The Fortieth AAAI Conference on Artificial Intelligence (AAAI-26)

Time-Series Anomaly Detection with Graph-Based Self-Supervised Learning and Foundation Models: Toward Real-World Applications

Thi Kieu Khanh Ho

Department of Electrical and Computer Engineering, McGill University
Mila-AI Quebec Institute, Canada

Abstract

Time-series data, which represent the evolution of one or more variables over time, are ubiquitous across domains such as finance, medicine, industry, and security. Time-Series Anomaly Detection (TSAD) is essential for identifying irregular events such as equipment failures, fraudulent activities, and neurological disorders. Despite significant progress, TSAD remains challenging due to the complexity of time-series signals, the diversity of anomaly types, and the scarcity of high-quality labeled data. This thesis contributes: (i) the first comprehensive surveys of Graph-based TSAD (G-TSAD) and Self-Supervised Learning for Anomaly Detection (SSL-AD), showing how graph modeling and SSL proxy tasks yield robust representations for TSAD while mapping limits and future directions; (ii) EEG-CGS, a contrastive–generative SSL framework that encodes fine-grained subgraph structure without anomaly labels, improving multivariate TSAD and localizing anomalous sensors and regions; (iii) TSAD-C, which integrates graph representations with diffusion models to capture long-range temporal and spatial dependencies while explicitly handling contaminated training data; and (iv) extending TSAD beyond benchmark datasets into other impactful domains, and developing foundation models specialized for biosignals to detect novel anomalies in drug-resistant epilepsy patients.

Related Work

Existing TSAD approaches face several limitations. First, many methods primarily model intra-variable dependencies while overlooking inter-variable relationships, which are critical in multivariate sensory systems. While several recent graph-based approaches have proposed modeling inter-sensor relationships, they rely on static graphs or overlook local structural patterns (Deng and Hooi 2021).

Second, existing TSAD studies can be broadly categorized into supervised, semi-supervised and unsupervised methods. Supervised and semi-supervised methods rely on both labeled normal and anomalous data during training, while unsupervised methods assume that only normal data is available. However, the two former approaches depend on labeled anomalies, which are scarce, and impractical in real-world applications. Unsupervised approaches avoid this limitation by assuming that the training data is anomaly-free, and they rely on reconstruction errors as anomaly scores, which do not guarantee generalizable representations. Self-supervised learning (SSL) has emerged as a promising alternative, leveraging proxy tasks to learn robust representations without labeled anomalies (Chen et al. 2020), though its application to time-series remains underexplored.

Another key challenge in unsupervised TSAD lies in the assumption that training data are purely normal, which is rarely valid in practice due to distribution shifts or labeling errors (Jiang et al. 2022). Such contamination causes unsupervised methods to overfit and misclassify anomalies encountered during training. At the same time, most TSAD models are benchmarked on small, curated datasets that suffer from mislabeled or overly simplistic anomalies (Wu and Keogh 2021), limiting their ability to generalize to real-world, long-duration, and complex settings. For instance, detecting epileptogenic zones (EZ) in electroencephalogram (EEG) requires modeling spatial dependencies across electrodes, long temporal windows, and heterogeneous anomaly types, which are factors not captured by current benchmarks. Additionally, recent time-series foundation models (Jin et al. 2024) demonstrate strong forecasting capabilities, yet they fall short of transferring effectively to other signal domains, such as biosignals, thus remain limited in detecting clinically meaningful anomalies.

Research Questions

Building on the challenges identified above, my thesis addresses the following core research questions:

- **Modeling time-series complexity:** How can I design methods that effectively capture both intra- and inter-variable dependencies to detect anomalies in complex, multivariate time-series?
- **Detecting complex anomaly types:** Beyond point or contextual anomalies in benchmark datasets, can I develop approaches capable of identifying more realistic anomaly categories (e.g., sensor-level or regional anomalies) critical to real-world applications?
- **Reducing dependence on labeled anomalies:** Since anomaly labels are scarce and costly, can I develop methods that operate without labeled anomalies while maintaining high detection accuracy?

- **Robustness to contaminated data:** In practice, training sets often contain mislabeled or hidden anomalies. How can models remain reliable under such contamination?
- **Towards domain-specific foundation models:** Can I develop foundation models specialized for industrial signals or clinical biosignals, such as EEG, to enable impactful applications, e.g., localizing the epileptogenic zone (EZ) in drug-resistant epilepsy patients?

Research Contributions to Date

1. **(Neural Networks, Co-first Author)**: One of the earliest comprehensive surveys on SSL for anomaly detection (Ho†, Hojjati†, and Armanfard 2024). This work reviewed proxy tasks, categorized SSL methods, and analyzed their strengths and weaknesses across domains such as computer vision, text, and signal processing.
2. **(IEEE TPAMI, First Author)**: The first in-depth survey on graph-based approaches for TSAD (Ho, Karami, and Armanfard 2025). This work introduced a unified taxonomy for graph construction and learning methods, analyzed their suitability for modeling intra- and inter-variable dependencies, and highlighted research gaps.
3. **(AAAI 2023, First Author)**: The novel EEG-CGS framework integrates graph construction with SSL to capture local graph patterns and improve robustness in multivariate TSAD (Ho and Armanfard 2023).
4. **(UAI 2025, First Author)**: The novel TSAD-C framework is designed to handle contaminated training datasets, where anomalies are mixed with normal samples (Ho and Armanfard 2025), which addressed a major limitation of existing unsupervised methods.
5. **(Target: Nature Medicine, First Author)**: Currently developing EZ-TNet, the first foundation model for EEG, targeting EZ identification in drug-resistant epilepsy patients. Trained on large-scale clinical datasets, this model aims to capture both local and regional anomalies across long-term recordings. The clinical significance is profound, i.e., accurate EZ identification can directly inform surgical planning to achieve seizure freedom.
6. **(ECAI 2024, Second Author)**: MOSAD, the first framework to address open-set TSAD, where only a small number of anomaly samples are available during training (Lai, Ho, and Armanfard 2024). The framework demonstrated improved adaptability and robustness to unseen anomaly types, compared to existing supervised and unsupervised methods.
7. **(WACV 2025, Second Author)**: GiCiSAD, a novel framework combining graph learning and diffusion models for anomalous human activity recognition (Karami, Ho, and Armanfard 2025). This work extended SSL and graph-based principles into video analysis, showing strong performance on challenging datasets and broadening the scope of anomaly detection beyond signals.

Proposed Research Plan

My doctoral research plan is organized into three main phases leading up to the completion of my dissertation.

- **By September 30, 2025**: (1) Extend my earlier contributions on SSL, graph-based methods, and foundation models by applying them to diverse real-world TSAD domains, including vehicle and aircraft predictive maintenance, industrial system monitoring, and financial anomaly detection. (2) Conduct comparative studies to assess the transferability and limitations of these approaches across different time-series contexts.
- **By January 20–21, 2026**: (1) Finalize the development of a large-scale EEG foundation model in collaboration with neurologists. (2) Focus on detecting complex anomaly categories (sensor-level and regional anomalies), with an emphasis on identifying the EZ in drug-resistant epilepsy patients. (3) Submit this work to Nature Medicine, ensuring that results are rigorously validated both computationally and clinically.
- **By August 31, 2026**: (1) Integrate all completed research into a unified dissertation. (2) Defend the thesis and prepare additional manuscripts for publication in AI, clinical, and industrial venues. (3) Engage with medical and engineering collaborators to evaluate the real-world performance of the proposed models and identify pathways for further improvement and deployment.

References

Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *ICML*.

Deng, A.; and Hooi, B. 2021. Graph neural network-based anomaly detection in multivariate time series. In *AAAI*.

Ho, T. K. K.; and Armanfard, N. 2023. Self-Supervised Learning for Anomalous Channel Detection in EEG Graphs: Application to Seizure Analysis. In *AAAI*.

Ho, T. K. K.; and Armanfard, N. 2025. Contaminated Multivariate Time-Series Anomaly Detection with Spatio-Temporal Graph Conditional Diffusion Models. In *UAI*.

Ho†, T. K. K.; Hojjati†, H.; and Armanfard, N. 2024. Self-supervised anomaly detection in computer vision and beyond: A survey and outlook. *Neural Networks*.

Ho, T. K. K.; Karami, A.; and Armanfard, N. 2025. Graph anomaly detection in time series: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Jiang, X.; Liu, J.; Wang, J.; Nie, Q.; Wu, K.; Liu, Y.; Wang, C.; and Zheng, F. 2022. Softpatch: Unsupervised anomaly detection with noisy data. In *NeurIPS*.

Jin, M.; Wang, S.; Ma, L.; Chu, Z.; Zhang, J. Y.; Shi, X.; Chen, P.-Y.; Liang, Y.; Li, Y.-F.; Pan, S.; et al. 2024. Time-llm: Time series forecasting by reprogramming large language models. In *ICLR*.

Karami, A.; Ho, T. K. K.; and Armanfard, N. 2025. Graph-Jigsaw Conditioned Diffusion Model for Skeleton-based Video Anomaly Detection. In *WACV*.

Lai, T.; Ho, T. K. K.; and Armanfard, N. 2024. Open-Set Multivariate Time-Series Anomaly Detection. In *ECAI*.

Wu, R.; and Keogh, E. J. 2021. Current time series anomaly detection benchmarks are flawed and are creating the illusion of progress. *IEEE TKDE*.