2025 European Conference on Computing in Construction

CIB W78 Conference on IT in Construction

Porto, Portugal
July 14-17, 2025

GRAPH DEVIATION NETWORK FOR FAULT DETECTION AND DIAGNOSIS USING BUILDING AUTOMATION SYSTEM DATA

Farivar Rajabi and J.J. McArthur
Toronto Metropolitan University, Toronto, ON, Canada

Abstract

Automated fault detection and diagnosis (FDD) is critical to maintain high building energy efficiency and good indoor environmental quality. Existing unsupervised methods typically lack the capability to identify the root cause of faults or determine the specific variables responsible for these faults. To address these limitations, this paper presents a Graph Deviation Network (GDN)-based FDD. The extent of faults (i.e. equipment-level, zone-level, or building-level) is determined by applying GDN to Fan Coil Units (FCUs) across zones in a building, and rule-based fault diagnosis identifies fault types. Experiments demonstrate that the proposed method outperforms baseline approaches in anomaly detection.

Introduction

Building Heating, Ventilation, and Air Conditioning (HVAC) systems, frequently encounter faults such as sensor or equipment failures and improper operation, resulting in energy waste, higher maintenance costs, and poor indoor air quality. Studies estimate that these faults can account for 15%–30% of energy waste (Katipamula and Brambley, 2005). The integration of cost-effective data sensing technologies, combined with advancements in smart building systems, cloud computing, and AI, has revolutionized automated fault detection in the building sector. These technologies enable real-time data collection and analysis at a reduced cost, making them highly accessible for HVAC systems in modern buildings. Smart buildings utilize interconnected systems and cloud platforms to store and process vast amounts of data, while AI enhances the ability to detect, diagnose, and predict faults with precision and scalability. Together, these innovations provide a cost-effective, efficient, and intelligent approach to improving system reliability, reducing energy waste, and optimizing building operations (Chen et al., 2023). Utilizing these technologies, Fault detection and diagnostics (FDD), or automated FDD (AFDD), can be implemented to maintain reliable system operation and reduce energy wastage. FDD implementation in office and higher education sectors in the United States has reportedly achieved a 10% median annual energy savings with a two-year payback period, highlighting its competitiveness as a cost-effective investment in the building sector (Kramer et al., 2020).

Despite recent advances in FDD research, there is still a need for an unsupervised methodology capable of handling unlabeled Building Automation System (BAS) data for both fault detection and diagnosis. This study applies and validates an unsupervised self-learning approach using Graph Deviation Network (GDN) for building HVAC FDD. The key objectives are: (a) to design a robust and adaptable algorithm capable of detecting faults across different equipment without relying on labeled training datasets; (b) to determine the extent of faults using GDN and identify fault types through a rule-based diagnosis step; and (c) to evaluate the algorithm's performance against various fault scenarios. Validation is conducted using real-world case study data.

Literature review

Extensive research has been conducted in the field of data-driven FDD, with several reviews summarizing the state-of-the-art advancements, including notable works (Bi et al., 2024; Chen et al., 2023; Mirnaghi and Haghighat, 2020; Zhang et al., 2023). Recent supervised approaches for HVAC fault detection have utilized deep learning and ensemble strategies to improve classification accuracy. Matetić et al. (2022) applied CNN, LSTM, and GRU models for detecting common FCU faults in hotel buildings, demonstrating high accuracy using real and simulated data. You et al. (2024) proposed an ensemble learning-based diagnosis system for air conditioning units, leveraging GANs to augment limited fault datasets and employing multi-label classification to identify concurrent faults. Gao et al. (2024) developed a building-level diagnostic method using Bayesian networks derived from system topology and metadata. While these methods show promising results, their performance heavily depends on the availability of labeled fault data, which is typically scarce, difficult to generalize, and expensive to obtain in real-world building environments.

Several recent studies have explored unsupervised learning approaches for fault detection in HVAC systems, aiming to overcome the reliance on labeled data. Rajabi and McArthur (2024) applied OPTICS clustering with

and without PCA to detect FCU anomalies using BAS data, demonstrating its effectiveness in large-scale building operations. Parsaei et al. (2024) used Autoencoders (AEs) and PCA to identify zone-level anomalies in VAV terminal units, offering interpretable results from multivariate BAS data. El Mokhtari and McArthur (2024) proposed an AE-based method capable of generalizing across FCUs. Tra et al. (2024) introduced an adaptive adversarial AE for AHU fault detection, achieving high detection accuracy with compact latent spaces. Dey et al. (2020), used clustering-based strategies for unsupervised FCU anomaly detection, validating outcomes with expert knowledge. Despite these advancements, a major challenge of unsupervised methods remains that fault diagnosis is not directly achievable and requires additional post-processing steps—such as rule-based classification or expert analysis—to interpret and classify the detected anomalies.

Graph-based methods for HVAC fault detection remain limited in the literature. Fan et al. (2023) proposed a semi-supervised fault diagnosis method using graph convolutional networks (GCNs) to improve HVAC fault classification performance in data-scarce environments. By constructing graphs from operational data using k-nearest neighbors, the model utilizes both labeled and unlabeled samples to boost classification accuracy. However, the method still requires a portion of labeled data for training. Gu et al. (2023) introduced a graph embedding-based approach that constructs dynamic graphs using Pearson correlation coefficients and performs anomaly detection through a two-stage process—first identifying abnormal time intervals and then localizing faulty sensors. This method relies on rule-based graph construction and requires separate models for different tasks. GDN (Deng and Hooi, 2021) is a graph neural network-based approach developed for anomaly detection in multivariate time series data. It addresses limitations of traditional methods, such as the inability to model complex, nonlinear relationships between variables. GDN learns a graph structure that captures the dependencies between sensors and detects deviations from these learned patterns, making it particularly suitable for systems with highly interdependent sensor data, like HVAC systems or industrial monitoring. Despite its potential, GDN has been rarely applied in the context of HVAC fault detection and diagnosis. Sun et al. (2024) proposed an improved version of GDN for chiller fault classification by incorporating sparse cointegration analysis and attention mechanisms; however, the method is supervised and depends on labeled datasets. However, no prior study has explored the application of GDN in a fully unsupervised framework for both fault detection and diagnosis across different system levels in building HVAC systems. The present study applies GDN in a fully unsupervised manner with learned graph structures, offering a unified framework for fault detection. In addition, it introduces a novel capability to diagnose the root cause of faults by identifying the extent of the fault—whether it originates from individual equipment, a thermal zone, or the broader building system—without requiring labeled data.

Methodology

GDN operates through four main components: (1) Sensor Embedding, which uses embedding vectors to represent the unique behavior of each sensor; (2) Graph Structure Learning, which constructs a directed graph to model dependency relationships between sensors; (3) Graph Attention-Based Forecasting, which predicts each sensor's future behavior based on an attention mechanism that assigns varying weights to neighboring nodes; and (4) Graph Deviation Scoring, which identifies and explains anomalies by comparing observed data with predicted behavior. These components collectively enable GDN to not only detect anomalies but also provide insight into the root causes of the deviations, enhancing its interpretability.

The training dataset consists of multivariate time-series data from N sensors over T_{train} time steps, denoted as $s_{train} = [s_{train}^{(1)}, \dots, s_{train}^{(T_{train})}]$. The goal is to identify anomalies in test data from the same N sensors across T_{test} time steps, denoted as $s_{test} = [s_{test}^{(1)}, \dots, s_{test}^{(T_{test})}]$. The algorithm outputs a binary label $a(t) \in \{0,1\}$ for each time step, where $a(t) = 1$ indicates an anomaly at time t. To capture the relationships between sensor measurements, each sensor is represented by an embedding vector $v_i \in R^d$, for $i \in \{1,2, \dots, N\}$, initialized randomly and trained with the model.

To model sensor relationships, a directed graph is constructed with nodes representing sensors and edges denoting dependencies. The graph is represented by an adjacency matrix A, where A_{ij} indicates a directed edge from node i to j. Candidate relations for each sensor i are defined as:

$$C_i = \{1,2, \dots, N\} \backslash \{i\} \quad (1)$$

which includes all other sensors except itself. The dependency between sensor i and its candidates $j \in C_i$ is computed using the normalized dot product of their embedding vectors:

$$e_{ji} = \frac{v_i v_j}{\|v_i\| \|v_j\|} \quad (2)$$

To ensure sparsity, the top-k dependencies are selected based on e_{ij}, where k is a user-defined parameter controlling sparsity. The adjacency matrix is updated as:

$$A_{ij} = 1 \text{ if } j \in TopK(\{e_{ki} : k \in C_i\}), \quad (3)$$

where $TopK$ selects the indices of the top-k values. This learned adjacency matrix A is then utilized in the graph attention-based model for anomaly detection.

GDN uses a forecasting-based approach to explain anomalies by predicting each sensor's behavior based on historical data. At time t, the input to the model is defined as a sliding window of size w over past time-series data:

$$x^{(t)} = \left[s^{(t-w)}, s^{(t-w+1)}, \dots, s^{(t-1)}\right], \quad (4)$$

where the model predicts $s^{(t)}$, the sensor data at time t.

To model sensor relationships, a graph attention-based feature extractor aggregates information from neighboring nodes. The aggregated representation of node i at time t is:

$$z_i^{(t)} = ReLU\left(\alpha_{i,i} W x_i^{(t)} + \sum_{j \in N(i)} \alpha_{i,j} W x_j^{(t)}\right), \quad (5)$$

where $N(i) = \{j \mid A_{ji} > 0\}$ denotes the neighbors of node i from the adjacency matrix A, W is a learnable weight matrix, and $\alpha_{i,j}$ are attention coefficients computed as:

$$g_i^{(t)} = v_i \oplus W x_i^{(t)},$$

$$\pi_{(i,j)} = LeakyReLU\left(a^{\top}\left(g_i^{(t)} \oplus g_j^{(t)}\right)\right), \quad (6)$$

$$\alpha_{i,j} = \frac{exp\left(\pi(i,j)\right)}{\sum_{k \in N\{i\} \cup \{i\}} exp\left(\pi(i,k)\right)}$$

The extracted representations $\{z_i^{(t)}\}$ for all N nodes are combined with sensor embeddings v_i through element-wise multiplication ($\circ$) and passed to fully connected layers:

$$\hat{s}^{(t)} = f_\theta\left(\left[v_1 \circ z_1^{(t)}, \ldots, v_N \circ z_N^{(t)}\right]\right) \quad (7)$$

where $\hat{s}^{(t)}$ is the predicted sensor data at time t.

The Mean Squared Error (MSE) between the predicted $\hat{s}^{(t)}$ and observed $s^{(t)}$ is used for training:

$$L_{MSE} = \frac{1}{T_{train} - w} \sum_{t=w+1}^{T_{train}} \left\| \hat{s}^{(t)} - s^{(t)} \right\|_2^2 \quad (8)$$

The anomalousness score for each sensor at time t is computed by comparing the expected and observed behavior:

$$Err_i(t) = \left| s_i^{(t)} - \hat{s}_i^{(t)} \right| \quad (9)$$

where $s_i^{(t)}$ is the observed value and $\hat{s}_i^{(t)}$ is the predicted value. To account for different scales across sensors, a robust normalization is applied using the median $\tilde{\mu}_i$ and interquartile range $\tilde{\sigma}_i$:

$$a_i(t) = \frac{Err_i(t) - \tilde{\mu}_i}{\tilde{\sigma}_i} \quad (10)$$

The overall anomalousness at time t is calculated by aggregating the sensor scores using the max function:

$$A(t) = max\, a_i(t) \quad (11)$$

To smooth abrupt changes in error values, a Simple Moving Average (SMA) is applied to produce smoothed scores $A_s(t)$. A time tick t is labeled as an anomaly if $A_s(t)$ exceeds a threshold, which is set as the maximum value of $A_s(t)$ over the validation data.

To evaluate the proposed method, Fan Coil Units (FCU) data collected from the BAS was utilized, with the equipment located in the Daphne Cockwell Complex at Toronto Metropolitan University (TMU), Toronto, Canada. FCUs were selected for their widespread presence, providing a large dataset, and their diverse configurations, such as heating-only, cooling-only, combined heating and cooling, and cooling with in-slab heating. As a commonly used terminal unit, FCUs incorporate key HVAC components like fans, heating and cooling coils, and control valves, making them representative of broader HVAC systems. This setup allowed for evaluating the scalability and adaptability of the proposed method, with the potential for widespread application across other buildings.

Figure 1 shows a schematic of an FCU, consisting of a heating coil, cooling coil, and fan working to maintain zone temperature (ZN-T) at the setpoint.

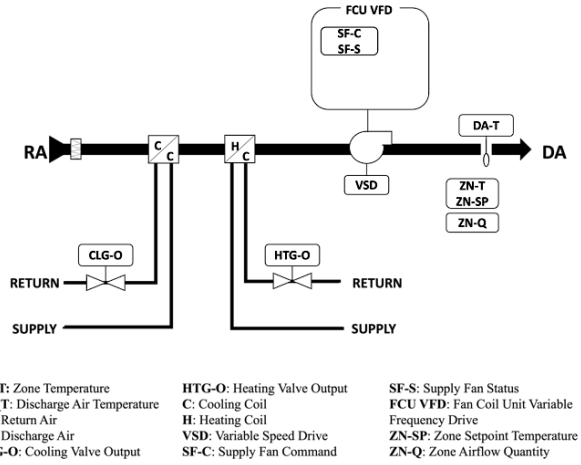

Figure 1: Schematic Diagram of a Fan Coil Unit (FCU)

Data was collected over a two-year period, from March 15, 2022, to March 15, 2024. The dataset includes room temperature (T), discharge air temperature (DA-T), command points like cooling valve output (CLG-O), heating valve output (HTG-O), and outdoor air temperature (OA-T). To detect the extent of the fault, GDN is applied to the data from four identical FCUs within the building, which are distributed across two distinct thermal zones. Table 1 and Figure 2 presents the FCU case studies, their corresponding thermal zones, and the spaces they serve.

Table 1: FCU case studies

FCU ID	Thermal zone	Level	Space served
FCU-1-12	2	1	Classroom
FCU-1-13	2	1	Classroom
FCU-3-1	3	3	Room
FCU-3-2	3	3	Gallery

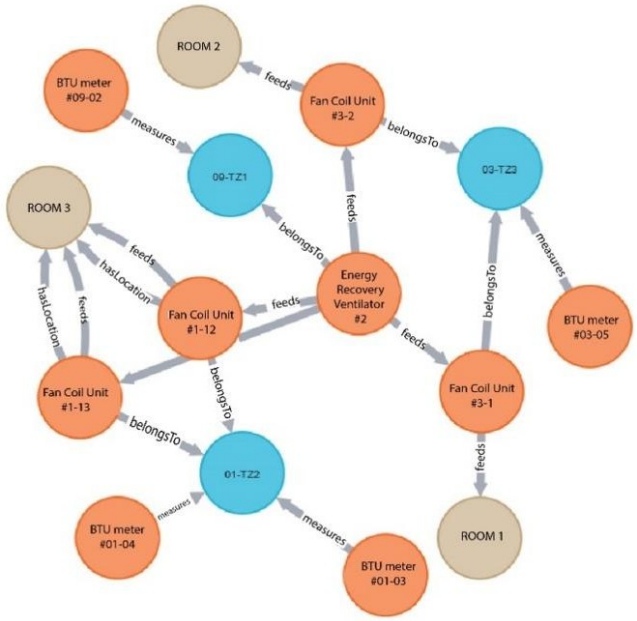

Figure 2: FCU configuration

GDN was implemented following (Deng and Hooi, 2021). The model was trained on the first half of a two-year dataset with a 10-minute interval and tested on the second half. The sliding window size was set to 5-time steps. The embedding size was chosen as 64. The graph sparsity was controlled by limiting the maximum number of neighbors for each sensor to 3. A learning rate of 0.001 was selected. A batch size of 128 was used. The model was trained for a maximum of 50 epochs, with early stopping applied after 10 epochs of no improvement in validation loss to avoid overfitting. The attention dimension was set to 64. The anomaly detection threshold was determined as the maximum value of the smoothed anomaly score observed during the validation phase. This approach ensures that any deviation exceeding the highest score seen under normal conditions is flagged as an anomaly. In this study, the final threshold value was 3.47 for all FCUs, corresponding to the peak smoothed score during validation. The performance of the GDN is compared against AEs which is a commonly used anomaly detection technique. Evaluation is conducted using standard metrics, including Precision, Recall, F1-Score, and Accuracy, calculated based on the test dataset and its associated ground truth labels.

Root Cause Diagnosis

Since the fault types could be generated by a variety of different equipment failures, a means to detect the extent of the fault – namely, whether it is occurring at the equipment, zone, or central plant level – is critical. This is achieved by applying the GDN to data from multiple FCUs across different thermal zones. GDN identifies anomalies in time-series data by capturing the relationships between sensors and pinpointing the responsible variables. The algorithm determines whether the fault is isolated to a single unit, impacts an entire zone, or originates from the central system. Once the extent of the fault is identified, a rule-based fault diagnosis step uses predefined rules to classify the specific fault type. Table 2 outlines the rules for diagnosing faults at the equipment, zone, and building levels.

Table 2: Rule-based fault diagnosis considering the extent of fault

Detected Fault	Fault Extent	Fault Type
Under-Cooling	Equipment Level	Cooling coil leak or cooling valve stuck closed.
	Zone Level	Low flow of chilled water.
	Building Level	Cooling plant malfunction (e.g., cooling tower, chiller, or pumps).
Under-Heating	Equipment Level	Heating coil leak or heating valve stuck closed.
	Zone Level	Low flow of hot water.
	Building Level	Central heating plant malfunction (e.g., boiler burner failure, pump failure, or loss of fuel supply).
Over-Cooling	Equipment Level	Cooling valve stuck open, causing very cold air circulation.
	Zone Level	n/a
	Building Level	n/a
Over-Heating	Equipment Level	Heating valve stuck open.
	Zone Level	n/a
	Building Level	n/a

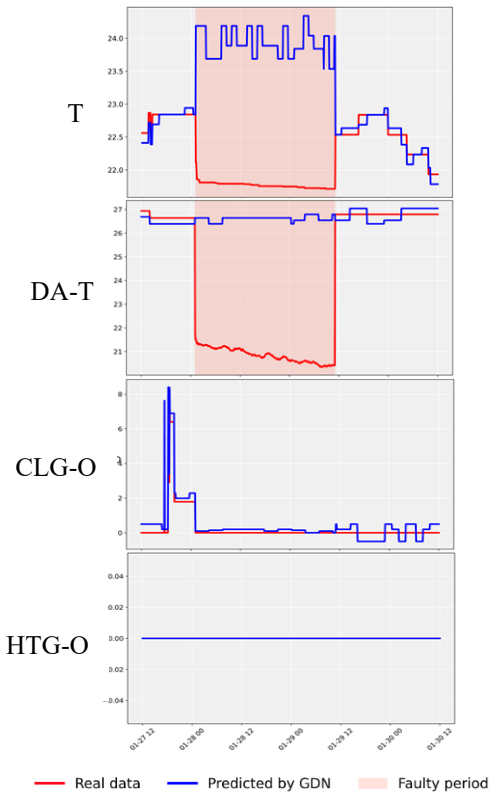

Figure 3: Over-cooling fault detected in FCU-3-2

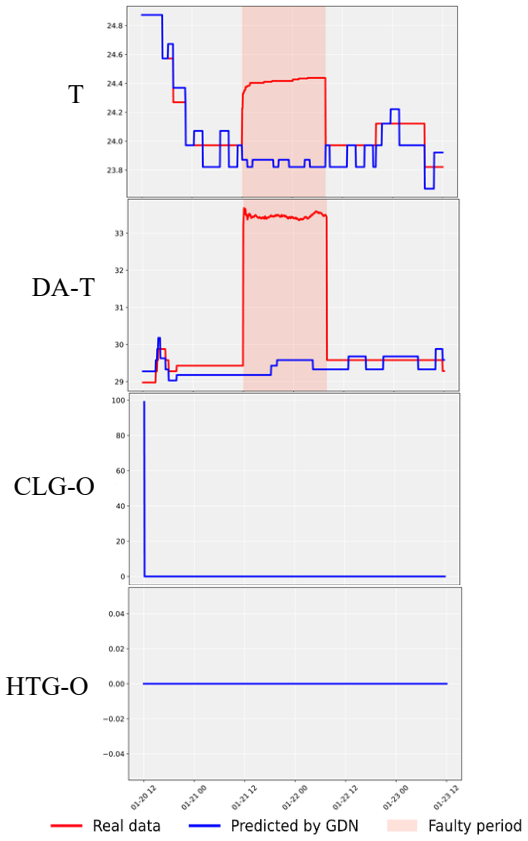

Figure 4: Over-heating fault detected in FCU-3-1

Results

The results of this study highlight the effectiveness of the GDN in detecting faults within four identical FCUs. The analysis focuses on the model's performance in identifying anomalies and pinpointing responsible variables. Following sections are organized based on the extent of the detected faults

Equipment-level faults

Figure 3 illustrates the output of GDN, highlighting one of the faulty periods detected. The red graph represents the actual sensor data, while the blue graph shows the variable predicted by GDN. The faulty period is marked in red for sensors with the highest deviation scores. As shown, the detected fault is over-cooling, occurring from 2024-01-28 01:05:00 to 2024-01-29 11:00:00. During this period, when CLG-O is 0, both DA-T and T are decreasing. This fault was identified only in FCU-3-2, while the other FCUs were determined to be operating normally. Consequently, the detected extent is at the equipment level. Based on the rules outlined in Table 2, the diagnosed fault type is a cooling valve stuck open, resulting in very cold air circulation within the space.

Similarly, as shown in Figure 4, an over-heating fault was detected in FCU-3-1 from 2024-01-21 12:00:00 to 2024-01-22 7:55:00. During this period, HTG-O is 0, while both DA-T and T were increasing, whereas GDN predicted lower values for both variables. The fault is classified as equipment-level since it was detected in only one unit. According to the diagnosis, the fault type is a heating valve stuck open.

Zone-level faults

Figure 5 presents the results of GDN for the detected fault occurring from 2024-02-03 20:00:00 to 2024-02-04 15:50:00 in FCU-1-12 and FCU-1-13. According to the results, GDN predicted higher values for T and DA-T compared to the actual data, while HTG-O remained at its maximum level of 100, indicating the presence of an under-heating fault. Since this fault was detected in two identical FCUs within the same thermal zone (zone 2: FCU-1-12 and FCU-1-13), while the other two FCUs (zone 3: FCU-3-1 and FCU-3-2) operated normally, the extent of fault is classified as a zone-level. Based on the diagnosis, the fault type is identified as low flow of hot water in zone 2.

Building-level faults

As shown in Figure 6, an under-cooling fault was detected simultaneously across all four FCUs, spanning both thermal zones 2 and 3. From 2023-05-09 13:25:00 to 2023-05-10 09:15:00, GDN predicted lower values for T and DA-T in FCU-1-12, FCU-1-13, FCU-3-1, and FCU-3-2, indicating that the extent of the fault surpasses the equipment and zone levels. During this period, CLG-O remained at 100, yet the circulation temperature increased, confirming an under-cooling fault at the

building level. Based on the diagnosis rules, the fault is identified as a malfunction of the cooling plant (e.g., cooling tower, chiller, or pumps).

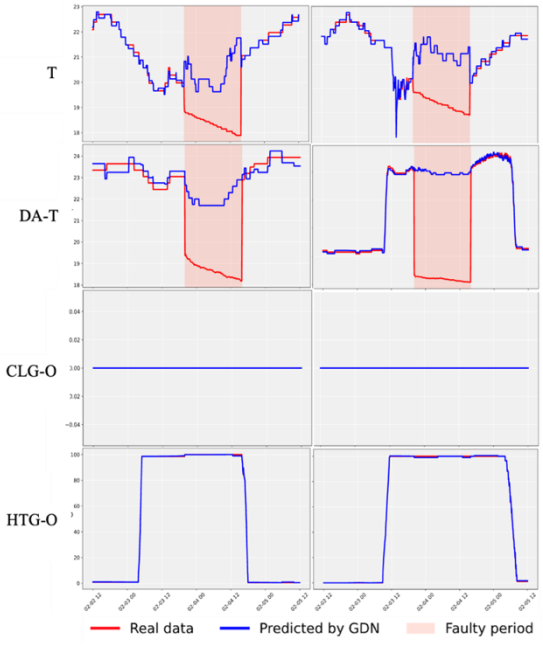

Figure 5: Under-heating fault detected in FCU-1-12 (left) and FCU-1-13 (right)

Table 3 provides a detailed comparison of the results for GDN and AEs across various fault types and FCUs. The results demonstrate that GDN outperforms AE in nearly all cases.

Discussion

The results study show that the GDN successfully detected various types of HVAC faults across different FCUs. Since fault diagnosis in this study is based on predefined crisp rules applied after anomaly detection, the classification performance was entirely dependent on the detection accuracy of the core algorithms—GDN and AE—allowing for a fair comparison focused on precision, recall, F1-score, and accuracy.

As shown in Table 3, GDN outperformed AE in nearly all cases, with F1-scores ranging from 0.62 to 0.88, compared to AE's range of 0.46 to 0.82. This improved performance can be attributed to GDN's ability to model inter-sensor dependencies using graph structures and to focus on the most influential variables via its attention mechanism, whereas AE treats each time series independently and lacks this relational awareness.

However, both algorithms exhibited lower detection performance for the under-cooling fault across all FCUs. This degradation is attributed to the presence of long-lasting under-cooling faults in the first half of the dataset (15 March 2022 to 15 March 2023), which was used for training. Since both GDN and AE operate under the assumption that the training data represents normal system behavior, they learned these faulty patterns as part of the normal distribution. Consequently, when the same under-cooling patterns occurred during testing, the models were less likely to detect them as anomalies. However, GDN experienced less performance degradation compared to AE. This is likely because GDN, unlike AE, uses temporal forecasting combined with learned graph-based relationships. Even if some faulty behavior is included in the training data, GDN can still detect the fault during testing if the related sensors show larger differences than what was seen before. Because GDN looks at how sensors are connected and makes predictions based on their relationships, it can notice unusual patterns that AE might miss, since AE only focuses on how well it can rebuild each sensor's data separately. These findings reinforce the robustness of GDN in real-world applications, where clean training data cannot always be guaranteed.

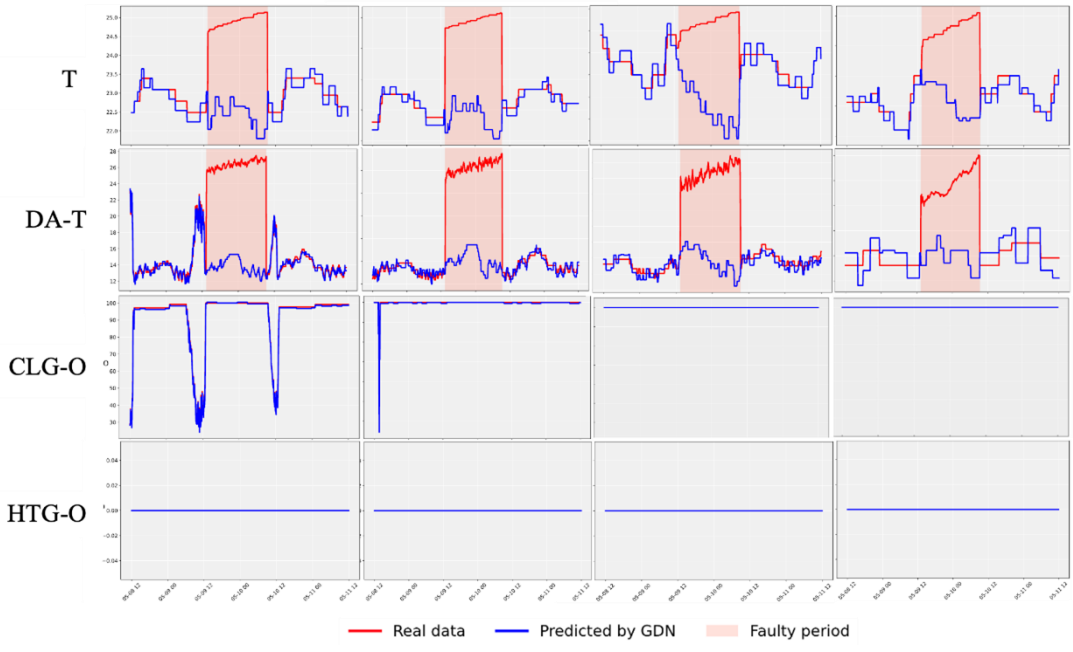

Figure 6: Under-cooling faults detected in (left to right) FCU-1-12, FCU-1-13, FCU-3-1, FCU-3-2

Table 3: Evaluation metrics for GDN and AE.

Fault Type	FCU	Method	Precision	Recall	F1-Score	Accuracy
Under-Cooling	FCU-1-12	AE	0.67	0.55	0.60	0.69
		GDN	**0.73**	**0.68**	**0.70**	**0.76**
	FCU-1-13	AE	0.48	0.45	0.46	0.57
		GDN	**0.65**	**0.60**	**0.62**	**0.70**
	FCU-3-1	AE	0.63	0.55	0.52	0.58
		GDN	**0.76**	**0.68**	**0.72**	**0.79**
	FCU-3-2	AE	0.64	0.58	0.61	0.69
		GDN	**0.78**	**0.74**	**0.76**	**0.82**
Under-Heating	FCU-1-12	AE	0.82	0.81	0.82	**0.9**
		GDN	**0.84**	**0.88**	**0.86**	0.89
	FCU-1-13	AE	0.81	0.79	0.8	0.78
		GDN	**0.88**	**0.87**	**0.88**	**0.84**
	FCU-3-1	AE	0.81	0.8	0.8	0.76
		GDN	**0.86**	**0.88**	**0.87**	**0.85**
	FCU-3-2	AE	0.79	0.84	0.81	0.83
		GDN	**0.86**	**0.85**	**0.86**	**0.87**
Over-Cooling	FCU-1-12	AE	0.79	0.8	0.8	**0.88**
		GDN	**0.85**	**0.88**	**0.86**	0.86
	FCU-1-13	AE	0.81	0.82	0.82	**0.94**
		GDN	**0.86**	**0.86**	**0.86**	0.91
	FCU-3-1	AE	0.77	0.79	0.78	0.74
		GDN	**0.87**	**0.88**	**0.88**	**0.94**
	FCU-3-2	AE	0.8	0.81	0.8	0.75
		GDN	**0.87**	**0.85**	**0.86**	**0.91**
Over-Heating	FCU-1-12	AE	0.81	0.8	0.8	**0.92**
		GDN	**0.87**	**0.85**	**0.86**	0.85
	FCU-1-13	AE	0.81	0.82	0.81	0.8
		GDN	**0.89**	**0.87**	**0.88**	**0.91**
	FCU-3-1	AE	0.78	0.79	0.78	0.73
		GDN	**0.87**	**0.85**	**0.86**	**0.85**
	FCU-3-2	AE	0.79	0.8	0.79	0.91
		GDN	**0.87**	**0.9**	**0.88**	**0.92**

Conclusion

The proposed GDN approach demonstrated reliable performance for FDD, surpassing AE in the majority of cases. GDN identified faulty periods and consistently highlighted the sensor with the highest deviation score as the responsible indicator, offering a precise and reliable approach to fault detection. In addition, the ability to simultaneously monitor multiple systems using GDN allows fault extent detection, which – combined with simple rules – can indicate the nature of the fault and support root cause analysis. This ability to not only detect but also diagnose faults using an unsupervised method makes GDN extremely valuable to support the field implementation of FDD.

A key limitation of this study includes the inability of the rule-based fault diagnosis approach to identify all possible fault types, especially complex or overlapping faults. A second limitation is that the method assumes that the training data represents normal operation, which may lead to false positives if minor faults are present in the training data. Future studies could explore training GDN on a variety of similar systems within a building, such as a wider range of FCUs, to incorporate diverse operational data into the training dataset and enhance the model's robustness and fault detection accuracy through transfer learning. Additionally, injecting noise into the training data – similar to the method used by Kandil and McArthur (2024) could improve fault tolerance and reduce the

likelihood of false positives when minimal faults are present in the training dataset. Finally, recent studies, including El Mokhtari and McArthur (2024), have shown that AEs perform well across different FCUs and fault types, making them a practical and commonly used baseline in HVAC FDD research. For this reason, AE was used as the benchmark in this study. However, stronger graph-based baseline methods—such as Graph Autoencoders (GAEs)—have not yet been explored for HVAC systems. Future work could focus on developing and testing these models to better compare and understand the performance of GDN.

Acknowledgments

This research was funded by the Natural Science and Engineering Research Council of Canada (RGPIN-2018-04105).

References

Bi, J., Wang, H., Yan, E., Wang, C., Yan, K., Jiang, L. & Yang, B. (2024) AI in HVAC fault detection and diagnosis: A systematic review. Energy Reviews, 100071, p.100071.

Chen, Z., O'Neill, Z., Wen, J., Pradhan, O., Yang, T., Lu, X., Lin, G., Miyata, S., Lee, S. & Shen, C. (2023) A review of data-driven fault detection and diagnostics for building HVAC systems. Applied Energy, 339, p.121030.

Deng, A. & Hooi, B. (2021) Graph neural network-based anomaly detection in multivariate time series. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp.4027–4035.

Dey, M., Rana, S.P. & Dudley, S. (2020) A case study based approach for remote fault detection using multi-level machine learning in a smart building. Smart Cities, 3, pp.401–419.

El Mokhtari, K. & McArthur, J.J. (2024) Autoencoder-based fault detection using building automation system data. Advanced Engineering Informatics, 62, p.102810.

Fan, C., Lin, Y., Piscitelli, M.S., Chiosa, R., Wang, H., Capozzoli, A. & Ma, Y. (2023) Leveraging graph convolutional networks for semi-supervised fault diagnosis of HVAC systems in data-scarce contexts. Building Simulation, 16, pp.1499–1517.

Gao, T., Marié, S., Béguery, P., Thebault, S. & Lecoeuche, S. (2024) Integrated building fault detection and diagnosis using data modeling and Bayesian networks. Energy and Buildings, 306, p.113889.

Gu, Y., Li, G., Gu, J. & Jung, J.J. (2023) Graph embedding-based anomaly localization for HVAC system. Journal of Building Engineering, 77, p.107511.

Kandil, M.S. & McArthur, J.J. (2024) The benefit of noise-injection for dynamic gray-box model creation. Advanced Engineering Informatics, 60, p.102381.

Katipamula, S. & Brambley, M.R. (2005) Methods for fault detection, diagnostics, and prognostics for building systems—A review, Part I. HVAC&R Research, 11(1), pp.3–25.

Kramer, H., Lin, G., Curtin, C., Crowe, E. & Granderson, J. (2020) Building analytics and monitoring-based commissioning: industry practice, costs, and savings. Energy Efficiency, 13, pp.537–549.

Matetić, I., Štajduhar, I., Wolf, I. & Ljubić, S. (2022) A review of data-driven approaches and techniques for fault detection and diagnosis in HVAC systems. Sensors, 23(1), p.1.

Mirnaghi, M.S. & Haghighat, F. (2020) Fault detection and diagnosis of large-scale HVAC systems in buildings using data-driven methods: A comprehensive review. Energy and Buildings, 229, p.110492.

Parsaei, A., Gunay, B., O'Brien, W. & Moromisato, R. (2024) Unsupervised identification of zone-level anomalies in VAV terminal units utilizing autoencoders and PCA. Science and Technology for the Built Environment, 30, pp.1196–1216.

Rajabi, F. & McArthur, J.J. (2024) Applying OPTICS with and without PCA for fault detection of fan coil units using building automation system data. Energy and Buildings, 319, p.114368.

Sun, B., Liang, D. & Zhang, H. (2024) An improved graph deviation network for chiller fault diagnosis by integrating the sparse cointegration analysis and the convolutional block attention mechanism. Energies, 17, p.4003.

Tra, V., Amayri, M. & Bouguila, N. (2024) Latent code description for unsupervised AHU fault detection using adaptive adversarial autoencoder. IEEE Transactions on Automation Science and Engineering, 2024, pp.1–14.

You, Y., Tang, J., Guo, M., Zhao, Y., Guo, C., Yan, K. & Yang, B. (2024) Ensemble learning based multi-fault diagnosis of air conditioning system. Energy and Buildings, 319, p.114548.

Zhang, F., Saeed, N. & Sadeghian, P. (2023) Deep learning in fault detection and diagnosis of building HVAC systems: A systematic review with meta-analysis. Energy AI, 12, p.100235.