An Anomaly Detection Using Multi-Modal Datasets in Industrial Control Systems*

Yunsung Kim[1], Gyeongdeok An[1], Jongbae Hwang[2], and Jaecheol Ha[1†]

[1] Hoseo University, Asan-si, Chungcheongnam-do, South Korea
veluv@naver.com, akg13078@gmail.com, jcha@hoseo.edu
[2] NNSP Inc., Seoul, South Korea
jbhwang@nnsp.co.kr

Abstract

Industrial Control Systems (ICS) are responsible for the secure operation of national critical infrastructures, but the growing threat of cyberattacks has highlighted the urgent need for stronger security. Anomaly detection in ICS is a key security measure to ensure system integrity and safety, and has traditionally relied on single-modal approaches using either sensor data or network data. However, sensor-based detection faces inherent limitations in distinguishing between attacks and simple malfunctions, while network-based detection lacks a direct connection to actual physical impacts. This paper proposes a multimodal anomaly detection model that simultaneously analyzes sensor time-series data and network packet data. A 3-layer LSTM is applied to the sensor modality, while a Transformer-based architecture is applied to the network modality to extract latent representations, which are then fused through a fully connected layer to detect anomalies. Experimental evaluation using the SWaT (Secure Water Treatment) dataset from the iTrust research center demonstrates that the proposed model achieves an F1-score of 0.89, showing significant improvements over single-modal models and existing approaches. These results provide experimental evidence that multimodal fusion effectively addresses the limitations of single-modal approaches and enhances ICS anomaly detection performance.

Keywords: Anomaly Detection, Multimodal SWaT Dataset, LSTM Model

1 Introduction

Industrial Control Systems (ICS) are responsible for ensuring the stable operation of physical processes across critical national industries, including smart factories, nuclear and power plants, water and wastewater treatment facilities, and transportation infrastructure. ICS collect and control data through sensors and actuators, while utilizing control devices such as Supervisory Control and Data Acquisition (SCADA), Distributed Control Systems (DCS), and Programmable Logic Controllers (PLC), along with Human–Machine Interfaces (HMI) and industrial network protocols such as Modbus and PROFINET, to maximize operational efficiency and automation [1]. However, ICS are directly exposed to cyberattack threats, which has rapidly increased the urgency of ICS security research. Given

* Proceedings of the 9th International Conference on Mobile Internet Security (MobiSec'25), Article No. 37, December 16-18, 2025, Sapporo, Japan.
† Corresponding author

the high availability and real-time requirements of ICS, cyber incidents can cause significant physical safety risks as well as financial damages. Consequently, anomaly detection in ICS has emerged as an essential security mechanism to ensure system integrity and safety.

Anomaly detection approaches in ICS are generally categorized into sensor-based and network-based methods. Sensor-based detection leverages numerical time-series data, including process variables and actuator values, to identify physical anomalies [2]. While effective in capturing physical deviations, it faces a structural limitation in distinguishing between cyberattacks and simple system faults. On the other hand, network-based detection analyzes industrial protocol traffic to uncover previously unknown attacks or data tampering, but it suffers from weak linkage to actual physical impacts [3].

To overcome these limitations, recent research has proposed multimodal detection techniques that fuse sensor and network data, aiming to improve both detection accuracy and root-cause traceability through integrated analysis. Despite these efforts, several challenges remain, such as data synchronization issues, modality imbalance, and real-time processing constraints. In this paper, we address these challenges by proposing a multimodal model that jointly analyzes numerical sensor data and categorical network data in ICS, and we evaluate its effectiveness through experiments.

2 Industrial Control System (ICS)

2.1 Overview of ICS

ICS are systems designed to monitor and control the physical processes of critical industrial infrastructures such as power plants, water treatment facilities, and chemical plants. An ICS is composed of physical-layer sensors and actuators, control devices such as PLC and DCS, HMI, and the industrial networks that connect them. Unlike traditional IT networks, ICS networks are designed with real-time responsiveness, determinism, and availability as top priorities, leading to significant differences in protocol design and operational priorities [4]. ICS also follow a hierarchical structure: at the bottom layer are field devices such as sensors and actuators, in the middle layer are controllers such as PLCs and DCS, and at the top layer are SCADA systems and HMIs.

However, the growing interconnectivity between organizational systems—a paradigm shift introduced by modern industrial practices—has significantly exposed industrial plants to unprecedented security risks. This vulnerability is primarily driven by the expansion of the attack surface due to the proliferation of the Industrial Internet of Things (IIoT) and increased integration with external communication networks. A representative case is the Stuxnet malware, which manipulated the control logic of Siemens PLCs to operate centrifuges abnormally, causing direct physical damage to industrial equipment. This incident demonstrated that cyberattacks can extend beyond data breaches to inflict severe impacts on the continuity and safety of industrial processes [5]. Similarly, the VPNFilter malware showed that infected networking devices could intercept SCADA protocol communications, block control commands, or manipulate them maliciously. Such incidents reveal that even without immediate physical destruction, the persistent exposure of critical operational data and control commands can lead to severe strategic damage [6]. Another notable example is the Maroochy Shire incident, in which unauthorized wireless access combined with insider knowledge was used to manipulate SCADA systems, resulting in the release of massive amounts of untreated sewage. This case highlighted how the absence of basic security controls in ICS environments can directly translate into physical harm [7]. Collectively, these representative incidents underscore that ICS security vulnerabilities and threats can produce profound consequences across technical, operational, and societal dimensions.

2.2 Anomaly Detection in ICS

Anomaly detection techniques for ICS security can largely be divided into sensor-based and network-based approaches, depending on the type of data utilized. Sensor-based anomaly detection is one of the most extensively studied approaches, where physical process variables such as water level, flow rate, pressure, and temperature are collected as time-series data and compared against normal patterns to identify anomalies. A common method is to train a prediction or reconstruction model and then it uses the error between predicted or reconstructed values and actual observations as the anomaly score. Since sensor data are directly connected to the physical behavior of processes, anomalies caused by sensor manipulation or process disruption can be detected relatively quickly. For example, in Singapore's SWaT testbed, a study combined an RNN-based prediction model with the CUSUM method and successfully detected multiple attack scenarios with low false alarm rates [8]. Moreover, deep learning models such as CNNs and LSTMs effectively learn the temporal dependencies and complex nonlinear patterns of sensor data, achieving higher detection performance compared to traditional statistical approaches.

However, sensor-based detection faces a structural limitation in distinguishing attacks from simple faults. Abnormal signals caused by sensor malfunctions or external environmental factors may resemble attack patterns, leading to increased false alarms and degraded detection performance. Sensor-based approaches are also sensitive to noise and data quality issues, and some studies have reported that disruptions in specific sensors can significantly reduce overall process detection accuracy [9].

Network-based anomaly detection, on the other hand, focuses on analyzing protocol traffic in ICS, such as Modbus, PROFINET, DNP3, and S7Comm, to detect anomalies. Unlike sensor-based methods, this approach directly identifies abnormal patterns at the packet, session, or protocol-command levels. It is particularly effective when adversaries infiltrate ICS networks to insert unauthorized commands or tamper with data, as such activities can be immediately detected at the network layer. Recent studies have demonstrated that deep learning methods such as 1D-CNNs, LSTMs, and Transformers can capture temporal characteristics of packet sequences and correlations among protocol fields, enabling the effective detection of even previously unseen attack types [10]. In addition, unsupervised anomaly detection research has actively explored the use of session-level statistics, flow features, and entropy-based characteristics.

Nevertheless, network-based detection also has structural limitations in explaining direct causal relationships with physical process anomalies. For instance, if an attacker replays legitimate commands or encrypts traffic, it becomes difficult to determine solely from network data whether such activity actually results in process disruption. Furthermore, the inherent real-time requirements and limited bandwidth of ICS networks impose constraints on deploying complex deep learning models.

In conclusion, sensor-based and network-based detection methods each excel in identifying anomalies in physical processes and threats at the communication layer, respectively. However, single-modal approaches alone are insufficient to comprehensively address the diverse attack scenarios and complex process environments in ICS. Accordingly, recent studies have begun to explore multimodal approaches that leverage both sensor and network data, but their application remains limited due to challenges such as data synchronization and real-time processing constraints [11].

3 Dataset

3.1 SWaT Testbed

The Secure Water Treatment (SWaT) dataset [12] was collected from a scaled-down water treatment plant testbed developed at the iTrust Lab of the Singapore University of Technology and Design

(SUTD). The testbed was designed to closely replicate the infrastructure of a modern water treatment facility, providing researchers with a realistic environment for studying cyber-physical security threats

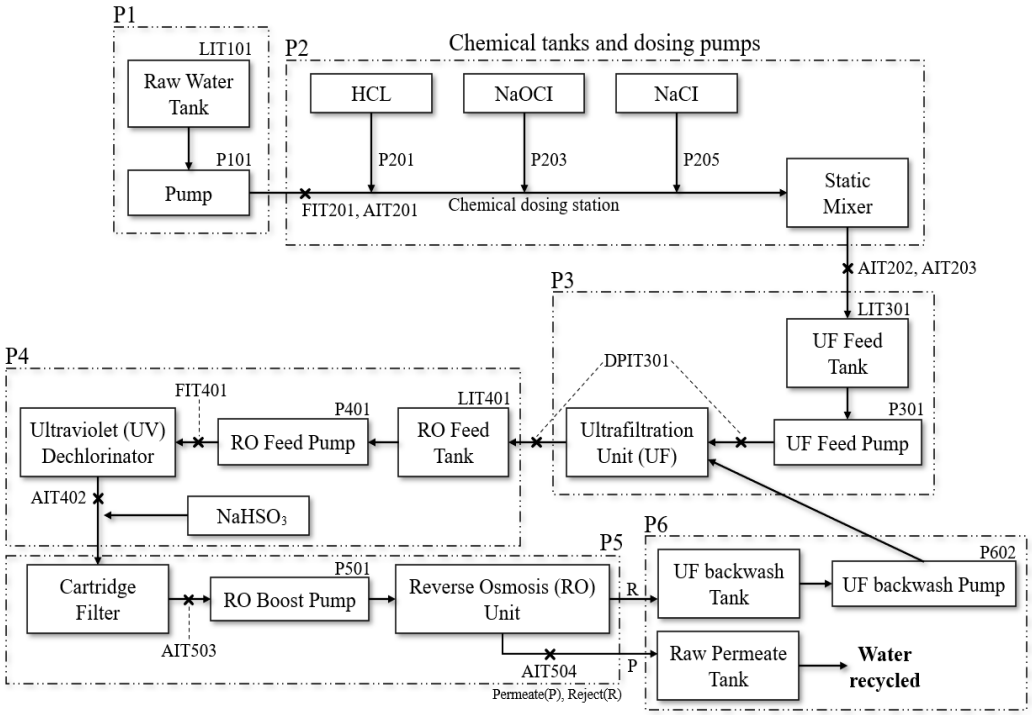

Figure 1: SWaT Testbed Process Flow

in ICS. Covering an area of approximately 90 m², SWaT simulates the complete water treatment process through six sequential stages, ranging from raw water intake to reverse osmosis and membrane backwash.

Figure 1 illustrates the six-stage filtration process of the SWaT testbed, showing its sequential design. Stage P1 involves pumping and storing raw water into the system. This stage contains the raw-water tank LIT101 that measures water level, the inlet pump P101, and the motorized valve MV101, which together regulate inflow to the treatment line. Stage P2 adjusts water quality by injecting chemicals to ensure purification and safety. It includes chemical dosing pumps P201–P206 and analyzers such as AIT201–AIT203, which measure conductivity, pH, and oxidation-reduction potential, respectively. Stage P3 removes impurities using ultrafiltration membranes and is equipped with a feed tank LIT301, feed pump P301, a differential-pressure sensor DPIT301, and flow meter FIT301 that monitor membrane performance and backwash cycles. Stage P4 eliminates chlorine and other contaminants through ultraviolet (UV) treatment, comprising the UV dechlorinator UV401, feed pump P401, and associated sensors AIT402, FIT401 for chemical control and flow regulation. Stage P5 further purifies the water via reverse-osmosis filtration, using RO boost and feed pumps P401–P404/P501, the reverse-osmosis (RO) unit, and multiple analyzers AIT501–AIT504, PIT501–PIT503 that monitor water quality and system pressure. Finally, Stage P6 discharges or recycles the treated water and contains the UF backwash pump P602, backwash tank, and flow meter FIT601 for cleaning the membranes and recycling permeate. Each stage is monitored and controlled by industrial sensors, actuators, and PLCs, coordinated through SCADA systems and HMIs. Together, these 51 sensors and actuators generate the physical-process data of the SWaT dataset, directly reflecting the operational state of the plant.

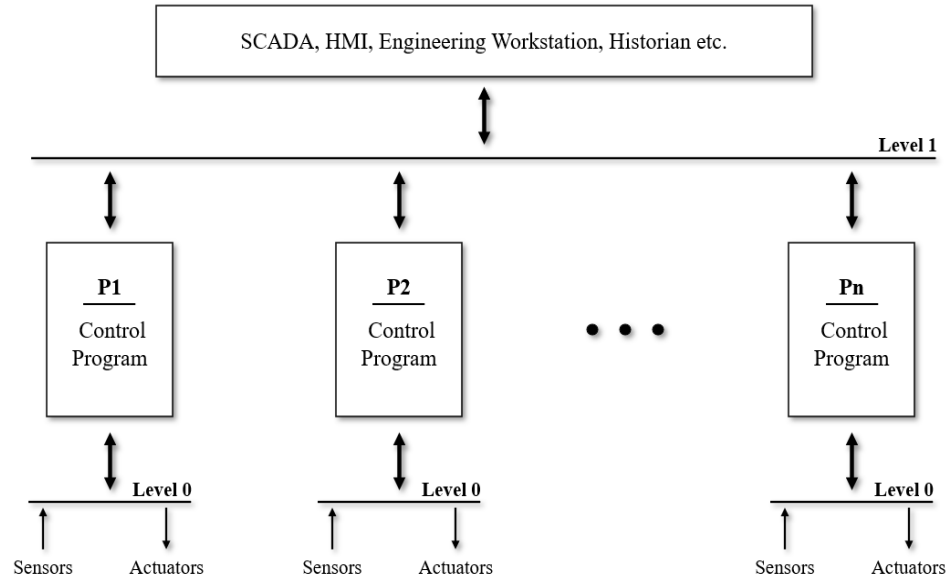

Figure 2: SWaT Testbed Architecture and Control Flow

Figure 2 depicts the control and communication architecture of the SWaT testbed. At Level 0, sensors and actuators operate the physical processes. Each stage is controlled by PLCs that communicate with one another and execute control programs. At Level 1, SCADA systems, HMIs, engineering workstations, and historians are connected to monitor and supervise the overall operation. This hierarchical design is similar to the Purdue model and enables the SWaT dataset to provide both physical data from sensors and actuators as well as network packet communication data.

Data collection began after the plant was initialized in an empty state and stabilized under normal operation. A total of 11 days of operational data were recorded, of which the first 7 days represent normal operation, while the subsequent 4 days include data collected under various cyber and physical attacks. In total, 36 distinct attack scenarios were designed, encompassing four major categories: Single Stage Single Point (SSSP), Single Stage Multi Point (SSMP), Multi Stage Single Point (MSSP), and Multi Stage Multi Point (MSMP) attacks. These scenarios targeted different process stages and control components, including sensors and actuators such as flow transmitters (FIT), level indicators (LIT), analytical instruments (AIT), motorized valves (MV), and pumps (P). The attacks involved behaviors such as valve and pump malfunctions, sensor data manipulation, and chemical dosing errors, all of which directly affect plant stability and water quality.

The dataset is provided in two main forms. First, the physical data include 51 variables from sensors and actuators—such as flow, water level, pressure, and chemical concentrations—recorded at one-second intervals. Second, the network data consist of Modbus/TCP communication logs between SCADA systems and PLCs, which can be used for packet-level feature extraction and anomaly detection research. Both modalities were collected simultaneously and shared synchronized timestamps, enabling direct correlation analysis between network activity and changes in physical process status. This multi-modality architecture enables a comprehensive assessment of cyber-physical interactions.

3.2 Dataset Selection and Partitioning

The SWaT dataset serves as a representative public benchmark that realistically emulates ICS environments. Its academic value lies in the fact that it provides both physical time-series data and network traffic logs, enabling comprehensive cyber–physical analysis. The dataset is unfortunately limited to a single plant configuration and the Modbus/TCP protocol, and therefore does not fully capture the diversity of industrial communication standards or heterogeneous network architectures found in real-world systems. In addition, because data collection was conducted in a laboratory-scale environment, practical factors such as noise, packet loss, and latency are only partially reflected. Nevertheless, the SWaT dataset remains a high-quality resource that quantitatively records cyber–physical interactions and has become a standard benchmark for multimodal anomaly detection research.

In this study, we designed a model for anomaly detection in ICS using the SWaT dataset (A1 & A2, 2015). Sensor data and network data from SWaT were first temporally aligned and synchronized, then individually preprocessed into forms suitable for model training. Dataset partitioning was conducted while preserving temporal order through sequential alignment without overlap, and the distribution of normal and attack data was adjusted to avoid excessive imbalance across the training, validation, and test sets.

Specifically, the proportion of attack data was approximately 5% in the training set, 4.9% in the validation set, and 4.3% in the test set. This ensured that attack data were neither overly sparse nor disproportionately concentrated, thereby allowing each dataset to realistically reflect imbalanced conditions in ICS environments while still enabling fair comparison of anomaly detection performance.

4 Proposed Multi-Modal Detection Model

In this paper, we propose a multimodal neural network anomaly detection model that fuses sensor data and network data collected simultaneously in ICS environments through precise synchronization and alignment. The proposed model independently learns modality-specific representations and integrates them through a late fusion mechanism, thereby enabling holistic cyber–physical anomaly detection.

4.1 Sensor Submodel

Sensor data in ICS are multivariate time-series generated continuously during process control operations. It is essential to capture both long-term dependencies and abrupt short-term variations in such data. Prior studies have demonstrated that LSTM models effectively address the long-term dependency problem in time-series data [13]. Therefore, we adopt an LSTM-based architecture as the sensor submodel.

In this implementation, the LSTM receives input of dimension 51, represented as a time-series sequence $x^{(sen)} \in R^{T\times 51}$. This sequence is processed by a three-layer LSTM with a hidden size of 256 to extract temporal patterns. Its output is passed through a fully connected (FC) block with layers ranging from 256 to 128 to 64 using the ReLU activation function, reducing the feature dimensionality and producing a latent vector $z^{(sen)} \in R^{64}$. Gradient clipping is applied during training to ensure stable convergence, and all input features are normalized using min–max scaling based solely on training-day statistics to prevent data leakage. This design enables the multilayer LSTM to effectively learn nonlinear dependencies within the time series while mitigating overfitting.

4.2 Network Submodel

Network data consist of packet sequences containing multiple categorical fields, where attack behaviors typically manifest not within individual packets but across extended temporal sequences. Due to this characteristic, Transformer architectures have attracted attention for their ability to model long-range dependencies in sequential data. The self-attention mechanism of Transformers enables the model to capture correlations across the entire input sequence simultaneously [14]. Therefore, we adopt a Transformer-based architecture as the network submodel.

The model receives input of dimension 16, represented as a categorical sequence $x^{(net)} \in Z^{P\times 16}$. Each categorical field is mapped through an independent embedding layer, with embedding dimensions adaptively determined according to the vocabulary size. The embedded sequence is then processed by a 2-layer Transformer encoder to learn sequence-level representations. Mean pooling is subsequently applied, followed by a 2-layer fully connected (FC) network to compress the feature space and produce a latent vector $z^{(net)} \in R^{64}$. This design leverages the global contextual learning capability of self-attention mechanism while maintaining manageable model complexity.

4.3 Multi-modal Fusion Model

The latent vectors extracted from the two submodels are concatenated to form a fused representation $\left[z^{(sen)}; z^{(net)}\right] \in R^{128}$. The fusion model is implemented as a shallow MLP to minimize latency and parameter count. All model components are trainable end-to-end, and inference complexity remains suitable for real-time analysis at the control-center level. For reproducibility, training employs BCEWithLogitsLoss with positive class weighting, the Adam optimizer, and gradient clipping. This fused vector is then passed through a 2-layer FC network to generate the final logits, followed by a sigmoid activation function to classify the input as either normal or attack.

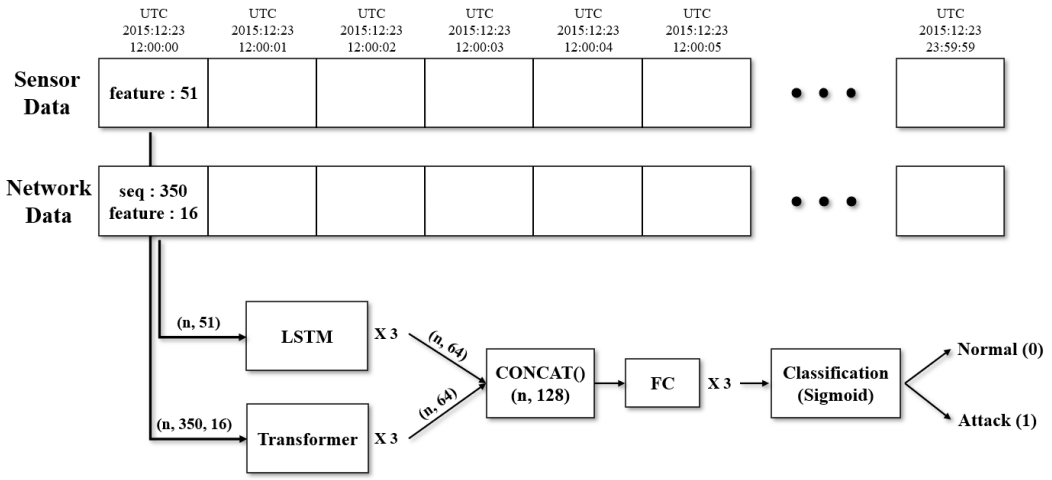

Figure 3: Multi-Modal Anomaly Detection Pipeline in One-Day Window: Sensor and Network Data Fusion

Figure 3 conceptually illustrates the overall workflow of the proposed multimodal anomaly detection model. The framework minimizes information loss by employing modality-optimized architectures and enhances complementarity through a late-fusion strategy. In the proposed workflow, synchronized sensor and network windows are first prepared, and each modality is independently encoded using its respective LSTM and Transformer submodels. The resulting latent vectors are then

concatenated, after which the fusion head computes an integrated anomaly score that is converted into a binary classification result through a fixed threshold.

This late-fusion structure provides a unified representation of both physical process behavior and network communication patterns within an ICS, thereby enabling the detection of complex cyber–physical anomalies. By jointly considering sensor-based anomalies and network-level attack patterns, the proposed model is expected to achieve superior detection performance compared to single-modal approaches.

5 Experiments and Evaluation

To evaluate the proposed model, we used the SWaT dataset (A1 & A2, 2015) collected by the iTrust research center at the Singapore University of Technology and Design. The input for the sensor-only model consists of 51 physical features, representing sensor and actuator values. The input for the network-only model is defined as 16 categorical features extracted from packets transmitted within ICS network traffic. Model training was conducted under a supervised learning setting, and the final output is a binary classification result indicating whether the input corresponds to a normal or attack instance. The input data for all experiments followed the time-ordered partitioning method described in Section 3, with no overlap between sets. Training was performed using Adam, a mini-batch size of 128, gradient clipping, and BCEWithLogitsLoss with positive class weights to mitigate class imbalance.

5.1 Data Preprocessing

For sensor data, Min–Max normalization was applied to all 51 continuous features based on the training set distribution to ensure numerical stability. Each time step was labeled as either normal (0) or attack (1) according to the iTrust-provided attack logs. The normalized data were segmented into fixed-length windows to construct input sequences that preserve temporal continuity.

Network data were processed in parallel and temporally aligned with the sensor data. Unlike sensor readings sampled at 1 Hz, network traffic contains on average about 350 packet logs per second. From the Modbus/TCP communication between SCADA and PLC, 16 key categorical fields were extracted, including source and destination IP addresses, ports, function codes, and SCADA tags. For each field, a dedicated vocabulary was built, and unique integer IDs were assigned to tokens; these IDs were later converted to continuous representations through embedding layers during training.

To synchronize both modalities, the 1 Hz sensor timeline was used as the reference axis. For each one-second interval, all packets within the same timestamp were grouped into a corresponding packet window. These windows were further summarized through statistical features such as packet counts, request–response ratios, and function-code distributions, while the raw categorical tokens were preserved for sequential modeling. This synchronization ensures a one-to-one correspondence between physical process states and communication activities, providing a temporally consistent multimodal input for joint anomaly detection.

5.2 Experimental Results and Analysis

Model performance was evaluated using precision (P), recall (R), and their harmonic mean, the F1-score, as the primary metrics. Precision represents the proportion of correctly identified anomalies among all detected anomalies, while recall indicates the proportion of actual anomalies correctly identified. The F1-score provides a balanced measure between precision and recall. Table 1 summarizes the evaluation results of the models.

Model	Precision	Recall	F1-score
Multi Modal	**0.79**	**0.86**	**0.89**
Sensor only	0.69	0.74	0.82
Network only	0.63	0.68	0.75
1D-CNN [15]	0.86	0.85	0.86
MAD-GAN [16]	0.98	0.63	0.77

Table 1: Evaluation Results and Comparison of the Proposed Multi-modal Anomaly Detection Model

The proposed multimodal model achieved P = 0.79, R = 0.86, and F1 = 0.89, demonstrating the highest performance among all evaluated models. This result outperforms both the sensor-only and network-only models, indicating that fusing information from the two modalities enhances detection capability. Furthermore, comparison with existing studies using the SWaT dataset confirms the superiority of the proposed multimodal model. A 1D-CNN-based model previously reported an F1-score of 0.86, while a MAD-GAN-based model achieved an F1-score of 0.77. By surpassing these approaches, the proposed multimodal model demonstrates effective anomaly detection performance on the SWaT dataset.

In summary, the proposed multimodal model not only outperforms single-modal approaches but also achieves competitive results compared to prior studies, while maintaining stable detection performance even under imbalanced data conditions.

6 Conclusion

To integrate sensor-based time-series data and network-based packet data for detecting cyberattacks in ICS environments, we proposed a multimodal anomaly detection model. The proposed model extracts latent vectors from sensor data using an LSTM and from network data using a Transformer, and then fuses them for final classification. Experiments were conducted using the SWaT dataset provided by the iTrust research center, and the proposed model achieved an F1-score of 0.89, showing significant performance improvements over both sensor-only and network-only models. Furthermore, it demonstrated superior detection performance compared to representative methods reported in prior studies. These findings suggest that multimodal fusion-based detection is a key research direction for overcoming the limitations of single-modal approaches in ICS anomaly detection. Looking forward, future work may focus on developing lightweight models suitable for real industrial environments and validating the approach across diverse datasets, thereby contributing to more practical and effective security enhancements.

Acknowledgment

This work was supported by the Industrial Technology Innovation Program (RS-2025-14842976, Development of industrial network protocol test and evaluation equipment to support smart factories) funded by the Ministry of Trade, Industry & Energy (MOTIE, Korea).

References

[1] C. Konstantinou and O. M. Anubi, “Resilient cyber-physical energy systems using prior information based on Gaussian process,” IEEE Trans. Ind. Inform., vol. 18, no. 3, pp. 2160–2168, Mar., 2022.
[2] M. A. Umer, K. N. Junejo, M. T. Jilani, and A. P. Mathur, “Machine learning for intrusion detection in industrial control systems: Applications, challenges, and recommendations,” Int. J. Crit. Infrastructure Protection, vol. 38, Sep., 2022.
[3] J. Gao, L. Gan, F. Buschendorf, L. Zhang, H. Liu, P. Li, X. Dong, and T. Lu, “Omni SCADA intrusion detection using deep learning algorithms,” IEEE Internet Things J., vol. 8, no. 2, pp. 951–961, Jan., 2021.
[4] K. Stouffer, V. Pillitteri, S. Lightman, M. Abrams, and A. Hahn, “Guide to Industrial Control Systems (ICS) Security,” U.S. Department of Commerce, Washington, D.C., USA, NIST-800-82 (R2), May, 2015.
[5] N. Falliere, L. O. Murchu, and E. Chien, “W32. Stuxnet dossier,” Symantec Corp., Tempe, AZ, USA, White paper, vol. 5, Feb., 2011.
[6] Z. Bederna and T. Szadeczky, “Cyber espionage through botnets,” Secur. J., vol. 33, no. 1, pp. 43–62, Mar., 2020.
[7] N. Sayfayn and S. Madnick, “Cybersafety analysis of the maroochy shire sewage spill (preliminary draft),” 2017.
[8] C. M. Ahmed, J. Prakash, and J. Zhou, "Revisiting anomaly detection in ICS: Aimed at segregation of attacks and faults," arXiv preprint arXiv:2005.00325, 2020.
[9] D. Abshari and M. Sridhar, “A survey of anomaly detection in cyber-physical systems,” arXiv preprint arXiv:2502.13256, 2025.
[10] D. Zhan, W. Zhang, L. Ye, X. Yu, H. Zhang, and Z. He, "Anomaly detection in industrial control systems based on cross-domain representation learning," IEEE Transactions on Dependable and Secure Computing, 2024.
[11] V. Berge and C. Li, “Enhanced anomaly detection in industrial control systems aided by machine learning,” arXiv preprint arXiv:2410.19717, 2024.
[12] J. Goh, S. Adepu, K. N. Junejo, and A. Mathur, “A dataset to support research in the design of secure water treatment systems,” in Proc. Int. Conf. Critical Inf. Infrastruct. Security, Cham, Switzerland, pp. 88–99. Springer International Publishing, 2016.
[13] L. D. Manocchio, S. Layeghy, W. W. Lo, G. K. Kulatilleke, M. Sarhan, and M. Portmann, “Flowtransformer: A transformer framework for flow-based network intrusion detection systems,” Expert Syst. Appl., vol. 241, p. 122564, 2024.
[14] Z. Long, H. Yan, G. Shen, X. Zhang, H. He, and L. Cheng, “A transformer-based network intrusion detection approach for cloud security,” J. Cloud Comput., vol. 13, no. 1, p. 5, 2024.
[15] M. Kravchik and A. Shabtai, “Efficient cyber-attack detection in industrial control systems using lightweight neural networks and PCA,” IEEE Trans. Dependable Secure Comput., vol. 19, no. 4, pp. 2179–2197, 2021.
[16] D. Li, D. Chen, B. Jin, L. Shi, J. Goh, and S. K. Ng, “MAD-GAN: Multivariate anomaly detection for time series data with generative adversarial networks,” in Proc. Int. Conf. Artif. Neural Networks, Cham, Switzerland, pp. 703–716. Springer International Publishing, 2019.