(INFN-CNAF)*

International Symposium on Grids & Clouds (*ISGC*) 2026 | Taipei, 20 March 2026

Case Study

Anomaly Detection in a Data Center

[1]

In mission-critical industries, continuous operation exposes equipment to wear that can escalate into costly failures and downtimes.

Raw data from IoT sensors enables real-time monitoring and predictive maintenance (PdM) of systems.

↓

The goal is to perform maintenance when the maintenance activity is most **cost-effective** and **before the system loses performance**, going below an acceptable threshold.

High Dimensionality

Hundreds of concurrent sensors recording temperature, pressure, and power.

Complex Interdependencies

Signals do not act in isolation; a subtle drop in voltage directly affects cooling pressures.

Subtle Deviations

Catastrophic failures often begin as minor fluctuations completely hidden within the noise of standard operations.

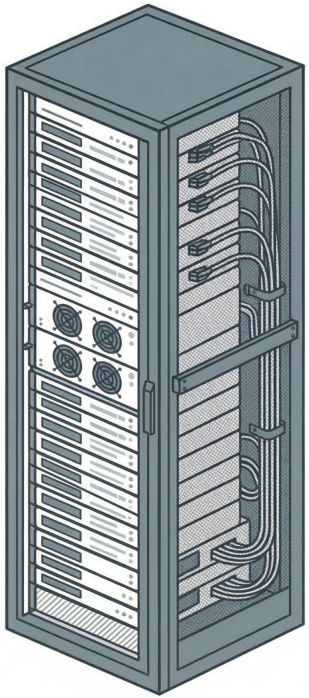

INFN CNAF data center:

Scale & Space: Over 2,000 m² of usable IT area, plus a dedicated expansion zone of up to 800 m² for future growth.

Power & Efficiency: Initial IT power of 3 MW, scaling up to 10 MW by 2026

High-Density Computing (200 m²): Utilizes Direct Liquid Cooling (DLC) to support extreme densities of up to 80 kW per rack.

Low-Density Resources (740 m²): Supports up to 16 kW per rack via air cooling; 150 racks are currently.

Massive Tape Storage (170 m²) Designed to host up to 4 tape libraries containing 32,800 slots, roughly 1.5 Exabytes (EB) of capacity.

Data Sources

71 multivariate time-series tracking the facility's electrical and mechanical systems, recorded at 15-minute intervals.

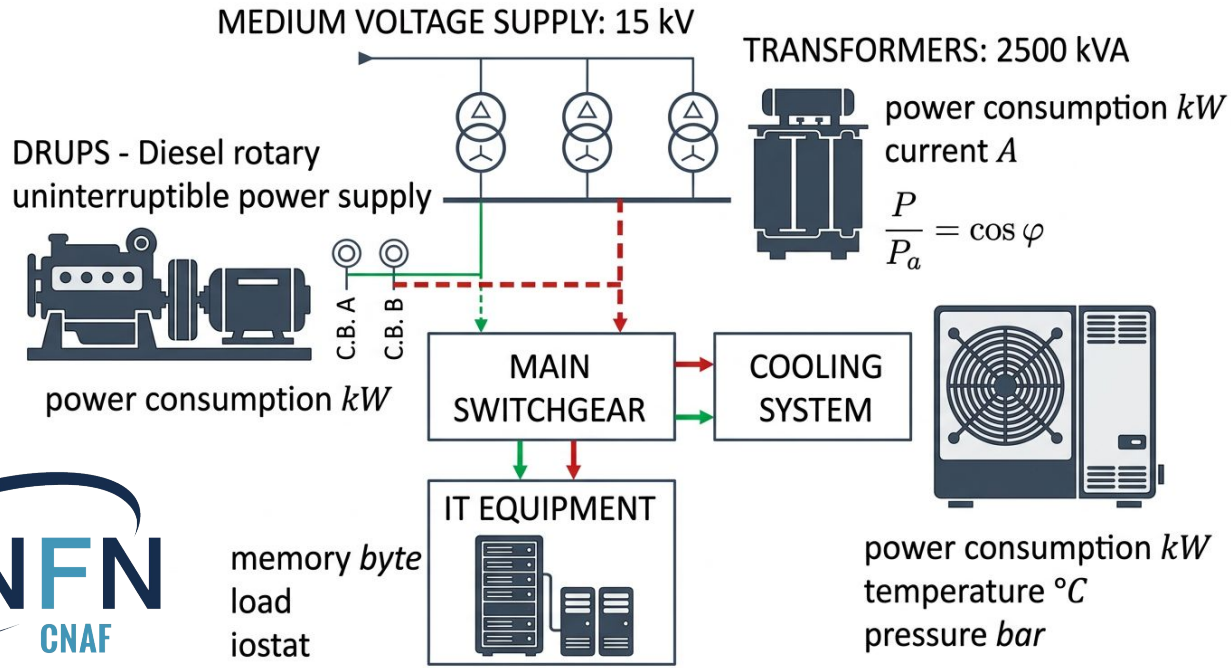

INFN CNAF

Previous Research on Graph Anomaly Detection

From Time-Series to Graphs

[2]

We followed a correlation network approach yielding to a single-layer of variations-based nodes

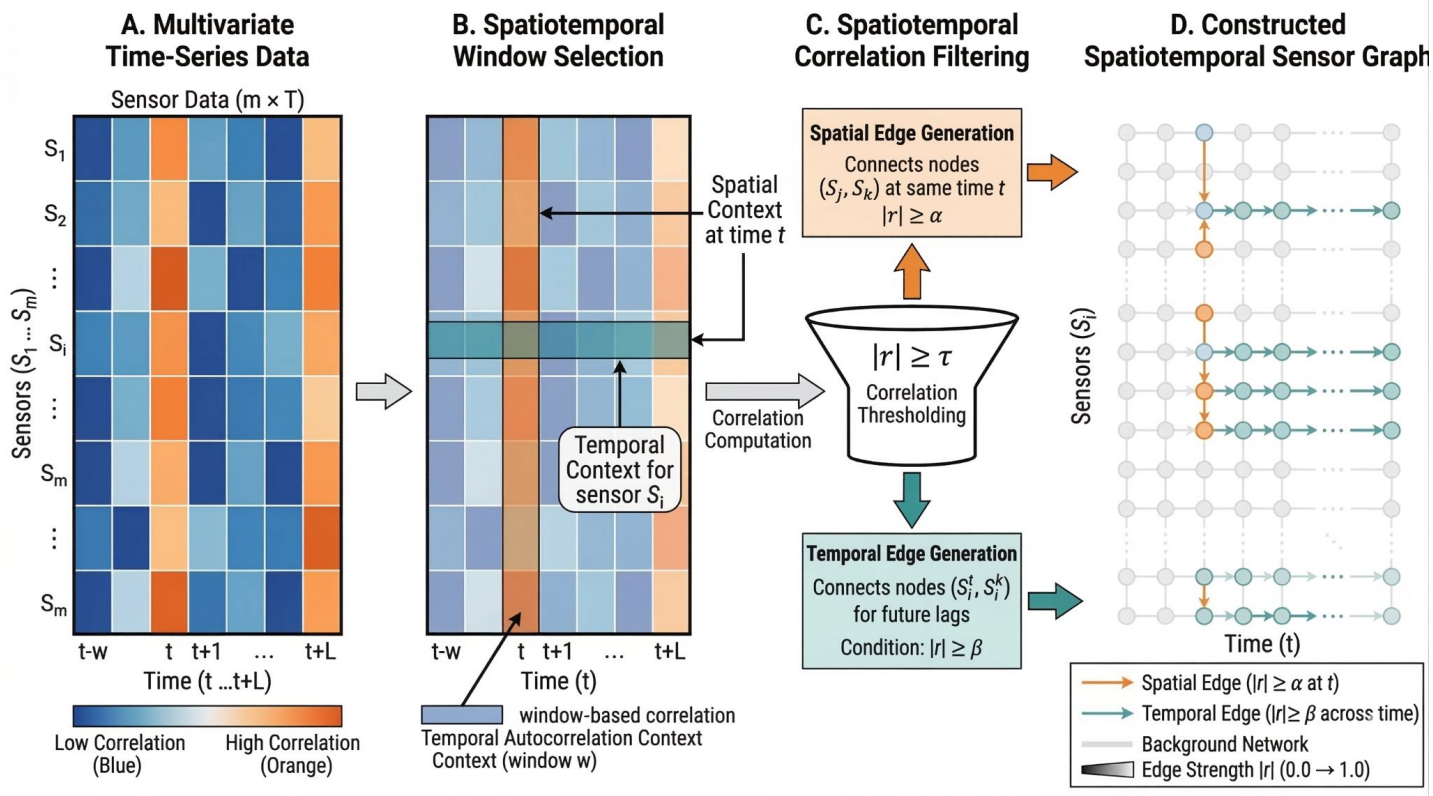

Node-Level GNN Model

The proposed hybrid architecture combines **GNNs Variational Autoencoding** to compress spatial and temporal relationships into a structured latent space
and a **Dual-GAN framework** to robustly regularize the latent representations and ensure highly realistic graph reconstructions

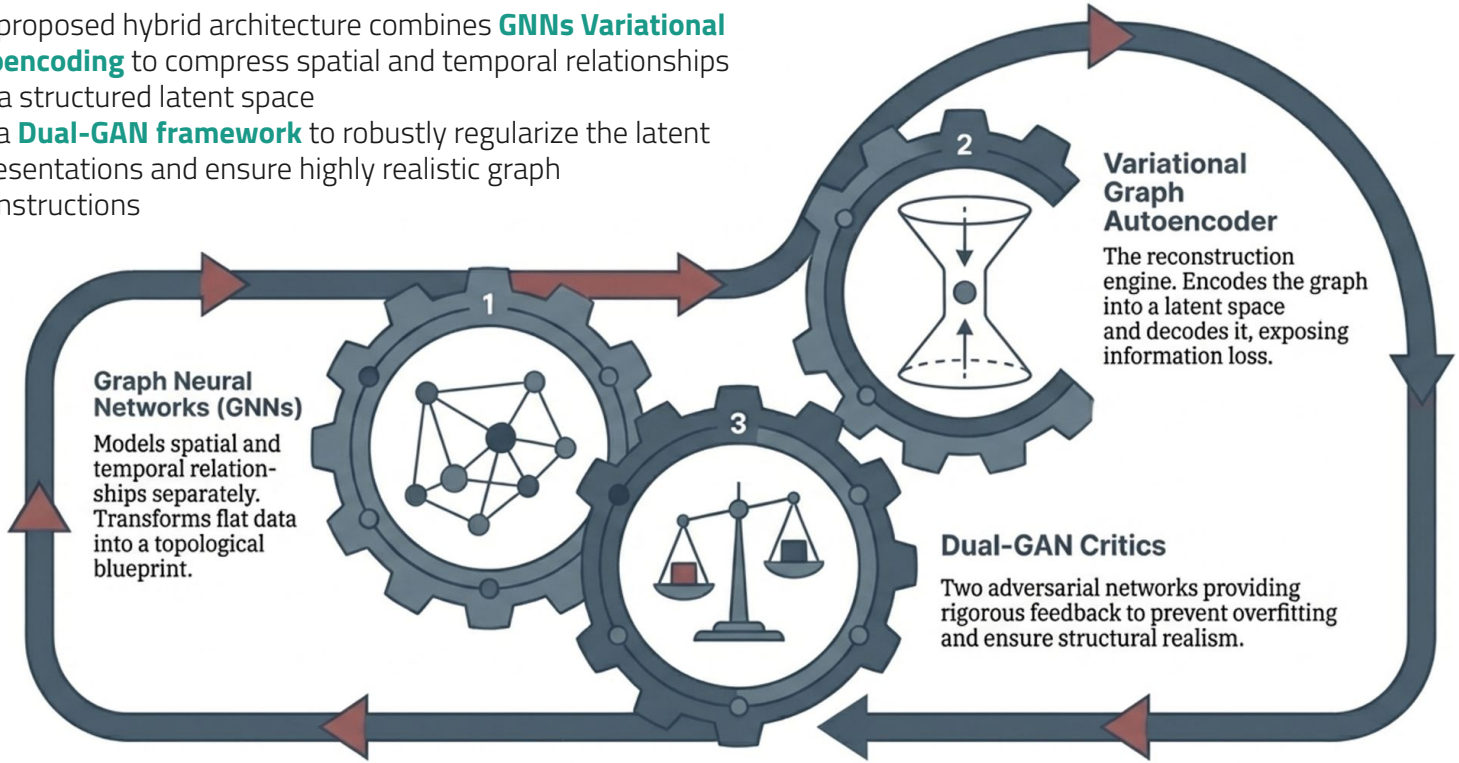

Node-Level GNN Model

[3]

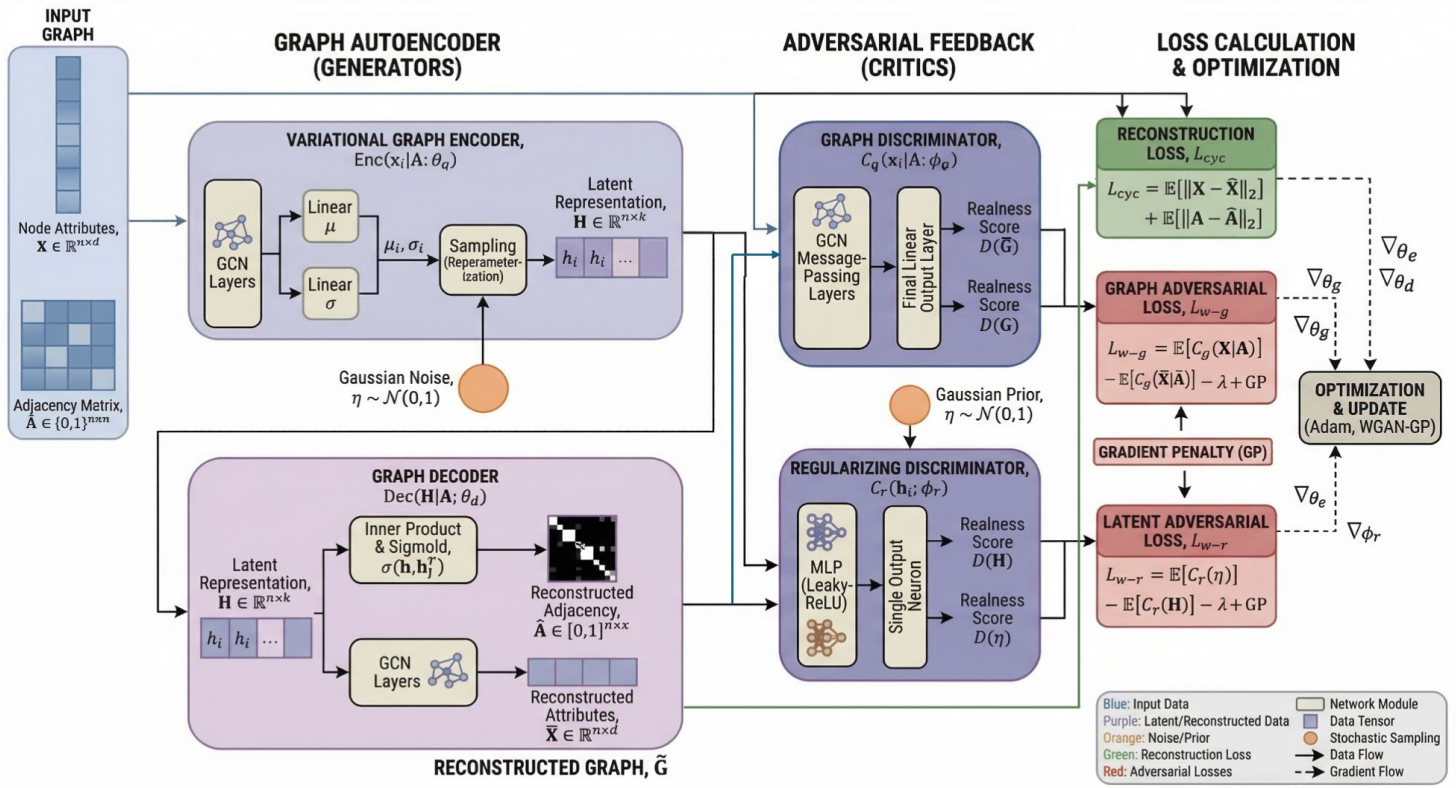

Anomaly Scores

Averaged Node Reconstruction Errors **(Top)** and Averaged Edges Reconstruction Errors **(Bottom)**

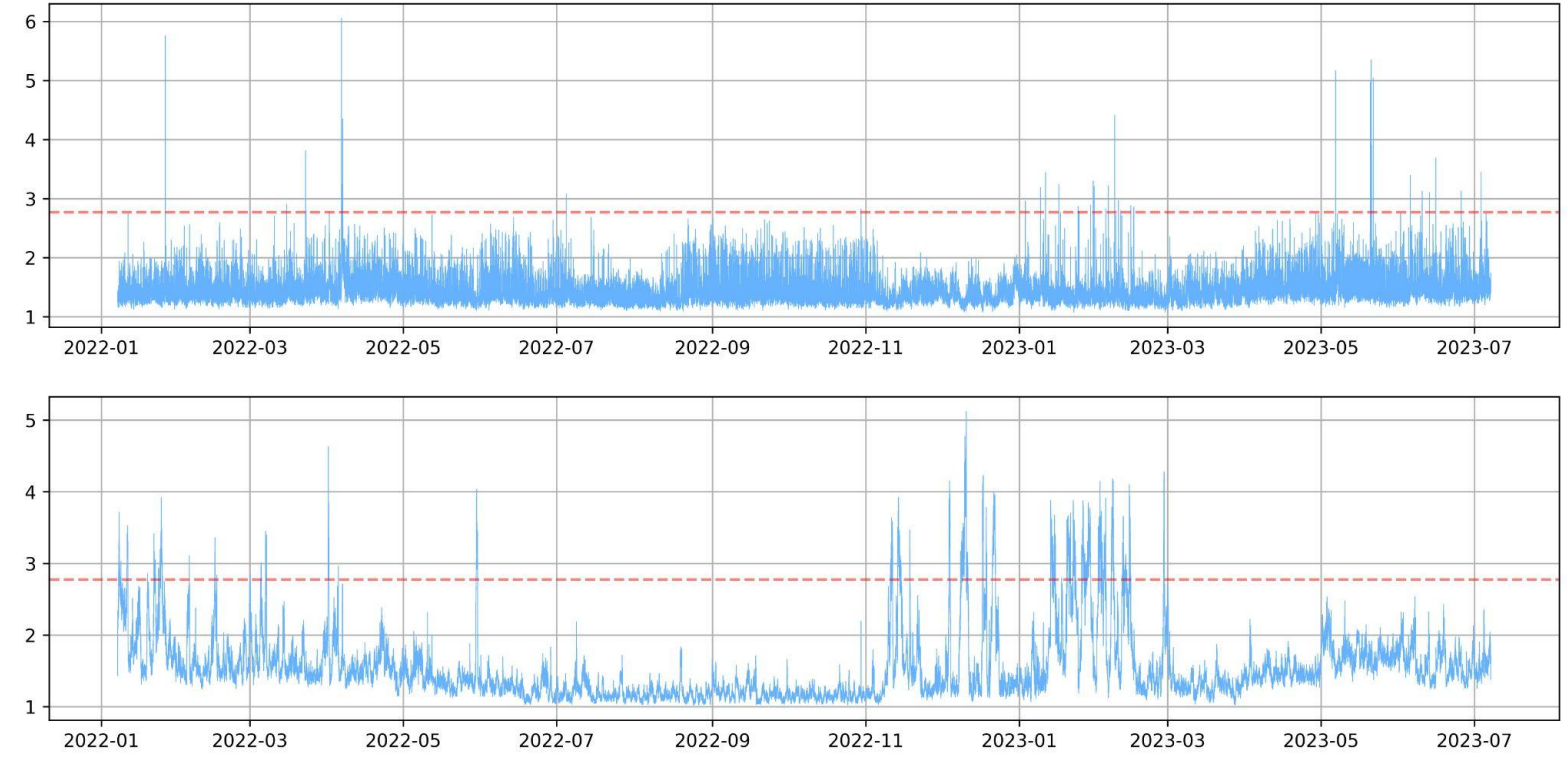

Model Predictions

Detected anomalies in the current measurements from three 2,500 kVA transformers operating under N+1 redundancy

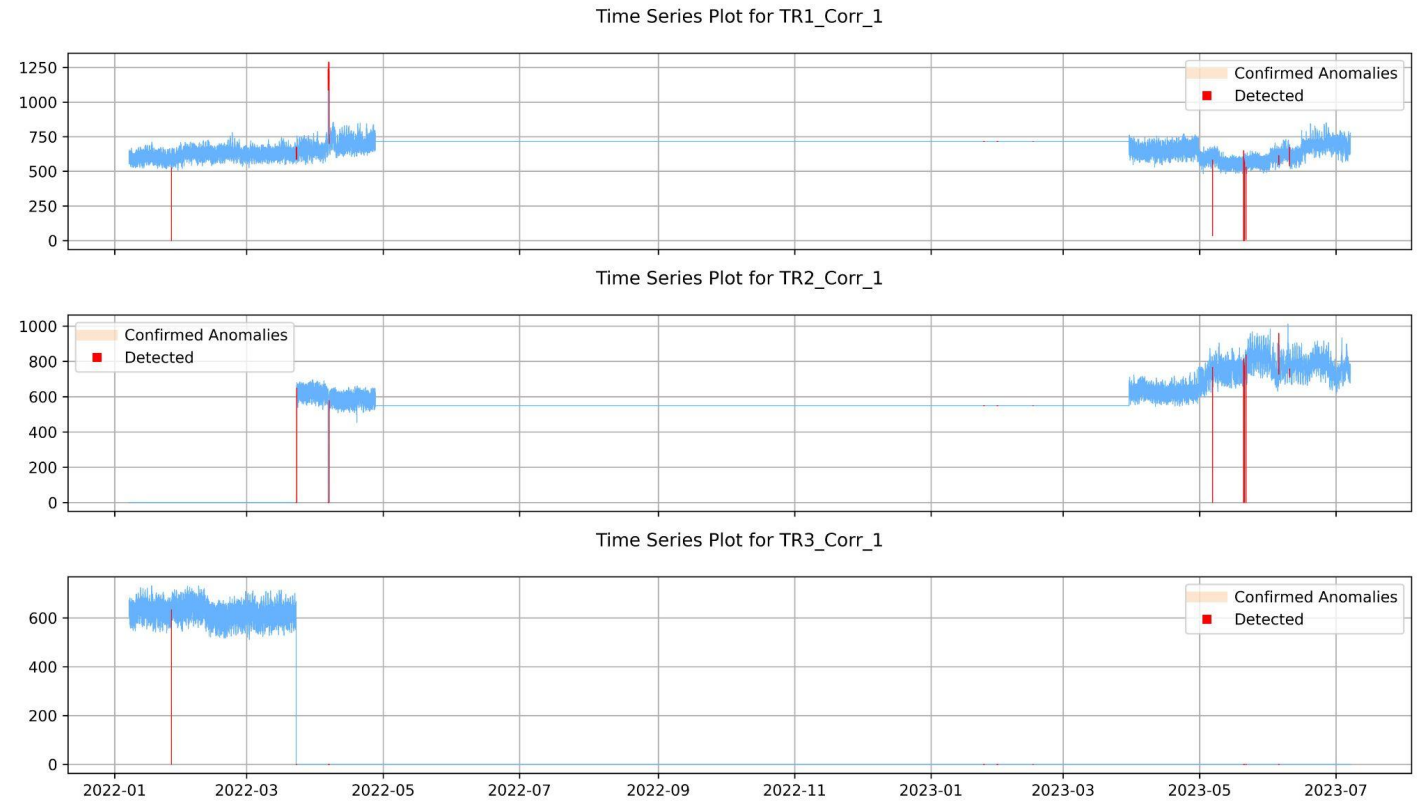

Interpretability via Subgraph Inspection

Each node represents an entry and its color depends on the signal it belongs to. Red nodes and edges are anomalous.

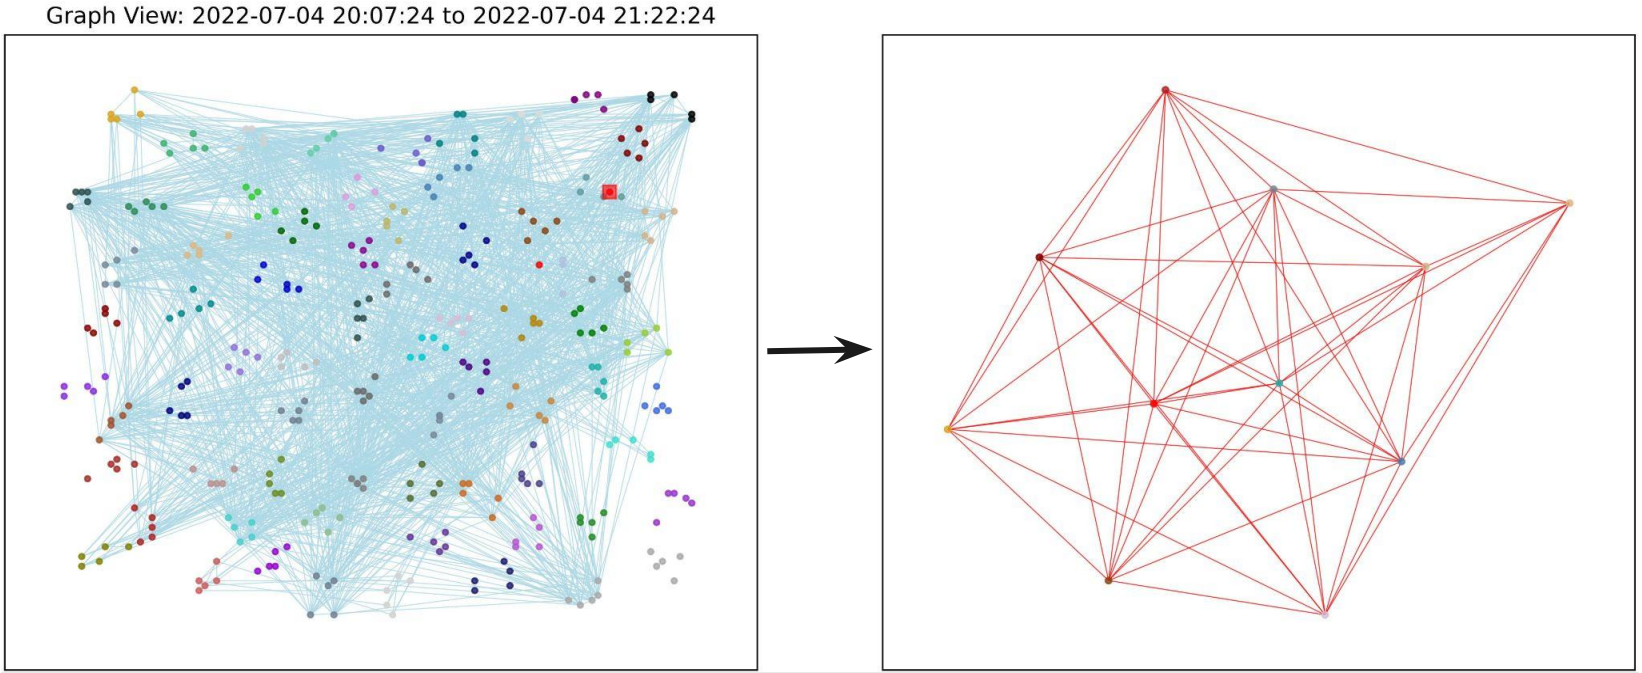

Interpretability via Subgraph Inspection

Each node represents an entry and its color depends on the signal it belongs to. Red nodes and edges are anomalous.

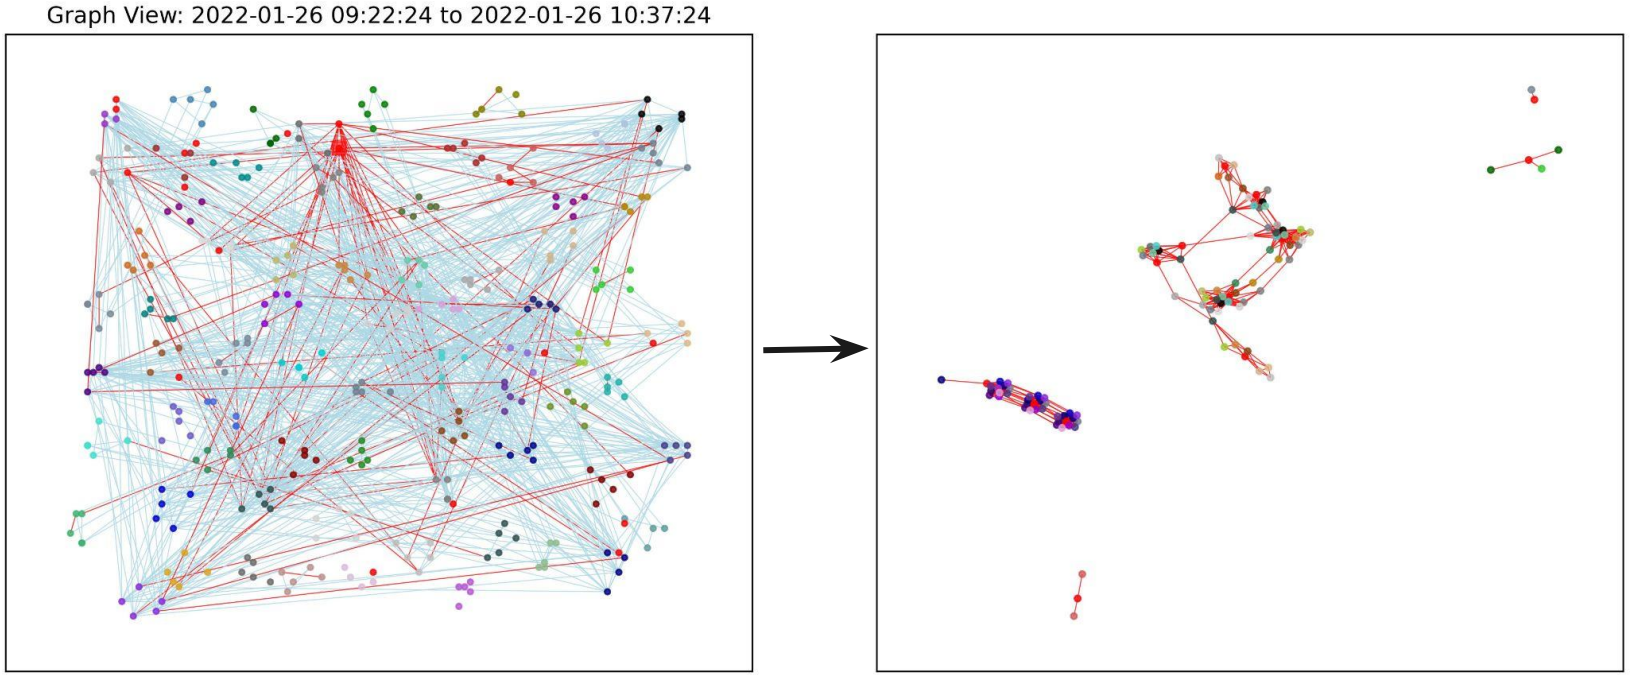

Anomaly Detection with Space-Temporal Graphs

Spatial-Temporal Dynamic Graphs

Temporal Graphs provide a dynamic representation of relationships and interactions between entities over time

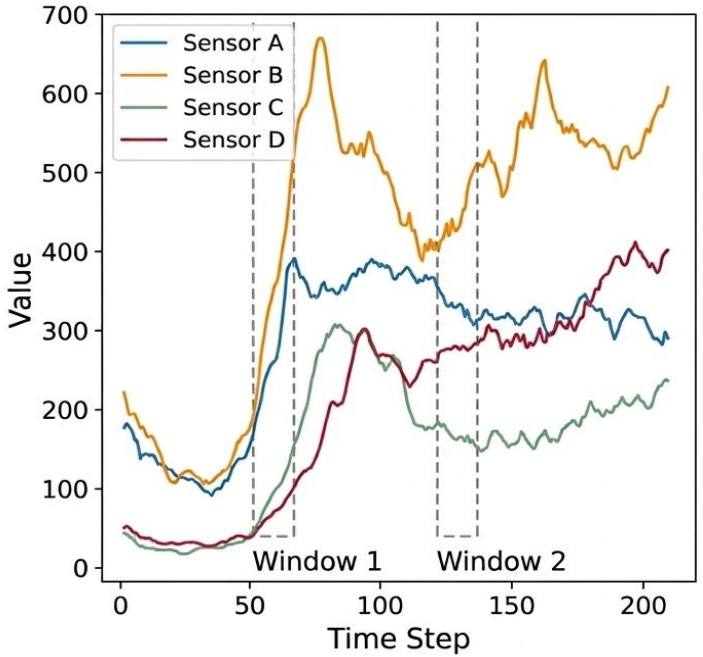

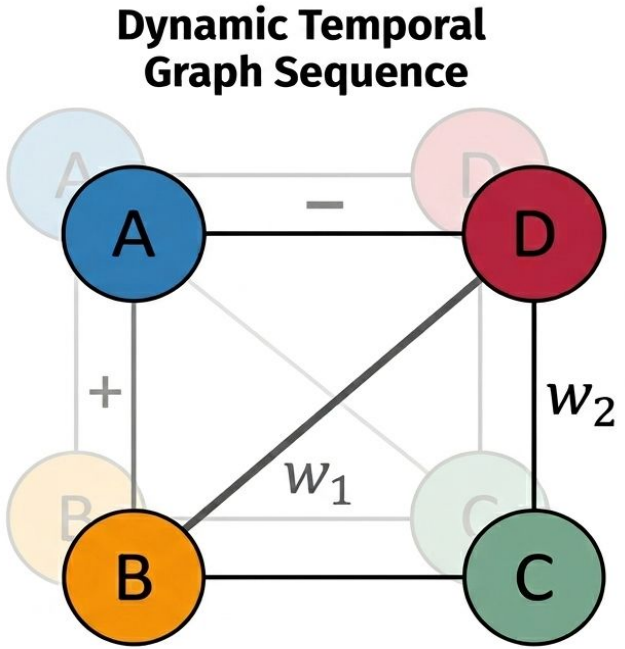

From Time-Series to Graphs

Each timestep's sensor readings are mapped onto the nodes of a dynamic graph whose topology evolves over time.

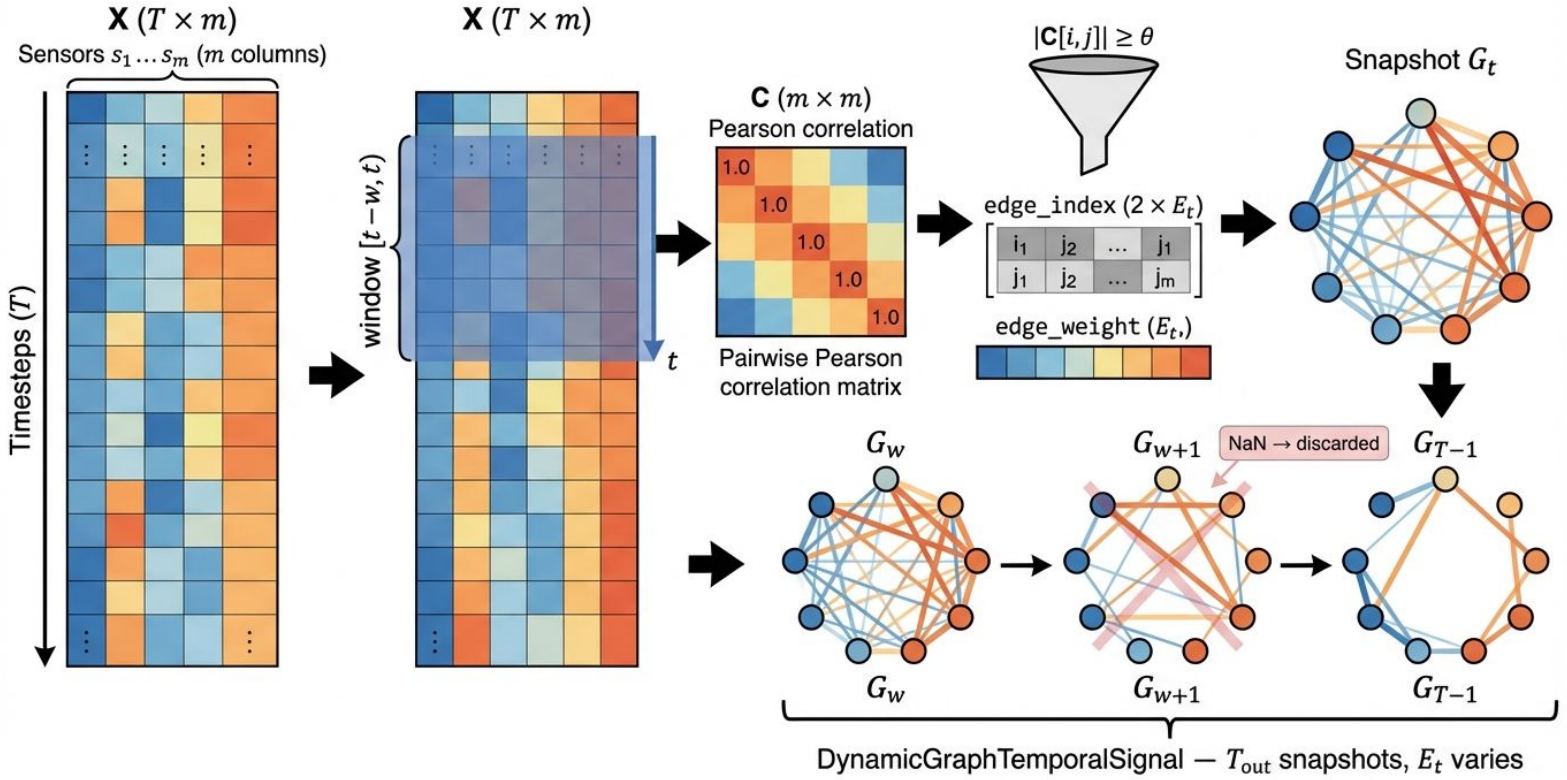

Spatio-Temporal Graph Neural Networks (STGNNs)

Combine graph convolutions to capture spatial structures and recurrent networks to learn temporal dynamics

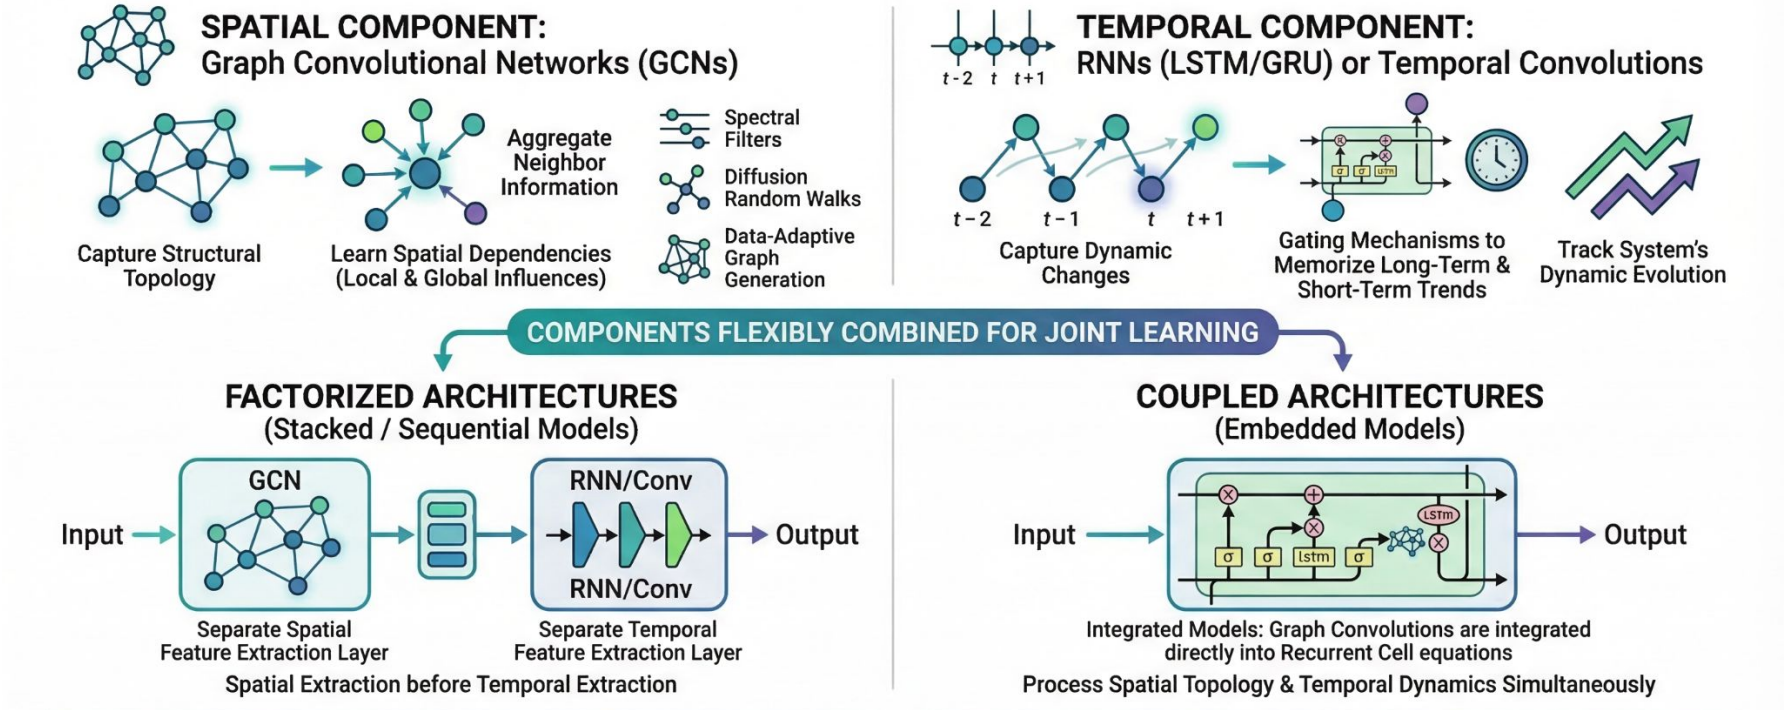

Temporal Graph Convolutional Network (T-GCN) cell [4]

A **decoupled** solution for spatio-temporal modelling. Uses GCN layers only for feature extraction at each GRU gate.

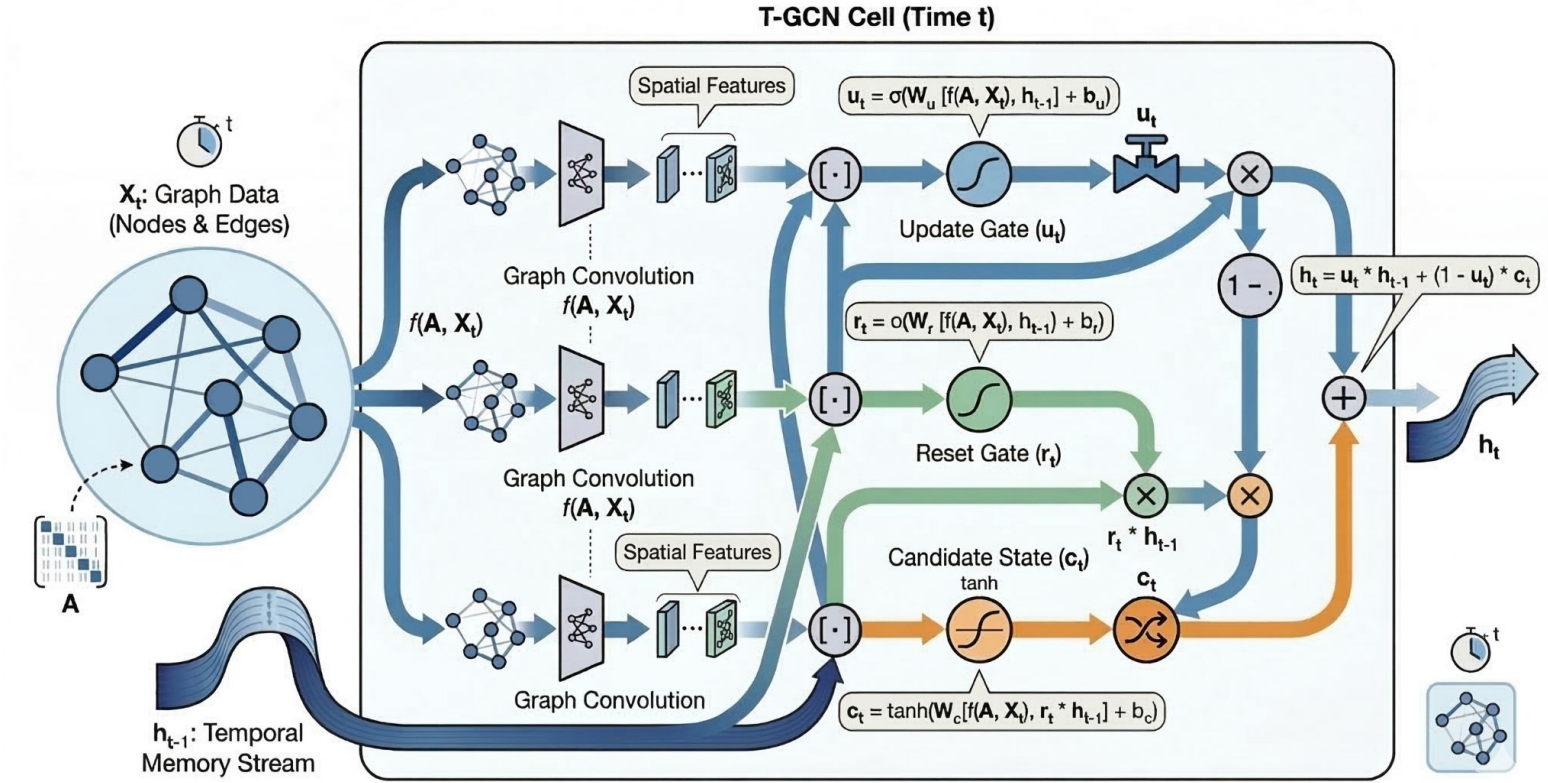

Graph-GRU (GConvGRU) Cell

[5]

Fuses spatial message-passing layers with temporal blocks in an **integrated** GRU cell.

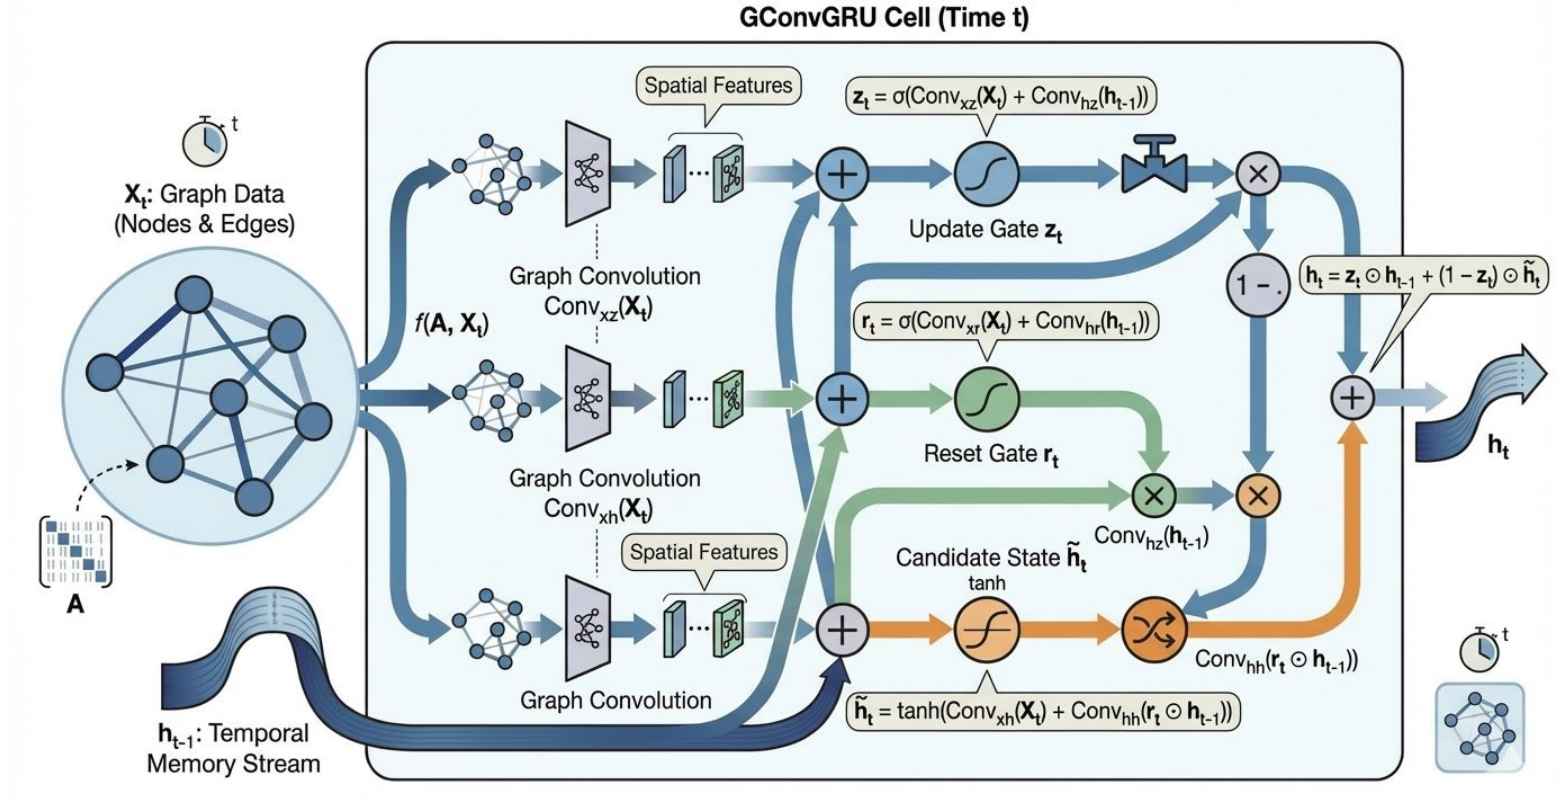

Graph-Level ST-GNN Model

[6]

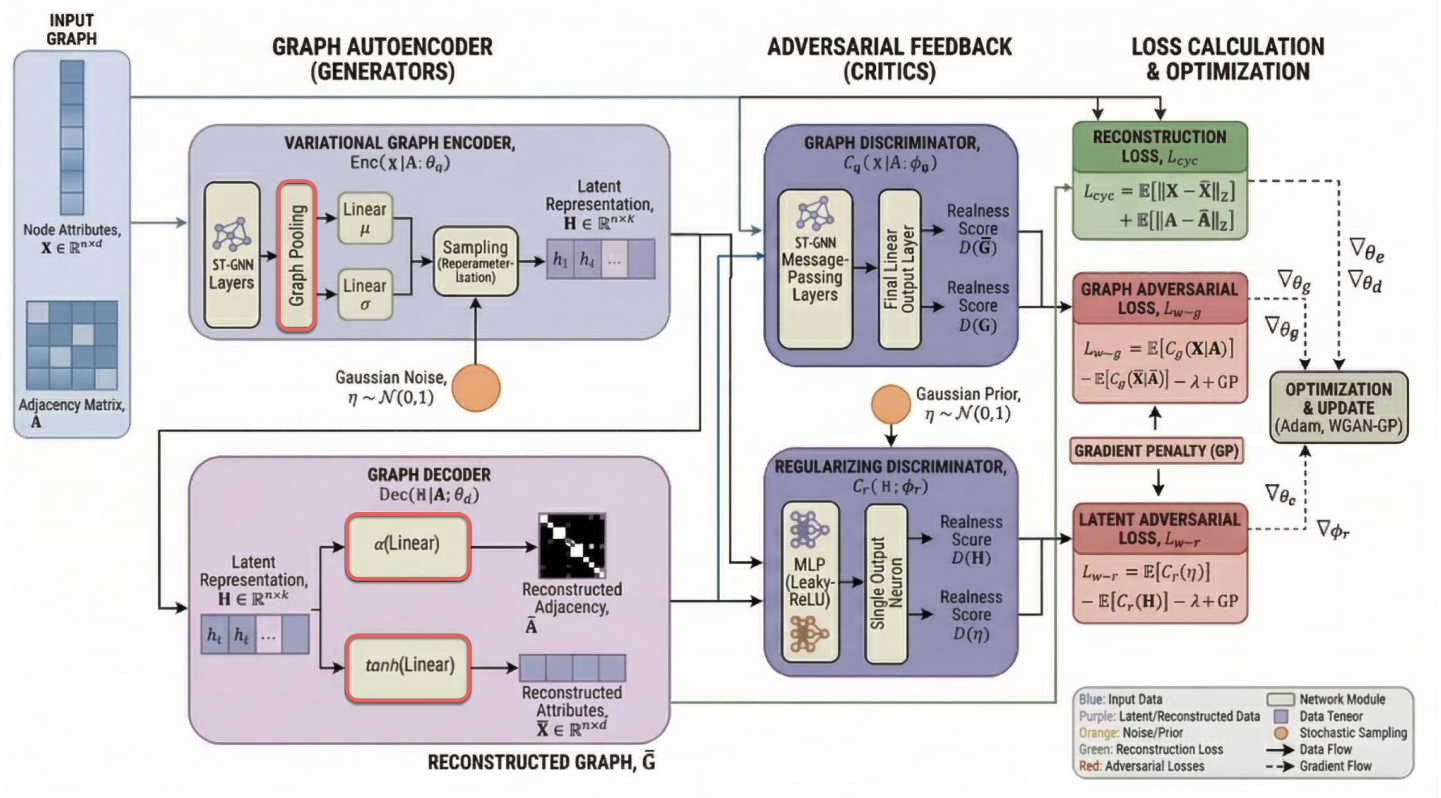

Nodes Reconstruction Errors

Averaged and Scaled for the T-GCN model **(Top)** vs. GConvGRU model **(Bottom).** Threshold at p=0.95.

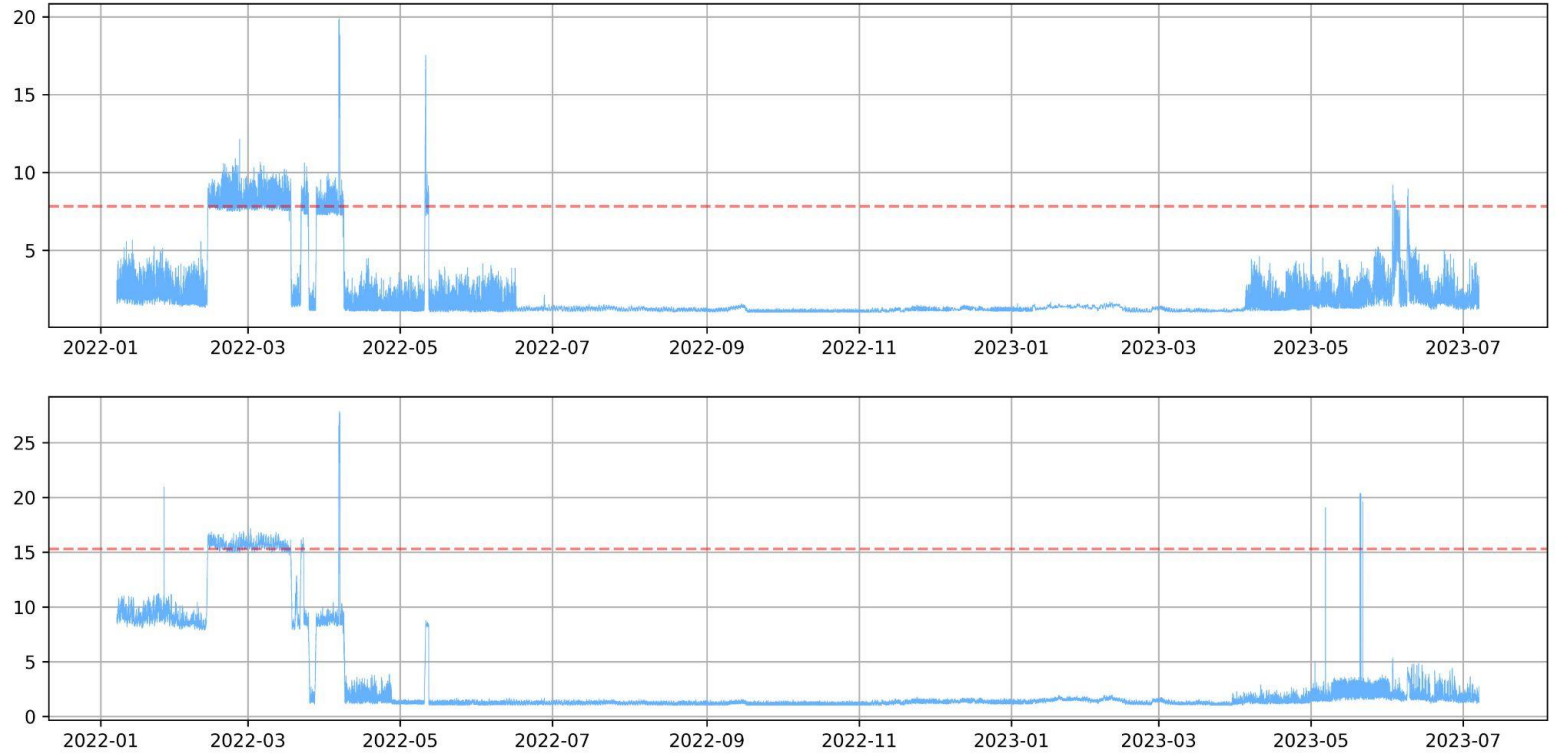

T-GCN Model Predictions

Anomalies in phase 1 current measurements from two active 2,500 kVA transformers (detected from nodes rec. errors).

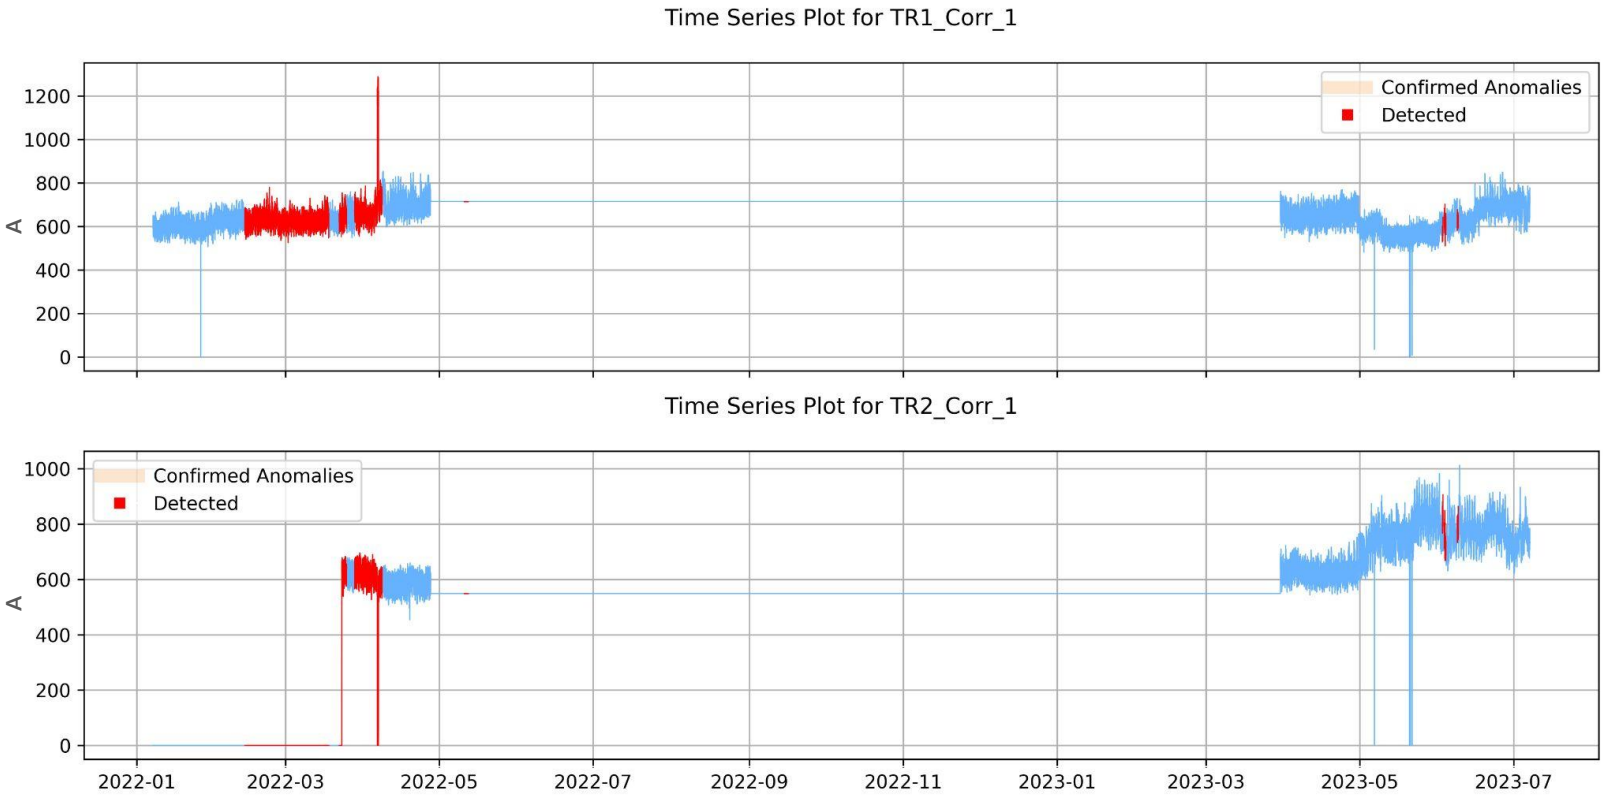

GConvGRU Model Predictions

Anomalies in phase 1 current measurements from two active 2,500 kVA transformers (detected from nodes rec. errors).

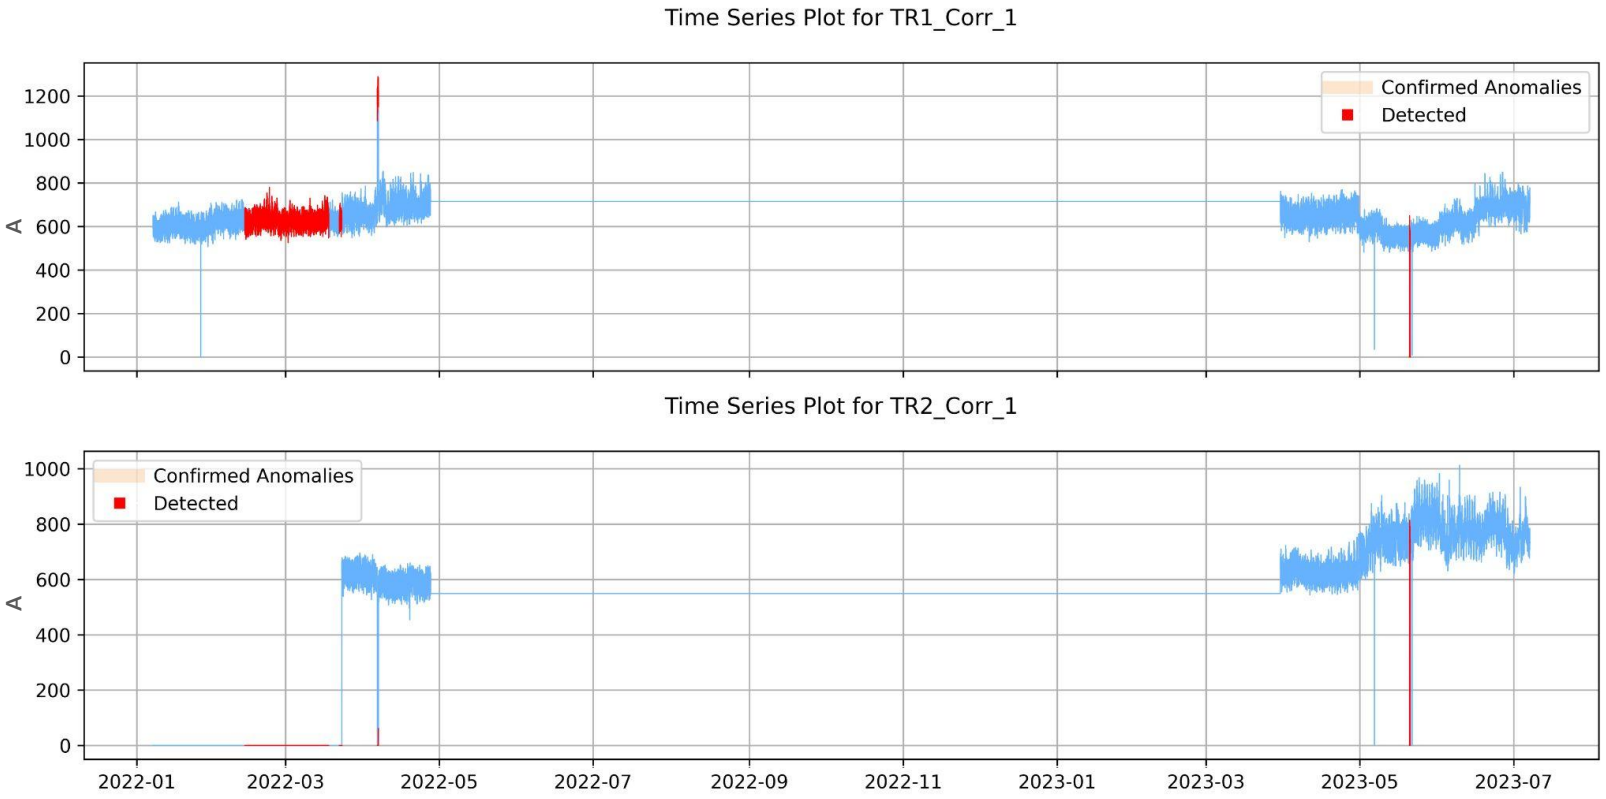

T-GCN Model Predictions

Detected anomalies in the power consumption of the two groups of pumps (detected from nodes rec. errors).

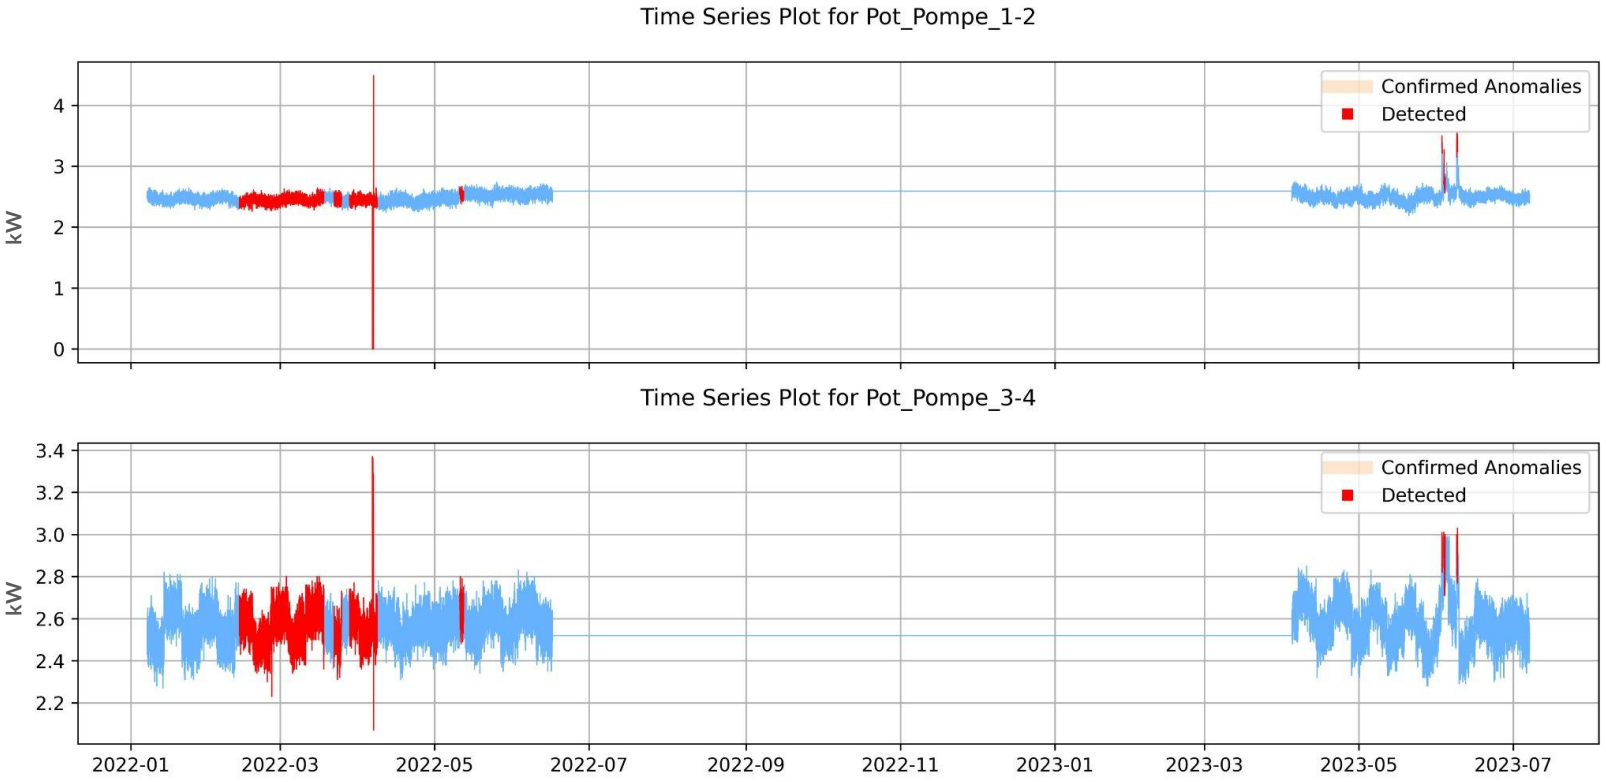

GConvGRU Model Predictions

Detected anomalies in the power consumption of the two groups of pumps (detected from nodes rec. errors).

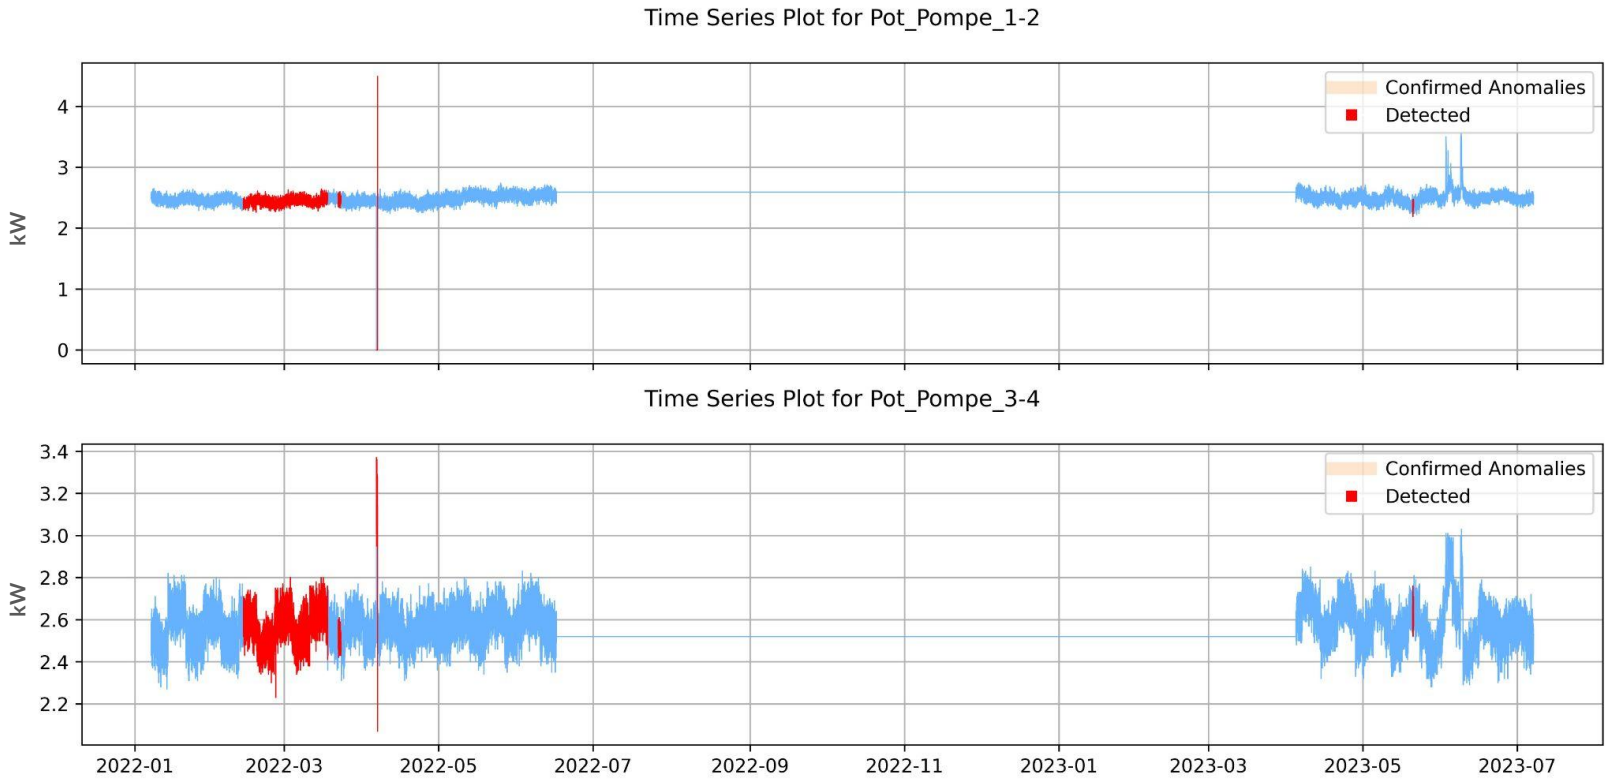

Edges Reconstruction Errors

Averaged and Scaled for the T-GCN model **(Top)**, GConvGRU model **(Middle)** and static graph GNN model **(Bottom).**

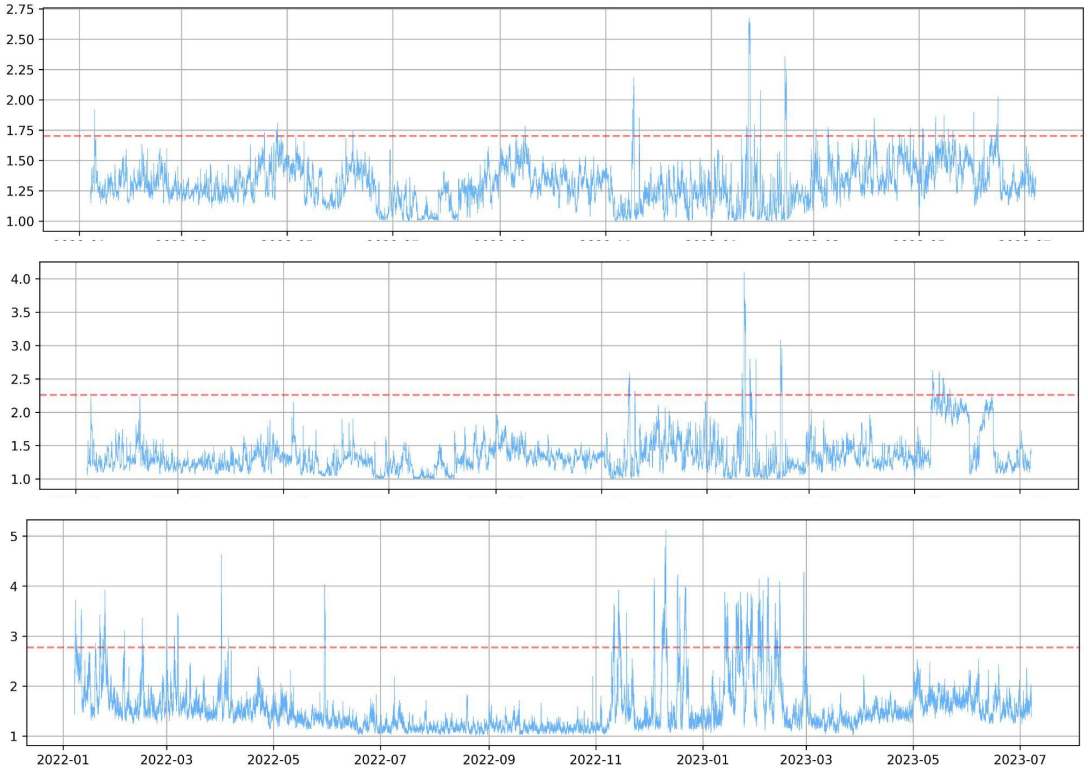

Future Research: Advancing ST-GNNs for AD

- **Dynamic Interpretability:** Extend our visual subgraph inspection tools to handle dynamic timestamp graphs, allowing operators to track how edges appear, disappear, or shift during an anomalous event.
- **Benchmarking**: Rigorously validate the new ST-GNN architectures against publicly available, labeled industrial datasets (such as SMD and PSM).
- **Performance Tuning**: Conduct thorough hyperparameter tuning and ablation studies on the spatio-temporal cells to optimize the precision-recall balance and minimize false positives.
- **Extension to Signed Graphs**: Implement signed GNN layers to explicitly model both positive and negative correlations, improving the system's ability to capture the true polarity of physical inter-dependencies.
- **Live Deployment Tests:** Evaluate the models' computational scalability, memory footprint, and inference speed in a live environment to transition from retrospective analysis to real-time predictive maintenance (PdM).

Thanks!

Any questions or comments?

Speaker: **Giovanni Zurlo** (INFN-CNAF)
giovanni.zurlo@cnaf.infn.it

Elisabetta Ronchieri (INFN-CNAF)
elisabetta.ronchieri@cnaf.infn.it

L. Scarponi,
D. Cesini,
A. Costantini (INFN-CNAF)

References

1. Cesini, D., Cavalli, A., Chierici, A., et al. (2024). **Migrating the INFN-CNAF datacenter to the Bologna Tecnopolo: A status update.** EPJ Web of Conferences, 295, 07030.

2. Yang, Y., & Yang, H. (2008). **Complex network-based time series analysis.** Physica A: Statistical Mechanics and Its Applications, 387(5–6), 1381–1386.

3. Zurlo, G., Ronchieri, E., Cornali, R., Scarponi, L. B., Cesini, D., & Dell'Agnello, L. (2026). **Graph neural network solutions for interpretable anomaly detection in IT infrastructure monitoring time series**. Journal of Business Analytics, 1–20.

4. Zhao, L., Song, Y., Zhang, C., Liu, Y., Wang, P., Lin, T., Deng, M., & Li, H. (2020). **T-GCN: A temporal graph convolutional network for traffic prediction.** IEEE Transactions on Intelligent Transportation Systems, 21(9), 3848–3858.

5. Seo, Y., Defferrard, M., Vandergheynst, P., & Bresson, X. (2016). **Structured sequence modeling with graph convolutional recurrent networks**. arXiv. https://doi.org/10.48550/arXiv.1612.07659

6. Rozemberczki, B., Scherer, P., He, Y., Panagopoulos, G., Riedel, A., Astefanoaei, M., Kiss, O., Beres, F., Lopez, G., Collignon, N., & Sarkar, R. (2021). P**yTorch Geometric Temporal: Spatiotemporal signal processing with neural machine learning models.** In Proceedings of the 30th ACM International Conference on Information and Knowledge Management (pp. 4564–4573).