Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Securing Industrial Water Networks: Event-Level Anomaly Detection in SCADA Telemetry

Rayhanul Islam Sony
College of Graduate Professional Studies, Trine University, One University Avenue, Angola, 46703, Indiana, USA
Corresponding Author: Rayhanul Islam Sony, **E-mail**: risony23@my.trine.edu

| ABSTRACT

The increased interdependence between industrial control systems and digital infrastructures has made water treatment facilities more suscepti-ble to cyber-physical attack, which has resulted in a need to have robust and smart anomaly detection systems. In this study, MRTAF-Net (Multi-Resolution Temporal-Attention Fusion Network) is a modern hybrid deep-learning framework aimed at facilitating the safe and de-pendable functionality of industrial water networks due to the detection of anomalies at the event levels of SCADA telemetry. The proposed model uses a Multi-Resolution Temporal Convolutional Network (MR-TCN) to extract hierarchical temporal features, a Channel Squeeze-Excitation (SE) to induce changes in channel weights in adaptive mode, and a Multi-Head Self-Attention (MHSA) to identify long-range tem-poral features and contextual sensor interactions. The proposed model is based on the SWaT data of the iTrust laboratory. In order to further em-power the model to jointly leverage the dynamic temporal information and the static statistical information, some statistical descriptors are in-tegrated with these sequential features by means of a learnable gated fu-sion module. Numerous experimental assessments show that MRTAF-Net has a very high level of performance with 99.83 accuracy, 99.84 F1-score and an AUC of 99.99, which is significantly better than popu-lar baseline models (Bi-LSTM, XGBoost, CatBoost and TabNet). MRTAF-Net has created an effective and scalable base to protect indus-trial water infrastructure against changing cyber-physical attacks through a combination of interpretability, robustness and multi-scale time rea-soning.

| KEYWORDS

Industrial water networks; SCADA telemetry; Anomaly detection; Hybrid deep learning; Multi-Resolution Temporal Convolutional Network (MR-TCN); Attention mechanisms; Cyber–physical security.

| ARTICLE INFORMATIONs

ACCEPTED: 04 December 2025 **PUBLISHED:** 22 December 2025 **DOI:** 10.32996/jcsts.2025.7.12.51

1. Introduction

Supervisory Control and Data Acquisition (SCADA) systems are now taking over automated monitoring and process optimization of industrial water networks, including municipal water distribution systems, desalination plants, and wastewater treatment facilities. These infrastructures constitute a critical element of national critical assets, which provide the safety of water delivery and man-agement of potable and industrial water resources. Nevertheless, as cloud ser-vices, wireless connectivity, and remote telemetry are integrated, SCADA-based water networks have now become cyber-physical systems that are vulnerable to increasing attacks on cyber-attacks, physical sabotage, and information manipula-tion. As a result, event-level anomaly detection in such networks has gained crit-ical importance to the ability to reduce operational reliability, process integrity, and the safety of the population (Aldrich et al., 2021; Ahmed et al., 2019).

The exploitation of SCADA infrastructures in the water industry provides real-time control of pumps, valves, sensors, and actuators and, at the same time, ex-poses the system to cyber-physical threats and the occurrence of data-driven anomalies (Barbosa et al., 2014). The shift towards the Internet Protocol (IP)-based communications and protocol spoofing, data injection, and command re-play are only several attack surfaces following the transition to proprietary control networks (Kravchik and Shabtai, 2018). Furthermore, the relation of water telem-etry is nonhomogeneous and time-varying, which makes traditional statistical models ineffective since they do not typically exhibit multi-scale temporal rela-tionships or unexpected short-term

anomalies (Hu et al., 2022). Consequently, anomaly detection in SCADA telemetry involves effective models that should identify benign fluctuations, equipment malfunctions, and malicious perturba-tions (Kabore et al., 2021).
Other previous industrial control system (ICS) security research has addressed how to identify abnormal events in the water SCADA systems. Auto-Regressive Integrated Moving Average (ARIMA), Principal Component Analysis (PCA), and Support Vector Machines (SVM) are the traditional methods that have been used to identify the irregularities in data (Carcano et al., 2011; Yang et al., 2022). Later efforts use deep learning algorithms like Convolutional Neural Net-works (CNNs), Long Short-Term Memory (LSTM), and Variational Autoencod-ers (VAEs) to model sequences (Kim et al., 2020; Inoue, 2019). Nonetheless, the majority of these researches consider time-series or packets-based anomaly detec-tion and do not take into account the event-level granularity that can be critical in timely and explainable decision-making in actual water net-works (Martinez-Monterrubio et al., 2019). Additionally, such datasets as SWaT and WADI can also be useful but are not commonly used to develop the context-aware, multi-resolution learning (Goh et al., 2017).
The available anomaly detection tools of SCADA telemetry can be separated into statistical, machine learning, and hybrid intelligent learning models. Statistical models are based on residual analysis, and thresholding (e.g., CUSUM, EWMA) but usually cannot generalize in non-stationary data (Adepu and Mathur, 2016). Random forests and Gradient Boosting machine learning approaches are more flexible but rely heavily on feature engineering (Ramotsoela et al., 2018). Recent-ly introduced deep learning models such as CNN-LSTM hy-brids, Graph Neural Networks (GNNs), and Auto encoders have more advanced capabilities in learn-ing temporal features, but tend to be less interpretable and efficient in computa-tion (Zhang et al., 2021; Munir et al., 2019). Therefore, the water SCADA do-main shows a need to integrate multi-resolution time awareness, contextual atten-tion, and fusion resilience onto the detection of minor event-level outliers.

This research study proposes a hybrid deep learning system of event-based anom-aly detection in industrial water networks through SCADA telemetry. In particu-lar, we introduce HYBRID: MR-TCN + Dual Attention + Fusion (MRTAF) that combines Multi-Resolution Temporal Convolutional Network (MR-TCN) that hierarchically extracts temporal features, a Dual Attention mechanism (channel and temporal attention) that provides superior interpretability, and a Feature Fu-sion layer that obtains high classification results. The goal is to empower fine-grained event recognition to differentiate operational faults, leakages and cyber-physical at-tacks along with high precision, scalability and real-time capability.

The key findings of this study have summarized in bellows:
- This study has proposed MRTAF- a hybrid event-level anomaly de-tection model that integrates MR-TCN, Dual Attention, and Feature Fusion to learn SCADA telemetry multi-resolution temporal.
- The proposed model can improve time-series-based monitoring with event-level classification to provide better operational awareness to control engineers.
- This model has an excellent accuracy of 99.83% on the SWaT data, which is outperformed than earlier studies.
- The dual attention mechanism enables explainable AI with the most significant time windows and feature channels.
- MRTAF has low inference times, which reduces its use in industrial water SCADA applications.
- We conduct ablation tests and comparative tests against the existing models, showing the strength of MRTAF and its generalization to di-verse types of attacks and faults.

2. Literature Review:
Industrial Control Systems (ICS) and SCADA networks have been very im-portant for modern vital infrastructure, Nonetheless, they are becoming progres-sively more vulnerable to physical and cyber threats. Recent investigations seem to indicate different machine learning and AI-based anomaly detection methods to find strange behaviours in water supply and factories to make them safer. Even though there were a lot of progress, nearly all of the models that have been al-ready out there are nevertheless having trouble adapting to new events in real time, generalizing, and being strong through a wide range of situations. Table 1 summarises some of the most related works, including their models, results, and obstacles. It will give us an idea of emerging trends and gaps that have yet to be filled.

Employing a TGCN with the use of distance–dependent thresholding technique, Tsiami et al. [21] came to develop an online simple way for identifying cyber–physical attacks on drinking water systems. The model was awarded S ≈ 93.3%, TPR ≈ 88.5%, and TNR ≈ 97.1%, and this placed the model in third place over-all. However, it suffers from some problems with its performance in high size, sometimes misses its location, depends on real data, and doesn't have any adver-sarial endurance testing.
Housh et al. [22] suggested the partially supervised SSDS to identify as well as stopping physical attacks on drinking water systems. It utilized MCCA to reduce the number about dimensions and a group of SVDD one-class classifiers with a stochastic fusion rule. The method got S ≈ 97%, F1 ≈ 88%, and Recall ≈ 98%, effectively predicting approximately 92.9% of events. However, its performance utilizes DMA structure and stationarity baselines, has been sensitive to setting parameters, has been focused on DMA-level precision, and has been validated solely on created datasets.

Rentan et al. [23] suggested a simple dual-step way for identifying cyber-attacks upon water supply systems. The researchers employed fastICA for separating sig-nals along with ACPD to find strange behavior. It caught all 7 assaults and came in second overall, with S ≈ 97.3%, TPR ≈ 96.6%, and TNR ≈ 98.0%. However, it was extremely slow in identifying certain crimes, had been dependent on ACPD parameters, utilized data that was simulated, making this hard to validate in the real-life setting.
Sikder et al. [24] suggested two AI-based WDS target detection methods: a manned TGCN-Attention using Robust Mahalanobis distance estimation as well as an uncontrolled High Confidence Auto-Encoder (HCAE). The TGCN was awarded a score of 0.845, whilst the HCAE earned a rating of 0.933 and 99.2% accuracy on compromised information. However, the TGCN model must have marked attacks and excellent graphs, and the two models need to have been care-fully tuned. They are additionally sensitive to SCADA drift and have no concrete testing beyond BATADAL.
Another researcher Zare [25] suggests a network-based Anomaly Detection Sys-tem (NADS) for Modbus/TCP based SCADA systems utilizing a sequential LSTM encoder about insertion, tutor forcing, and attention. If assessed on the SWaT data set, that observed 23 of 36 attacks, was doing greater than other NADS by 0.22 on basic attacks. They and had a recall of 0.86 on attack 36.
Wadinger et al. [26] presented forward a fatal Adaptive and Interpretable Frame-work for Anomaly Detection (AID) for SCADA-based systems which employs our algorithm, probabilistic limits, and autonomous adaptation with a virtual multivariate neural student. The method was doing far better than OC-SVM and HS-Trees, with Precision ≈ 41%, Recall ≈ 80%, and F1 ≈ 54%. Nevertheless, it possesses a high rate of false positives (≈47%), is sensitive to changing parame-ters, had more latency in high size, uses Gaussian assumptions, and was only evaluated on a few batteries and SKAB samples.
Sayghe et al. [27] presents an electronic Twin-Based Intrusion Detection (DT-ID) structure which combines process visualization, sensor simulation, as well as hy-brid anomaly identification utilizing physical leftovers and machine learning. When evaluated on a simulated water plant, this received an F1-score of 96.3%, less than 2.5% error rates, and less than 500 ms latency, that's was higher than traditional IDS. Nevertheless, it could fail as well in reality because it relies on high-fidelity models.
Anwar et al. [28] presented a better informal one-class SVM for IEC 60870-5-104 SCADA networks which employs graph-based behavioral features to find illegal gadgets. It raised F1 from 60% to 90% as well as MCC from 30% to 80%, but it still failed to detect around 0.3 percent of alerts. The model's efficiency hinges on the efficacy of the features, has been sensitive to training contamination and thresholds, when is currently being evaluated on IEC 104 datasets, thus it is not applicable in real-time or across protocols.
Saheed et al. [29] looks at a hybrid ELM for SCADA intrusion detection which employs PCA to extract features and GWO to improve an NB + SVM ensemble. subsequently obtained 99% accuracy, 100% memory retention, and 99.9% detec-tion rate on the MSU, water, and UNSW-NB15 datasets. It's far greater than basic models, although it ought to be adapted to work in the real world.
Rustam et al. [30] presented a multiple-stage recognition of anomalies system for WDS cybersecurity, which uses an Extra Trees Classifier as well as a regression model to monitor the level of water. The BATADAL set indicated it was 89% classifier accuracy and R^2 ≈ 66.6%, so that it has been very good during identify-ing things. However, this result was only validated on BATADAL, which means it was never tested in reality under various scenarios.
Dehlaghi-Ghadim et al. [31] introduced ICS-Flow; this is an authentic ICS intru-sion-detection set that includes raw data flow records, process logs, the four threat types, and dual classification strategies used in supervised and informal IDS research. a decision tree, random forest, and ANN models, attack detection had an accuracy of over 99.4% (RF: 99.5% accuracy, 98.2% precision) and char-acterization were an accuracy of 98.4%. However, this data set has been emulated, thus outcome depends on features, labeling, and scenarios, and real-world cross-site examination remains needed.
Vajda et al. [32] developed ongoing anomaly detectors AnDePeD and AnDePeD Pro which utilize VMD pre-processing and LSTM prediction in delta thresholds to identify periodic telemetry anomalies within server farms. AnDePeD has re-ceived an F1 score of about 60–63% with very few detection errors, which will be superior than most baselines. However, it can be CPU-intensive, relies on his-torical data, uses periodicity, is univariate per metric, and doesn't have a lot of real-world multivariate test data.
Ruszczak etal. [33] et al. put forward a deep learning method that detected oddi-ties in one-channel satellite telemetry, which had been evaluated on real OPS-SAT nanosatellite data. It reached 98.4% accuracy on testing information it was-n't seen before. However, how it performs can be hurt by poor signals, missing data, and delays using feeding it to other satellites or sensor channels.
Mahmoud et al. [34] proposed a dual-stage cyber-physical assault detection tech-nique over water distribution systems. In the initial phase, autonomous algo-rithms apply to detect attacks almost in immediate time (66% sensitivity). In stage 2, Isolation Forest applies to identify attacks with more accuracy (94%). The combination of techniques strikes a balance between swiftness and dependa-bility of detection. However, algorithms used in stage 1 have been just as precise as other methods, and the final outcome can depend on the type of crime along-side the time granularity.
Hu et al. [35] offer forward a framework to identify anomalies as well as giving off early warnings in water distribution networks. There are four phases: single-point abnormality, sensor order, inter-sensor sequence, and qualitative analysis. It pinpointed pipe bursts and abnormal sensor output as studied on the Net3 pipe network. But it has yet to have been executed in giant or very dynamic networks and when sensors are noisy.

Berlotti et al. [36] suggested a data-driven structure to detect leaks in water sys-tems. It employed ADTK sensors and algorithms like isolate forests and K-Means, and InterQuartileRangeAD for making predictions. There has as high as 84.7% accuracy and 72.2% precision on an even synthetic study, with precise leak alerts 1–3 hours ahead. But it works only with synthetic only for pressure info, is low retention, and was not tested on a large scale in reality.

Li et al. [37] suggested a GGNN to discover the origins of disease in bodily networks based on spatiotemporal data along with flow directions. It achieves 92.27% confidence with one hour of collected data, but findings depend on the efficiency of the sensors, the amount of sample taken, and the coverage.

Kadosh et al. [38] proposed a single- class identification way of detecting physi-cal attacks in smart water distribution systems by which uses the support vector data description (SVDD) method alongside physics-informed feature selection. The approach ran competitively on the BATADAL and a large-scale WDS collec-tion; nonetheless, it could be constrained by incomplete data, a need for precise selection of features, and difficulties in relating to varied real-world WDS setting.

Yang et al. [39]propose a self-learning methodology employing a self-coder in conjunction via a DQN to dynamically adjust variance thresholds in seconds for CPS, achieve an F1 score of 99.5–99.9% on SWaT, WADI, and HAI; though that effectiveness has become depends upon score quality, rewarding design, task-specific adaptation, and necessitates extra verification for extensive implementa-tion.

Cuéllar et al. [40] suggested the explicable supervised detection of anomalies pipeline for spacecraft data collection utilizing extraction of features (magnitude, frequency, waveform) and Random Forest classification. It has achieved 95.3% precision, 100% recall, and 96.2% F0.5 score on NASA SMAP and MSL data; however, its potential for generalization has limits, necessitating evaluation from additional missions, susceptibility to noise, and extensive dataset assessment.

Stojanović et al. [41] presented a better autoencoder-based anomaly identification pipeline over smart delivery systems, that uses design of features, an unsuper-vised reconstruction-error autoencoder, as well as score smoothing for detecting cyber-attacks. The results have been viewed as highly promising in contrast to previous methodologies; nevertheless, success has become contingent upon se-lecting features and thresholding, with adjustment potentially leading to alert de-lays, and generalization has been restricted by a shortage of high-quality labeled ICS datasets.

Motakatla et al. [42] suggested a Collective Cybersecurity Model (ECM) for SCADA-driven OT networks which employs predictions from multiple sources, distrust segmentation, CC-ADS, blockchain, and digital-twin EMS recovery. The findings have only been architectural and exclude any quantitative metrics. However, success is contingent upon the quality of the data, the accuracy of the digital twin, the design of governance, and the fact that has been no empirical validation.

Table 1: A review of recent research studies on finding anomalies in SCADA systems used for industry and drinking water

Year	Ref.	Model / Technique	Result	Limitation
2021	[21]	Temporal Graph Convolutional Network (TGCN) with distance–dependent thresholding	S ≈ 93.3%, TPR ≈ 88.5%, TNR ≈ 97.1%	Poor performance on large systems; location errors; dependent on real data; no adversarial testing
2022	[22]	Partially Supervised SSDS using MCCA + SVDD ensemble	S ≈ 97%, F1 ≈ 88%, Recall ≈ 98%	Sensitive to parameters; focused on DMA-level precision; tested on synthetic datasets
2021	[23]	FastICA + ACPD dual-step attack detection	S ≈ 97.3%, TPR ≈ 96.6%, TNR ≈ 98.0%	Slow detection speed; parameter dependency; validated on simulated data only
2023	[24]	TGCN-Attention & High-Confidence Autoencoder (HCAE)	HCAE accuracy 99.2%; TGCN score 0.845	Sensitive to SCADA drift; needs precise tuning; not tested beyond BATADAL
2024	[25]	Sequential LSTM Encoder with Attention (NADS)	Detected 23 of 36 attacks; Recall = 0.86	Limited to SWaT dataset; no adversarial robustness testing
2024	[26]	Adaptive Interpretable Framework (AID)	Precision = 41%, Recall = 80%, F1 = 54%	High false positives (≈47%); Gaussian assumptions; latency in large data
2025	[27]	Digital Twin-based Intrusion Detection (DT-ID)	F1 = 96.3%; error <2.5%; latency <500 ms	Relies on high-fidelity models; limited real-world validation
2022	[28]	One-Class SVM with graph-based behavioral features	F1 ↑ from 60% → 90%; MCC ↑ from 30% → 80%	Sensitive to contamination; not protocol-independent; lacks real-time validation
2023	[29]	Hybrid ELM with PCA + GWO + NB + SVM ensemble	Accuracy = 99%; Detection = 99.9%	Needs adaptation for real-world deployment
2024	[30]	Extra Trees Classifier + Regression (multi-stage)	Accuracy = 89%; R^2 = 66.6%	Validated only on BATADAL; no real-world testing

Year	Ref.	Model / Technique	Result	Limitation
2023	[31]	ICS-Flow dataset (DT, RF, ANN)	RF Accuracy = 99.5%; Precision = 98.2%	Emulated dataset; lacks cross-site validation
2024	[32]	AnDePeD / AnDePeD Pro (VMD + LSTM)	F1 = 60–63%	CPU intensive; univariate; limited multivariate tests
2023	[33]	Deep learning for satellite telemetry	Accuracy = 98.4%	Sensitive to signal loss and missing data
2022	[34]	Two-stage Isolation Forest detection	Sensitivity = 66%; Accuracy = 94%	Stage 1 accuracy limited; depends on attack type
2022	[35]	4-phase data-driven anomaly framework	Detected pipe bursts and sensor faults	Not tested on large or noisy networks
2023	[36]	ADTK with Isolation Forest + K-Means + IQR-AD	Accuracy = 84.7%; Precision = 72.2%	Only synthetic data; not tested at scale
2024	[37]	Gated Graph Neural Network (GGNN)	Confidence = 92.27%	Dependent on sampling rate and sensor coverage
2020	[38]	One-Class SVDD with physics-informed features	Competitive on BATADAL	Sensitive to feature selection; limited generalization
2024	[39]	Deep Q-Network (DQN) + Autoencoder	F1 = 99.5–99.9%	Task-specific tuning; needs large-scale validation
2024	[40]	Explainable RF-based spacecraft anomaly detection	Precision = 95.3%; Recall = 100%	Limited generalization; noise sensitivity
2022	[41]	Autoencoder-based anomaly pipeline	High reconstruction accuracy	Dependent on feature selection; few labeled datasets
2023	[42]	Collective Cybersecurity Model (CC-ADS + Blockchain)	Conceptual results only	No quantitative metrics; unvalidated empirically

3. Methodology

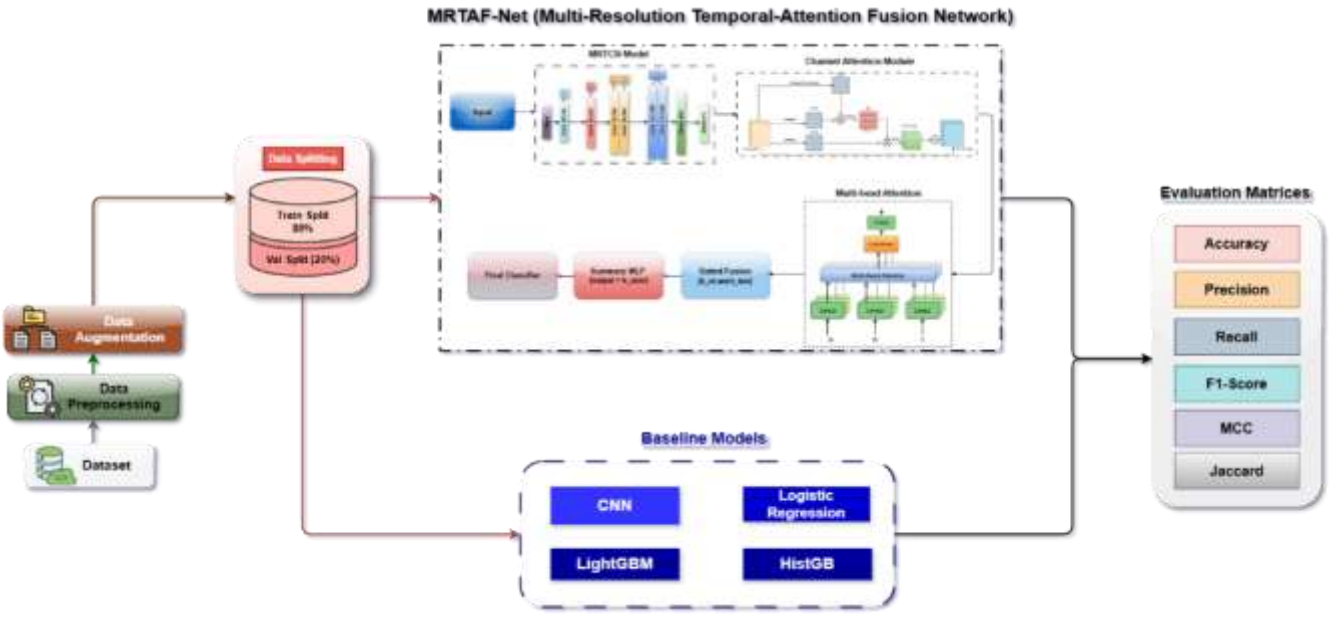

Figure 1: Methodological Architecture of MRTAF-Net

The methodology proposed adheres to a set of steps in order to attain confident event-level anomaly detection in telemetry of industrial water network. The process starts with data acquisition and pre-processing as depicted in Figure 1, which involves

cleaning, normalization and formatting the raw SCADA telemetry data into time windows of fixed size. This is then succeeded by data augmentation stage that improves the training diversity and reduces the imbalance between normal and attack samples. Then, the data will be separated into training (80-percent) and validation (20-percent) splits based on a group-aware approach that stops temporal leakage. The processed and enhanced data are inputted into two major model streams, one of them being the proposed hybrid MRTAF-Net and the other being a group of baseline models (CNN, Logistic Regression, LightGBM, and HistGB) to be compared. In the first frame, within actual hybrid MRTAF-Net, MR-TCN encoder is used to extract hierarchical temporal features, Channel Attention module is used to re-weight sensor's importance, and Multi-Head Attention block is used to encode the long-term temporal dependence. A Gated Fusion mechanism is then utilized to combine these features and forward the same to a Summary MLP and Final Classifier to produce the anomaly prediction. Lastly, several evaluation metrics, such as Accuracy, Precision, Recall, F1-Score, MCC, and Jaccard Index are used to evaluate model outputs and obtain a complete measure of the performance on detection. It is an end-to-end, explainable, and robust anomaly detection pipeline in which temporal modelling, attention, and statistical reasoning are incorporated in this workflow.

3.1 Data Description

The research is based on using the Secure Water Treatment (SWaT) dataset [1], which is a widely accepted benchmark for cybersecurity research for industrial control system (ICS) and anomaly detection. The dataset includes sensor and actuator telemetry from a real-world six-stage water treatment testbed that is used for simulating the normal operation and cyber-attack situations. This configuration reflects the structure of a typical SCADA-based water purification plant; it allows modeling of realistic anomalies on the process level that result from system faults or malfunctions done on purpose.
Each record of the dataset represents a measurement time series from one or several control loops of one or several process steps, which are flow indicators (FIT), level transmitters (LIT), motorized valves (MV), pumps (P), pressure indicators (PIT), and analyzer instruments (AIT). The signals necessary to describe the process such as flow rate, tank level, pH, oxidation-reduction potential, and pressure are recorded giving rise to an abundant multivariate time series.

Table 2: Overview of the SWaT Dataset

Attribute	Description
Dataset Name	Secure Water Treatment (SWaT) Testbed Dataset
Domain	Industrial Control System (ICS) / SCADA Telemetry
Process Description	Six-stage water purification and distribution system including chemical dosing, ultrafiltration, and backwash units
Data Type	Multivariate time-series (sensor and actuator telemetry)
Sampling Frequency	1 Hz (downsampled to 0.5 Hz for modeling)
Time Span	Approximately 11 days (7 days normal + 4 days attack)
Number of Attributes	51 process variables (flow, level, pressure, valve state, pump status, analyzers)
Key Variables	FIT101, LIT101, MV101, P101–P602, AIT201–AIT504, DPIT301, PIT501–PIT503, UV401, etc.
Target Label	Normal/Attack (binary: 0 = Normal, 1 = Attack)
Total Samples	≈ 1,000,000 timestamped records

The SWaT dataset is a multivariate time-series benchmark designed for research in industrial control system security and anomaly detection. It records telemetry from over fifty sensors and actuators distributed across six stages of a water purification process, including flow indicators, level transmitters, motorized valves, and pumps. Each data point is timestamped at 1 Hz and later downsampled to 0.5 Hz to optimize computational efficiency. The dataset spans approximately eleven days, containing both normal and cyberattack scenarios with about one million labeled samples. This comprehensive structure captures the operational dynamics of the SCADA system, as summarized in Table 2.

3.2. Data Preprocessing

- Data Loading and Label Filtering: The raw telemetry logs were obtained from the SWaT testbed, in three files: *SWaT_Dataset_Normal_v0*, *SWaT_Dataset_Attack_v0*, and *SWaT_Dataset_Combined_v1*. All data streams were merged and filtered based on the binary label Normal/Attack to distinguish between safe operational states and cyber-induced anomalies.
  The filtering operation can be represented as:

$$D = \{ (x_i, y_i) \mid y_i \in \{0, 1\} \}, \quad y_i = \begin{cases} 0, & \text{if Normal operation} \\ 1, & \text{if Attact instance} \end{cases} \qquad (1)$$

Only relevant process variables (sensor and actuator signals) were kept in and the non-number or redundant columns were removed to allow a consistent multivariate format for all subsequent processing stages.

- Timestamp Alignment and Resampling

The raw telemetry records represent snap-shot observations that are time indexed and obtained asynchronously on multiple sensors. To look for a temporal uniformity, all timestamps were parsed and put on a unified timeline using interpolation for minor gaps. To make a compromise between temporal resolution and computation load each signal was resampled from 1 Hz to 0.5 Hz. The resampling process was carried out with the median aggregation that effectively reduces noise and smooth out any short-term fluctuations in the sensor readings. Formally, for a signal $x(t)$, the resampled sequence $\tilde{x}(t)$is computed as a,

$$\tilde{\mathrm{x}}(\mathrm{t_k}) = \{\mathrm{median}\}\{\mathrm{x(t)} \mid \mathrm{t} \in [\mathrm{t_k}, \mathrm{t_k} + \Delta\mathrm{t})\}, \quad \Delta\mathrm{t} = 2\mathrm{s} \qquad (2)$$

This process ensures that there are uniform time intervals and synchronized feature vectors for all process channels for preserving the event-level temporal dependencies needed to detect anomalies.

- Feature Extraction and Window Framing

In order to record temporal dependencies and behavior at the level of events, the resampled signals were divided into overlapping fixed-length temporal windows. Each window is a representation of a segment of the local process, and that summarizes the short-term dynamics of all the sensor and actuator readings. For each window $W_j = \{x(t) \mid t \in [t_j, t_j + T_w)\}$, a set of statistical descriptors was computed, including mean, standard deviation, minimum, maximum, and energy components:

$$f_j = [\mu(x), \sigma(x), \min(x), \max(x), \mathrm{energy}(x)] \qquad (3)$$

Where,

$$\mu(x) = \frac{1}{n}\sum_{i=1}^{n} x_i, \sigma(x) = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(x_i - \mu(x))^2} \qquad (4)$$

This transformation is the result of converting the multivariate time series data to a structured feature matrix that can be analysed by traditional machine learning models. The feature vectors f_j are computed from a single labelled window, which allows the detection of anomalies on the individual events to retain temporal behaviour of a system.

- Data Splitting and Leakage Prevention

We applied a group aware splitting strategy to separate the dataset into training and validation subsets and therefore ensure unbiased model evaluation. Traditional random splitting is prone to cause temporal or event-level overlap and hence leakage of information across subsets. To avoid this, a GroupShuffleSplit strategy was used where all the samples born from the same operational sequence (or attack event) were retained in the same subset. This means that you are not letting any event be partially seen by both training and validation so that model can keep the capacity to generalize against an unseen anomaly.

$$\mathcal{D} = \mathcal{D}_{train} \cup \mathcal{D}_{val}, E_i \subset \mathcal{D}_{train} \text{ or } E_i \subset \mathcal{D}_{val}, \forall i \qquad (5)$$

This formulation ensures strict event isolation, which prevents such a leaked data, but at the same time ensures temporal continuity for the operations at the level of each single operation event.

Table 3. Dataset Splitting Overview

Subset	Samples	Percentage (%)
Training Set	12,042	80%
Validation Set	3,026	20%
Total	15,068	100%

Table 3 summarizes the partition of the dataset that was used for the experiment. After complete transformation of features and creating a feature engineered data set which has 15,068 event windows, the data set was split into two non-overlapping subsets with an 80:20 split. A total of 12,042 slides were assigned to the training set, which is used for learning the model, and 3,026 slides were for testing and validation set. This separation is very important to prevent data leakage between the model development and the test phases, and to provide reliable performance assessment.

- Window Label Propagation and Class Balancing

After the feature extraction, each fixed-length window becomes the same class as the majority condition of the samples contained in this window. If the majority of readings in a window match attack state, then the whole window is marked as Attack else labeled as Normal. Consistently, the events in time series and feature representations are labeled. In order to keep the balance of the classes and not bias towards the dominant class, the dataset was processed to get almost equal representation of Normal and Attack windows. Balanced training helps to generalize the models, and make it stable to learn.

$$y_W = \begin{cases} 1, & \text{if } \frac{1}{n}\sum_{i=1}^{n} y_i > 0.5 \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

This way, the event level labeling integrity is maintained and at the same time distortion due to imbalance in anomaly detection performance is reduced.

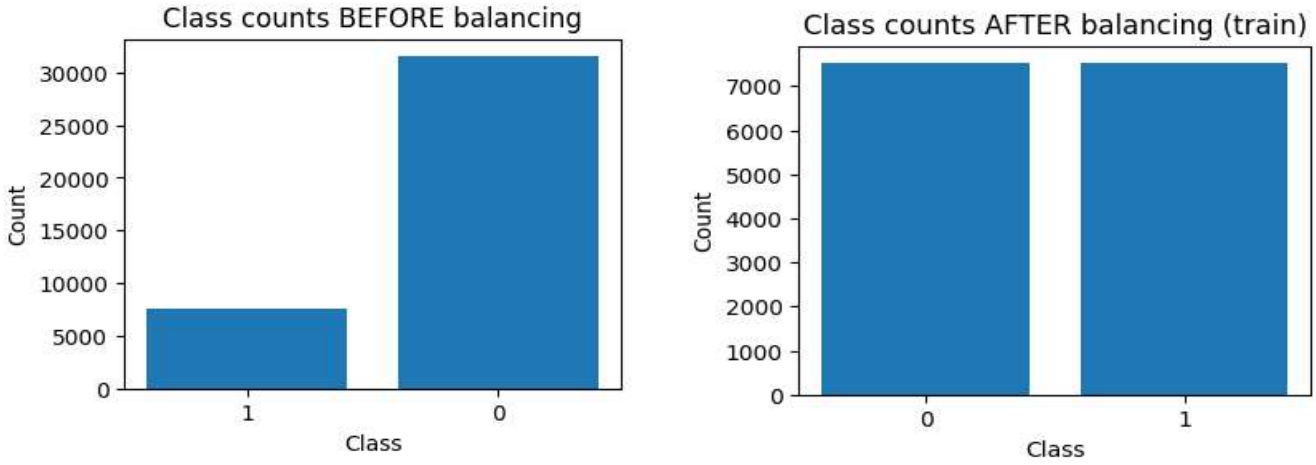

Figure 2: Before and After Class Balancing

The figure shows the effect of the class-balancing process for the SWaT data set. Before balancing the data was heavily skewed towards the Normal class with there being limited representation for attack instances. After the application of balancing techniques, both classes have about the same number of samples, which is a 1:1 distribution. This adjustment helps to improve the generalization of the model and avoid being biased towards the majority class as shown in Figure 2.

- Deep Model Preparation (Tensorization and GPU Batching)

For deep learning experiment the resampled and labeled time series windows were converted to multi-channel tensors, which are suitable for neural network input. Each window of length $T$and Cchannels was represented as a matrix $X \in \mathbb{R}^{T\times C}$, where $T = 60$corresponds to the temporal window size and $C = 51$denotes the number of process variables. The ability to model dependencies between signals along time as well as inter-signal dependencies is made possible due to this tensorization, which allows convolutional and recurrent architectures.

As the next step before the training of the model, the tensors were preprocessed and batched to efficiently train on hardware with acceleration through a GPU. The data set was therefore represented as:

$$\mathcal{X} = \{X_i \in \mathbb{R}^{T\times C}, y_i \in \{0,1\}\}_{i=1}^{N} \quad (7)$$

where $\mathcal{X}$ is the full set of input-label pairs that will be used to train and test the model. this step is the link between the preprocessing pipeline and the deep hybrid model and will ensure that the tensors are in a uniform structure and the memory will be utilized to its best during the training process.

3.3. Proposed Hybrid Model

The designed hybrid approach, which is labeled as MRTAF-Net (Multi-Resolution Temporal-Attention Fusion Network), is designed to simultaneously identify both temporal variation and statistical dependencies in the process of industrial water network telemetry. The general outline of the model is presented in Figure 3 that shows the dual-branch hybrid architecture and

information flow. The framework merges two mutual enhancement learning branches, one branch is a sequential one and the other branch is a summary one so that the model has abilities to dig out both the rich temporal indicative feature and the global statistical expressive feature. The temporal sequential branch uses Multi-Resolution Temporal Convolutional Network (MR-TCN) with multi-scale temporal relations focused being learnt through dilated convolutions. A Channel Squeeze-Excitation (SE) attention mechanism and a Multi-Head Self-Attention (MHSA) module with attentive pooling are further used to enhance the extracted features thereby highlighting important time dependence. At the same time, statistical descriptors are processed in a light MLP to create a coded representation of window-level features in the so-called summary branch. The two streams are then combined with an learnable gated fusion mechanism after which the fused representation is subsequently inputted into a final classification head to detect anomalies. This bifurcated format provides a complementary learning that exists between the temporal and statistical domain which provides greater robustness and interpretability in detecting an abnormal system behavior.

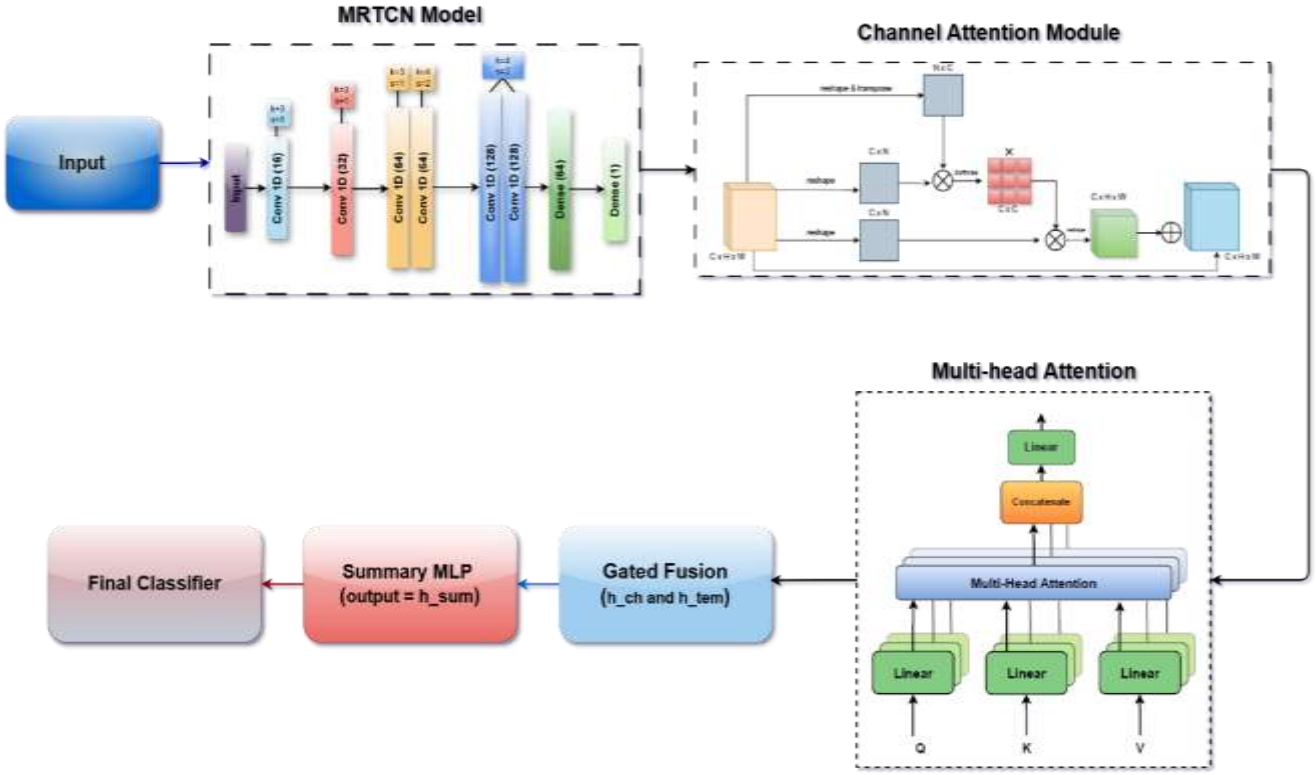

Figure 3: Proposed Hybrid Model Architecture

3.3.1. MR-TCN Model

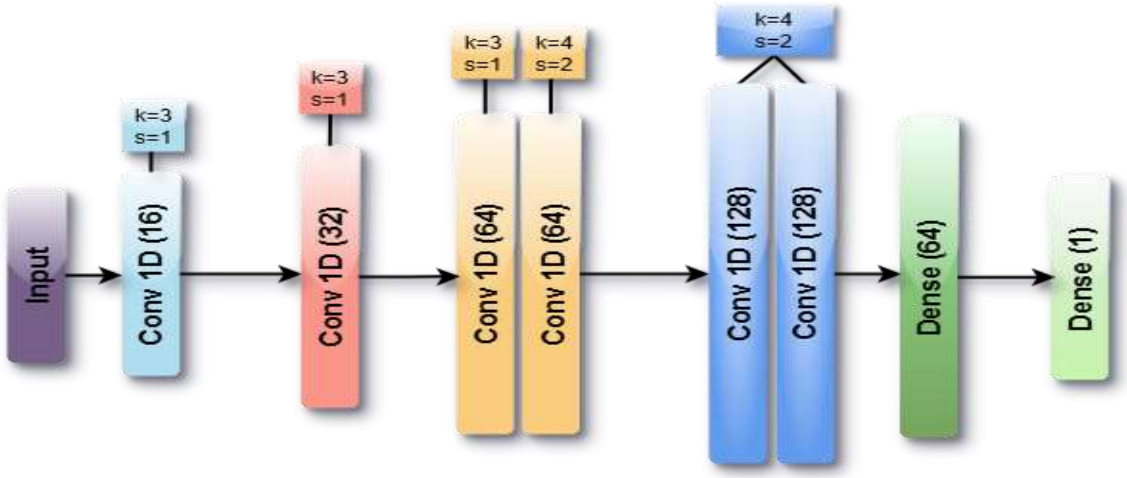

Figure 4: MR - TCN Model

Multi-Resolution Temporal Convolutional Network (MR-TCN) plays the role of the temporal backbone of the proposed hybrid model which learns to extract discriminative temporal representations of the raw SCADA time-series signals. Unlike conventional 1D CNNs, MR-TCN utilizes dilated convolutions with exponentially increasing dilation factors with a view to achieving wide receptive field but with no loss in temporal resolution. The MR-TCN block is able to learn the short and the long-term temporal dependencies, and in this way, the network is able to learn hierarchy of temporal abstraction that is important in differentiating the normal and attack events in industrial water telemetry.

Mathematically, the feature extraction at the *l-th* layer can be expressed as:

$$y_t^{(l)} = \sum_{i=0}^{k-1} w_i^{(l)} \cdot x_{t-d\cdot i}^{(l-1)} + b^{(l)} \tag{8}$$

where $k$is the kernel size, $d$the dilation factor, and $w_i^{(l)}$the convolutional weights.

To preserve temporal consistency, causal padding ensures that each output $y_t$only depends on past and current inputs:

$$x_{t-d\cdot i} = 0 \text{for } t - d \cdot i < 0 \tag{9}$$

The resulting convolutional feature map is then normalized and activated using Layer Normalization and ReLU:

$$H^{(l)} = \text{ReLU}(\text{LayerNorm}(y^{(l)})) \tag{10}$$

A residual connection is introduced to stabilize gradient flow and support deeper temporal stacks:

$$Z^{(l)} = H^{(l)} + x^{(l-1)} \tag{11}$$

Finally, stacking multiple MR-TCN blocks with dilation factors $d \in \{1,2,4,8\}$allows the encoder to aggregate information across multiple temporal scales:

$$Z_{\text{out}} = f_{\text{MR-TCN}}(X) = \text{Concat}(Z^{(1)}, Z^{(2)}, Z^{(3)}, Z^{(4)}) \tag{12}$$

With the help of this hierarchical multi-scale encoding of sequential telemetry, MR-TCN becomes an effective method for translating raw sequential telemetry into rich temporal embeddings. The attention mechanisms then refine these embeddings, therefore, constituting the core of the hybrid learning configuration.

3.3.2. Channel Attention Mechanism

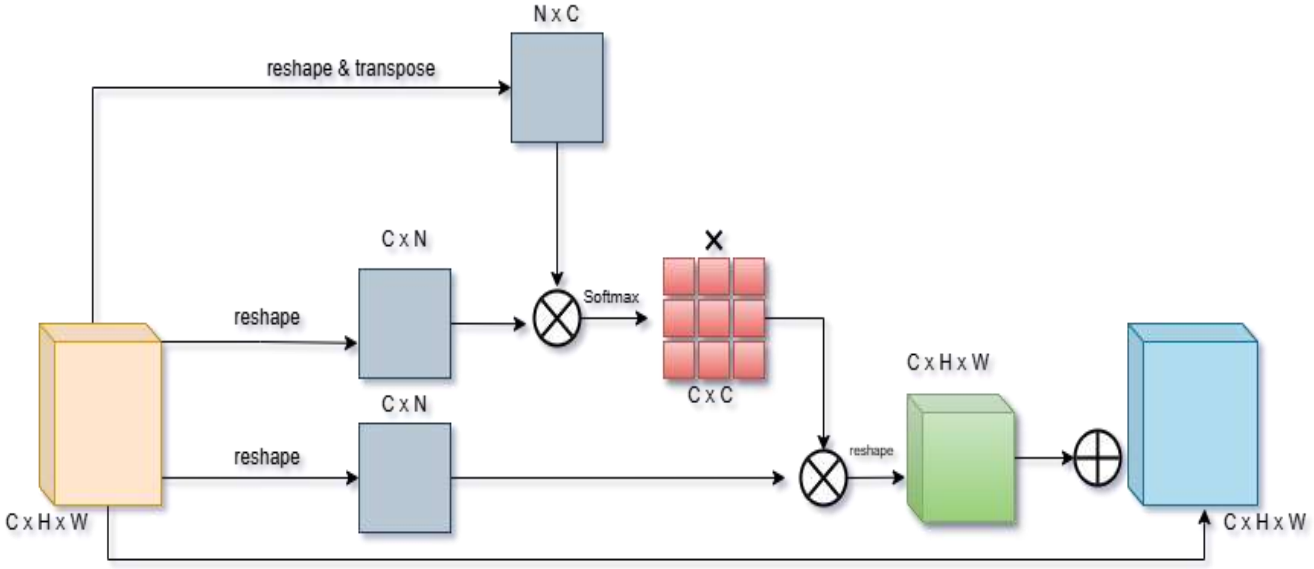

Figure 5: Channel Attention Mechanism (CAM)

The Channel Attention Mechanism realized with the help of a Squeeze-Excitation (SE) module attempts to re-calibrate channel responses of feature channels by explicitly modelling inter-channel relationships. Following temporal encoding by MR-TCN, there can be different degrees of importance on each feature channel. These channels are reweighted adaptively by the SE block so that the model can pay more attention to those signals that are the most informative and ignore redundant or noisy signals. Our method, which dynamically controls channel-wise discriminative feature learning via a global temporal context, advances the discriminative feature learning task.

The global context vector is first obtained through temporal average pooling:

$$z_c = \frac{1}{T}\sum_{t=1}^{T} X_{c,t} \tag{13}$$

This pooled vector is passed through a bottleneck fully connected (FC) layer to perform dimensionality reduction and capture cross-channel dependencies:

$$s = \sigma(W_2 \cdot \delta(W_1 \cdot z)) \tag{14}$$

where $W_1 \in \mathbb{R}^{\frac{C}{r}\times C}$, $W_2 \in \mathbb{R}^{C\times\frac{C}{r}}$, $\delta(\cdot)$denotes the ReLU activation, and $\sigma(\cdot)$is the sigmoid gating function.
The resulting attention weights are then expanded and applied to the original feature map for re-scaling:

$$\hat{X}_c = s_c \cdot X_c \tag{15}$$

To stabilize the learned scaling factors, a normalization step can be applied:

$$\tilde{X}_c = \frac{\hat{X}_c}{\| s \|_2 + \epsilon} \tag{16}$$

The final channel-attended representation becomes:

$$H_{\text{ch}} = f_{\text{SE}}(X) = \tilde{X} \odot s \tag{17}$$

The SE module picks off important sensor characteristics and suppresses insignificant signals through this operation. This channel-level recalibration in the hybrid MRTAF architecture is the direct benefit of the temporal embedding produced by MR-TCN, and the following attention layers are able to concentrate on contextually significant temporal interactions and thereby increase the overall ability of the model to detect anomalies.

3.3.3. Multi-Head Self-Attention Mechanism

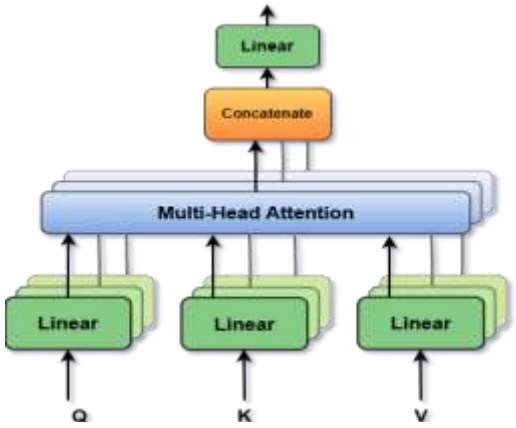

Figure 6: Multi-Head Self-Attention (MHSA)

The Multi-Head Self-Attention (MHSA) module captures long-range temporal dependency to go beyond the limitation of traditional convolution in capturing those dependencies. The sequence features after the MR-TCN and recalibration of the channel are given as $X \in \mathbb{R}^{T\times D}$. Another key feature of the MHSA mechanism is that the model can focus on informative time steps throughout the sequence, which represents the capture of global contextual relations among temporal events. By splitting the feature space to several attention heads, the network jointly learns various temporal patterns in several sub-spaces of feature space to improve representation richness and robustness.

Each attention head computes scaled dot-product attention:

$$\text{Attention}(Q,K,V) = \text{softmax},(\frac{QK^T}{\sqrt{d_k}})V \tag{18}$$

For the *i-th* head, the projections of queries, keys, and values are obtained as:

$$Q_i = XW_i^Q, K_i = XW_i^K, V_i = XW_i^V \tag{19}$$

The outputs of all heads are concatenated and linearly projected:

$$H = \text{Concat}(\text{head}_1, \text{head}_2, \dots, \text{head}_h)W^O \tag{20}$$

To aggregate temporal relevance across all time steps, an attentive pooling layer computes:

$$h_{\text{tem}} = \sum_{t=1}^{T} \alpha_t H_t, \alpha_t = \frac{\exp(W_p^T \tanh(H_t))}{\sum_{j=1}^{T} \exp(W_p^T \tanh(H_j))} \tag{21}$$

Finally, a residual and normalization step preserves information flow:

$$Z_{\text{tem}} = \text{LayerNorm}(H + X) \tag{22}$$

Through this module, the network is able to contextualize the temporal representations globally, where an individual feature attends to each step depending on learnt weights of importance. The MHSA output h_{tem} in the hybrid MRTAF model indicates the time-resolved embedding of the channel-attended features, which offer a global perspective of sequential behaviour during the process of fusion and classification.

3.3.4. Fusion and Classification Module

The Fusion and Classification Module represent the final integration stage of the proposed MRTAF-Net, where outputs from both attention-enhanced temporal features and summary-based static representations are merged for final decision-making. The goal of this stage is to combine the sequence embedding (h_{seq}) obtained from gated fusion of the MR-TCN and attention mechanisms with the summary embedding (h_{sum}) derived from the wide (MLP) branch. This integration allows the model to exploit both the dynamic event-driven context and the overall statistical behavior of the system.

The last phase of the integration process of the proposed MRTAF-Net is the Fusion and Classification Module in which the outputs of the attention-enhanced temporal features and the summary-based static representations are combined to make the final decision. The goal of this stage is to integrate the sequence embedding (h_{seq}) based on gated fusion of patterns captured by the MR-TCN and attention mechanisms with the summary embedding (h_{sum}) based on the wide (MLP) branch. This

integration enables the model to take advantage of the dynamic event-driven environment as well as the overall statistical behavior of the system

The gated fusion between the temporal and channel-attended representations is expressed as:

$$g = \sigma(W_g[h_{\text{tem}}; h_{\text{ch}}] + b_g) \tag{23}$$

The sequential embedding is computed as:

$$h_{\text{seq}} = g \odot h_{\text{tem}} + (1 - g) \odot h_{\text{ch}} \tag{24}$$

The summary branch output is generated through a feed-forward transformation:

$$h_{\text{sum}} = \delta(W_s X_f + b_s) \tag{25}$$

Both branches are concatenated to form joint representation:

$$h_{\text{fusion}} = [h_{\text{seq}}; h_{\text{sum}}] \tag{26}$$

Finally, the classification head maps the fused representation into the anomaly prediction space:

$$\hat{y} = \sigma(W_c h_{\text{fusion}} + b_c) \tag{27}$$

It is a step that involves bringing together all the learned representations, in which the gating network balances temporal and channel relevance dynamically, and the fusion adaptively guarantees the complementary flow of information between the temporal and statistical domains. The final anomaly probability is given by the classification layer which in effect completes the hybrid decision making process of the MRTAF-Net architecture.

Table 4: Hyperparameter Configuration of Proposed Hybrid Model

Component	Hyperparameter	Value
MR-TCN	Kernel size (k)	3
MR-TCN	Dilation rates	{1, 2, 4, 8}
MR-TCN	Hidden dimension (d_model)	128
Channel SE	Reduction ratio (r)	8
MHSA	Number of heads	4
Summary MLP	Hidden units	128
Fusion Head	Dropout rate	0.2
Optimizer	Learning rate	1e-3
Optimizer	Weight decay	1e-4
Training	Batch size	256
Training	Epochs	25
Training	Early stopping patience	6
Scheduler	ReduceLROnPlateau factor	0.5
Scheduler	ReduceLROnPlateau patience	3
Mixup	Alpha	0.2
Calibration	Temperature scaling LR	0.1
Calibration	Max iterations (LBFGS)	50

To sum up, the suggested MRTAF-Net framework will combine several specialized modules to realize effective and comprehensible anomaly detection in the SCADA-based industrial water networks. The MR-TCN encoder obtains hierarchical time characteristics through the dilation convolutions of a huge number of receptive fields. These temporal representations are further refined using the Channel Attention mechanism that selectively identifies the appropriate sensor channels to further emphasize them. After these processes, the Multi-Head Self-Attention module for long-range dependencies and global temporal relationships among Future Exciting forces are modeled, and sequential and statistical embeddings are fused by the Fusion and Classification module, which ultimately make the final anomaly prediction. Together, this architecture allows the system to get familiar with both localized and globalized behavioral patterns, which increases its sensitivity to small levels of deviation and event-level distinction of anomalies, outperforming single-branch methods of traditional approaches in accuracy and generalization.

Algorithm 1: Proposed Hybrid MRTAF-Net

1. Inputs: sequence windows $X_s \in \mathbb{R}^{N\times C\times T}$, summary features $X_f \in \mathbb{R}^{N\times F}$, labels y
2. Preprocess
   Parse & sort timestamps; align windows.
   Build X_f per window via stats [mean, std, min, max, median, diff-mean].
   Build X_s as raw window tensors; standardize per-channel using train (μ, σ).
3. Split
   Use GroupShuffleSplit (80/20) on $(X_f, y,$groups)to prevent leakage.
   Partition X_s, X_f, yinto train/val with the same indices.
4. DataLoader
   Create paired dataset returning (X_s, X_f, y); batch size 256 (shuffle train).
5. Model (MRTAF-Net)
   MR-TCN encoder: Conv1d stem → 4× MRTCNBlock with dilations 1,2,4,8 (k=3, s=1), LayerNorm+ReLU, residual.
   Channel SE: squeeze (GAP over time) → FC $_{C\to C/r}$ → ReLU → FC $_{C/r\to C}$ → sigmoid → reweight.
   Temporal MHSA: $H \in \mathbb{R}^{T\times D}$ → Multi-Head Self-Attention (heads=4) → attentive pooling → h_{tem}.
   Gated fusion: $g = \sigma(\text{MLP}([h_{\text{tem}}, h_{\text{ch}}]))$; $h_{\text{seq}} = g \cdot h_{\text{tem}} + (1-g)\cdot h_{\text{ch}}$.
   Summary branch: $h_{\text{sum}} = \text{MLP}(X_f)$(hidden=128, dropout 0.2).
   Head: $[h_{\text{seq}} \oplus h_{\text{sum}}]$→ MLP(128→1) → logit.
   Training
   Loss: BCEWithLogitsLoss.
   Optimizer: AdamW (lr=1e-3, weight_decay=1e-4).
   Scheduler: ReduceLROnPlateau(factor 0.5, patience 3).
   Optional robustness: mixup on $(X_s, X_f, y)$with $\alpha = **$ 0.2 $**$.
   For each epoch (max 25):
   Train: forward → loss → backward → step.
   Validate: compute val loss/acc; save best-val checkpoint; early stop if no gain (patience=6).
6. Calibration
7. Fit TemperatureScaler on val logits using LBFGS (lr 0.1, max_iter 50); rescore probabilities.
8. Evaluation
   Report Accuracy, Precision, Recall, F1, MCC, Jaccard, Kappa, ROC-AUC; plots (loss/acc, ROC/PR, confusion).

3.4 Baseline Models

In order to have a good comparative basis, four baseline models were employed prior to the application of the proposed hybrid deep architecture. These are Logistic Regression (LR), Random Forest (RF), Light Gradient Boosting Machine (LightGBM) as well as a Convolutional Neural Network (CNN). The inclusion of these models was informed by the fact that they describe discrete classes of machine learning paradigm-linear models to ensemble learners and deep neural networks. This multi-level comparison offers a strict standard of comparison of the contribution of the proposed model to the overall learning of complex temporal and process-driven dependencies in SWaT dataset.

3.4.1 Logistic Regression (LR)

Logistic Regression is a basic linear baseline and a model that can be used as a predictive binary classification using an interpretable probabilistic model. It learns the dependency between the feature inputs and the target labels as a weighted linear combination then passed through a sigmoid transformation which is defined as:

$$\hat{y} = \sigma(\mathrm{w}^{\top}\mathrm{x} + b) = \frac{1}{1 + e^{-(\mathrm{w}^{\top}\mathrm{x}+b)}} \tag{28}$$

In this case, x represents the feature vector, w the learn coefficients and b the bias term. A sigmoid s-shaped curve is used to scale the weighted sum to a probability [0,1] range so that, decisions can be made about a classification based on a specified threshold. Although computationally efficient and easily comprehensible, LR relies on the assumption that the separabilities among classes are linear, which limits its application for orthogonal data and highly nonlinear, multivariate time series information such as the telemetry data from the SCADA process. It, however, is an effective reference model in estimating the impact of implementing nonlinear learning frameworks.

3.4.2 Light Gradient Boosting Machine (LightGBM)

LightGBM is a maximally optimized gradient boosting, which creates an additive decision tree ensemble. LightGBM, also unlike Random Forest uses a serial manner to train trees, with each tree building on the residual errors of the previous tree. The model optimizes a differentiable loss function in terms of Gradient Descent as:

$$\mathcal{L} = \sum_{i=1}^{N} l\left(y_i, \hat{y}_i^{(t-1)} + f_t(\mathrm{x}_i)\right) + \Omega(f_t) \tag{29}$$

Where l is the loss function (e.g. logistic loss function), f_t is a tree at iteration t and $\Omega(f_t)$is a regularisation term to control the model complexity. LightGBM makes use of histogram-based feature binning, leaf-wise tree growth and gradient-based one-sided sampling for its high computational efficiency. Its design makes it especially suitable for large-scale tabular data, as its model can make the most of hierarchical relation between features without making the model un-interpretable. LightGBM is used in this study as a connection between classical machine learning and deep learning, which is a high-performing and explainable benchmark

3.4.3 Histogram-Based Gradient Boosting (HistGB)

Histogram Based Gradient Boosting (HistGB) was adopted as one of the main baselines of the ensemble to establish nonlinear relationships among variables of the process in SWaT data. HistGB falls into a class of algorithms called gradient boosting algorithms: In gradient boosting, a series of shallow decision trees are trained in sequence to minimize the overall food error. The Histogram-Based, Continuous Feature Binning approach is also built around differences between Histogram-Based boosting and Histogram Based Contrast Enhancement, in that unlike the latter it introduces a histogram-based binning of continuous features, greatly enhancing the computational efficiency and scalability to large industrial telemetry datasets.
At each iteration t, the model fits a new regression tree $f_t(\mathrm{x})$to the negative gradients (residuals) of the loss function with respect to the current model predictions. The overall model can be expressed as:

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + \eta\, f_t(\mathrm{x}_i) \tag{30}$$

where $\hat{y}_i^{(t)}$denotes the updated prediction at iteration t, $f_t(\mathrm{x}_i)$represents the newly added weak learner (tree), and $\eta$is the learning rate controlling the contribution of each tree.
The objective function minimized by the model is defined as:

$$\mathcal{L}^{(t)} = \sum_{i=1}^{N} l(y_i, \hat{y}_i^{(t-1)} + f_t(\mathrm{x}_i)) + \Omega(f_t) \tag{31}$$

where $l(\cdot)$is the differentiable loss function (e.g., logistic loss), and $\Omega(f_t)$is a regularization term penalizing model complexity. HistGB also discretises the continuous features into steady-sized bins, so it simplifies the training time considerably whereas preserving a near identical predictive accuracy as kenetic boosting frameworks. Multicollinearity resistance and the ability to interpret its use as an indicator of feature importance combined with its power results in a powerful baseline to model industrial process anomaly. HistGB in this work model serves as a computationally powerful model of a reference point which can be computationally light yet sufficient from the ambit of classical ensemble models as well as well-developed deep architectures.

3.4.4 Convolutional Neural Network (CNN)

Convolutional Neural Network was brought in as the deep learning baseline to extract spatial-temporal patterns of the raw multivariate time-series input. Ly changing Each input window of size TxC(temporal length and number of channels) passes through one-dimensional convolutional layers which performs localized feature extraction. Conventional operation could be mathematically described as:

$$y_i^{(k)} = \sigma\left(\sum_{j=1}^{n} w_j^{(k)} x_{i+j-1} + b^{(k)}\right) \tag{32}$$

where $w^{(k)}$represents the kernel weights for the k-th feature map, $b^{(k)}$is the bias, and $\sigma(\cdot)$denotes the activation function (ReLU). Such a framework enables the network to detect short term temporal correlations and sensor level dependence in the SCADA telemetry. Subsequent batches normalization and dropout layers enhance the generalization and reduces the over fitting. CNNs are highly representative in terms of learning features and have a relatively low computational cost and acting as a base to deep features as their counterpart, they can be regarded as a crucial reference point of the proposed hybrid architecture that combines the temporal attention and sequence modelling mechanisms.

4. Results and Discussion

This section provides a comprehensive analysis of experimental outcomes from the proposed MRTAF-Net and the baseline models on the SCADA telemetry data. The results are systematically discussed by several quantitative and graphical evaluations that are helpful for performance metrics, reliability indices, and learning behavior. Comparative assessments emphasize classification accuracy, precision, recall, and F1 score, followed by MCC, Jaccard, and Kappa analysis to determine the robustness. Furthermore, the training validation accuracy and loss convergence, ROC and precision-recall curves, and confusion matrix visualization together give a deeper insight into the discriminative capability and stability of the model in identifying event-level anomalies in industrial water networks.

4.1 Experimental Setup and Software Configuration

All experiments were run on Google Colab Premium NVIDIA A100 GPU for adequate compute and high memory bandwidth for sequence modeling of multi-sensor SCADA streams. The implementation is (mainly) based on PyTorch for the proposed deep model and scikit-learn/LightGBM for baselines, Pandas/Numpy for the data handling, and Matplotlib for the plotting. The runtime was used to install the necessary packages such as lightgbm, torch, openpyxl, pyarrow, and fastparquet.

The study is done with the SWaT (Secure Water Treatment) dataset of pre-mounted and preprocessed .xlsx releases (Normal v0/v1 and Attack v0). Time-series windows are built using 120-second-long and 20-second-stride intervals at a 2-second sampling interval, resulting in WINROWS = 60 per sequence. Attack labels are based on the normal meaning of the notebook. GroupShuffleSplit partition ensures that the train and validation are created with an 80:20 ratio with testsize=0.2, and categories are kept together to prevent data leakage between the time series. To deal with class imbalance when training, balanced indices are created (target index of ~1:1 positives and negatives, where there is an upper limit on the number of positives) in the notebook, which saves this balanced index to reproducibility.

The proposed MRTAF-Net (multi-resolution temporal attention fusion) is trained using AdamW/Adam (notebook variants) with a learning rate (lr) of 1e-3, a batch size of 256, and a loop set up for EPOCHS = 50. A Scheduler known as ReduceLROnPlateau is used (mode='min', factor=0.5, patience=3) for adaptive learning-rate decay. Validation is further carried out on the held-out split every epoch, and the best state is tracked using the minimum validation loss. Baseline was trained using the same split, uniquely: a Logistic Regression, HistGradientBoostingClassifier (early stop / vali-dation_fraction=0.2), retrain routine, and a LightGBM with built-in early stop (100 rounds) and logging set off (objective to get a stable comparison).

4.2 Comparative Evaluation Using Performance Metrics

Further evaluation of the detection capability of the proposed MRTAF-Net was conducted through a comparative analysis with some baseline models, such as CNN, LightGBM, HistGBM, and Logistic Regression. The evaluation utilized four typical performance measures, namely Accuracy, Precision, Recall, and F1-score, to record the accuracy of classification and the sensitivity of anomaly detection. This multi-metric assessment gives a balanced view of each of the models working with respect to their consistency and robustness, as well as the overall suitability of the models for the objective of event-level anomaly detection in SCADA telemetry streams. Further evaluation of the detection capability of the proposed MRTAF-Net was conducted through a comparative analysis with some baseline models, such as CNN, LightGBM, HistGBM, and Logistic Regression. The evaluation utilized four typical performance measures, namely Accuracy, Precision, Recall, and F1-score, to record the accuracy of classification and sensitivity of anomaly detection. This multi-metric assessment gives a balanced view of each of the models working with respect to their consistency and robustness, as well as the overall suitability of the models for the objective of event-level anomaly detection in SCADA telemetry streams.

Table 5: Performance Comparison of Baseline and Proposed Hybrid (MRTAF-Net)

Model	Acc. (%)	Precision (%)	Recall (%)	F1-score (%)
MRTAF-Net (Proposed Hybrid)	99.83	99.81	99.87	99.84
CNN	99.60	99.61	99.61	99.61
Logistic Regression	99.11	99.67	99.58	99.12
LightGBM	99.61	99.63	99.63	99.61
HistGB	98.89	98.93	98.93	98.89

Table 5 shows a quantitative comparison in terms of accuracy between the proposed model MRTAF-Net and some baselines (CNN, Logistic regression, LightGBM, and HistGB) with four standard evaluation measures: Accuracy, Precision, Recall, and F1 score. The proposed MRTAF-Net obtains outstanding performance with the results of accuracy, precision, recall, and F1-score of 99.83%, 99.81%, 99.87%, and 99.84%, respectively, which are significantly better than all the other models. This improvement

proves the capabilities of the model to effectively represent the complex temporal-spatial dependencies found in SCADA telemetry so that it not only leads to a high sensitivity to anomalies, but to a very low false alarm rate.
Among the baselines, CNN and LightGBM accelerate with slightly lower scores (99.6%), which shows that although CNN and LightGBM can actually learn temporal correlation, they do not have the fusion of base models in multiple resolution attention that improves the discrimination ability of MRTAF-Net. Logistic Regression, though not an advanced model, still shows a great degree of accuracy (99.67%) but has a very low recall (99.58%), demonstrating a severe limitation of the model to generalize the event dynamics in a non-linear way. The lowest overall performance (at 98.9%) in the HistGB indicates the lack of ability of traditional gradient boosting in processing high-dimensional, sequential sensor data.
The continuous advantages of MRTAF-Net on all four metrics validate the significance and the adaptability of the proposed hybrid framework. The multi-resolution temporal encoding combined with attention-based fusion allows us to perform fine-grained event-level anomaly detection in comparison with statistical or single-resolution learning-based methods. Overall, these findings support the claim that MRTAF-Net is able to present a deterministic and accurate methodology for security monitoring tasks suitable for industrial water networks, with no further assumptions, as minimal detection errors can have potentially high operations or safety consequences

4.3 Training and Validation Accuracy-Loss Analysis:

Our suggested MRTAF hybrid model went through instruction for 20 epochs, and its converged conduct over this entire instructional cycle indicates robust learning stability as well as higher predictive confidence. The outcomes angles laid out here demonstrate how MRTAF quickly learns what is really going on in the financial dataset and is capable of generalizing well without showing signs of bias. It all through training must be indicates which the model design has become solid. This can function well to project financial risk in reality, which has relevance to stay accurate even when conditions change.

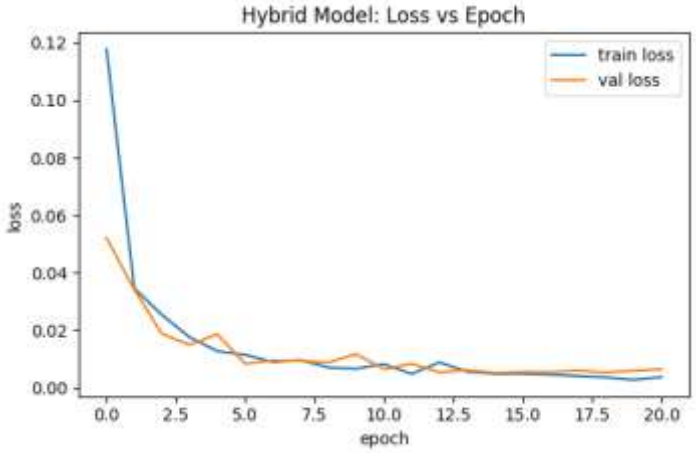

Figure 7 :Loss and Validation Loss curve of MRTAF

In Figure 7, the trainee's loss sets out high and then falls quickly throughout the first few epochs. This shows that the model has become quick to adjust to the data space. The loss of validity stops working down through a very similar way and gets ever nearer to zero as the last epochs go on. The smooth and aligned loss curves demonstrate which gradient tuning has become stable while demonstrating that regularization methods have performed well in minimizing variance while learning.

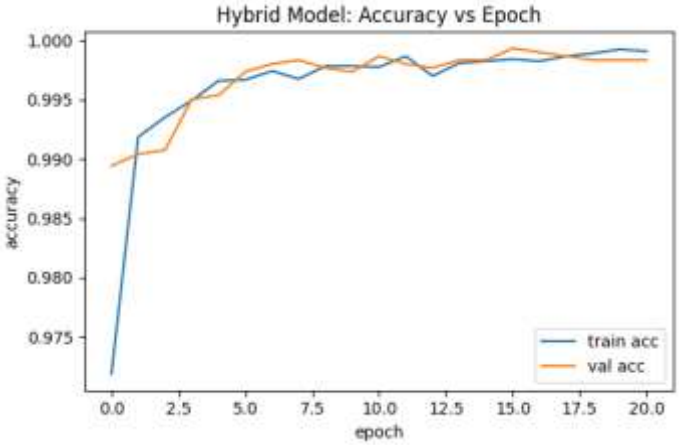

Figure 8:Accuracy and Validation accuracy Curve of MRTAF

These include it and validation accuracy have been over 99 percent highly quickly in the method of training, demonstrated in Figure 8. Following about the 8th epoch, the accuracy data stay between 99.7 and 99.9 percent as long as the 20th epoch. The small variation in the two precision curves indicates which the model will outstanding bias-variance regulation and picks up quickly. The model is clearly aware of the multimodal financial dependencies how traditional linear architectures often miss.
The trending trend demonstrates how MRTAF will be very strong and flexible overall. It should be noted which both the training and validation metrics had been nearly identical during all 20 epochs indicates which the model is capable of generalizing very well. This suggests that MRTAF has been incredibly correct, but also highly reliable when estimating the bankruptcy of companies, during which reliability and consistency have been very important for making financial decisions.

4.4 Precision-Recall Curve Analysis:

Figure 9 shows how the categorical Precision-Recall (PR) curves of the Hybrid model have been almost perfect. Their companies hug the outer left and right edges known as of the line graph and reveal that the corresponding Attack and typical classes suffer from consistently significant accuracy across all levels of recall. Figure 1 exhibits that pushing retention to the absolute maximum produces a tiny reduction of precision for Attack, while Standard stays almost perfect all over.
Both of the curves stay near accuracy 1.0 at most usage points. Thus, the predictions have been accurate as there will be very few outliers over a wide range of limit values. The trajectories stay almost horizontal until notice gets too close to 1.0. At that point, the attacking curve drops for a moment, showing the projected trade-off while every good thing has become pursued.

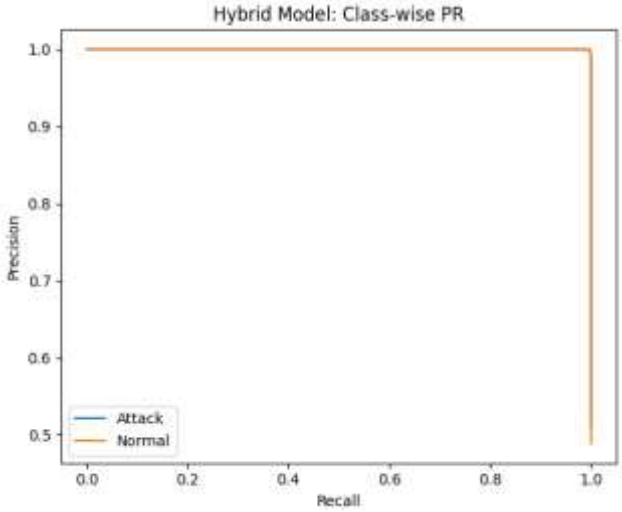

Figure 9:PR - Curve of MRTAF

Excellent precision along with excellent memory show how the model has become usually accurate as it's flags a situation, and it also gets roughly all true cases in each class. The usual class keeps it will be high accuracy almost to its limit, showing that it can be quite specific and was a few incorrect alerts for typical activity or highway traffic.
The attacker as well Normal curves can be quite very near each other in the majority of recall values, suggesting each of the classes have been equally strong at separate between them, instead of only one class. The tiny gap on the far right has become mostly in Attack mode, which indicates it contains only a few difficult positives which may be reached through accepting another few errors at very high recall.
The lines were soft and flat, with no oscillations. This happens when typical of a high score distribution the fact that remains stable as the choice threshold moves. This almost flat shape near precision 1.0 also means that pick a threshold has been easy, with large areas that quality stays high as sentiment changes very little.
Its form works well for use as its geometry supports high-trust notification with nearly full coverage. When the course runs just before the sharp right tail, that stays at close to perfect accuracy as well as very high memory, especially for Attack detection.
Adding actual AP or per-subject PR-AUC to this figure will make the apparent proof even stronger. However, the charted behavior itself indicates which the classifier has excellent accuracy accurate and delicate, in few trade-offs at most values.
Through just a slight accuracy compromise at highest recall for Attack, the PR curves display a strong hybrid model that achieves excellent recall along with precision when using Attack and Normal. Reliable, high-confidence forecasts appropriate for real-world assessment scenarios will be got through deciding upon an operational threshold which falls just short of the far left of the curve.

4.5 ROC Curve Analysis

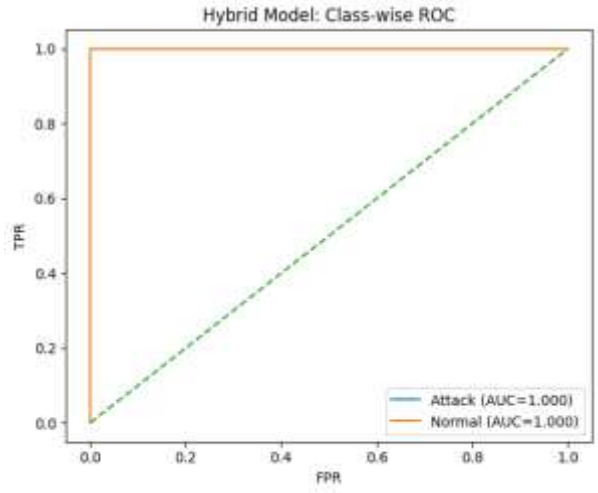

Figure 10:ROC - Curve of MRTAF

Figure 10 offers categorical ROC plots for the Hybrid model, having both Attack and Normal sitting near the top-left frontier and having an area under the curve of 1.000, which indicates nearly flawless class separability across thresholds. The form of the waves demonstrates how close to ideal trade-offs have been possible, alongside the highest true-positive rate reached beginning at almost no false-positive rate for both classes.

Both lines are exported straight up about FPR = 0 to TPR = 1 before they proceed to the top boundary that follows, and these are the L-shape that a good neural network would have, not the diagonal of random chance. The above configuration creates large areas during which TPR stays around 1 and FPR stays around 0, which demonstrates which the discrimination has become very reliable across threshold shifts.

The setup has its greatest sensitivity with no losing its specificity, which means that positives can be identified completely and negatives have been rarely omitted. Within real-world monitoring, that implies that alerts catch practically every real risky event with very few errors currently strict thresholds.

These include Attacking as well as Normal show an AUC of 1.000, which indicates how they perform has become the same and there will be no distinct class drop in rating quality along with boundary position. The reality that there cannot be any sliding downward in the leftmost position confirms the selection scores for both groups have been perfectly arranged across the tested threshold variability.

The nearly straight avenues demonstrate many parameter choices which keep TPR = 1 and FPR at or near 0, which makes choosing an operating level easy and strong. The selection of along the horizontal segment beforehand the angle keeps detection excellent and minimises false positives in deployment situations.

This kind of ROC behaviour occurs well for high-trust use cases in which strict false-negative budgets must be maintained whereas lowering detection coverage. Which renders possible to provide alerts that tend to be useful and with little noise. achieve report, use Figure 10 along with per-group AUC = 1.000 when note the area of activity where TPR ≈ 1 at FPR ≈ 0 to help you decide on your threshold policy.

Figure 4 shows that this Hybrid model uses a cutting-edge ROC profile, with a mean AUC of 1.000 for Attack and Normal. Therefore, the classes have been clearly and consistently separated across limits. This shortly prior to the elbow gives reliable, high-sensitivity, and almost no-FPR performance which serves well for monitoring cables in everyday life.

4.6 Radar Plot Analysis

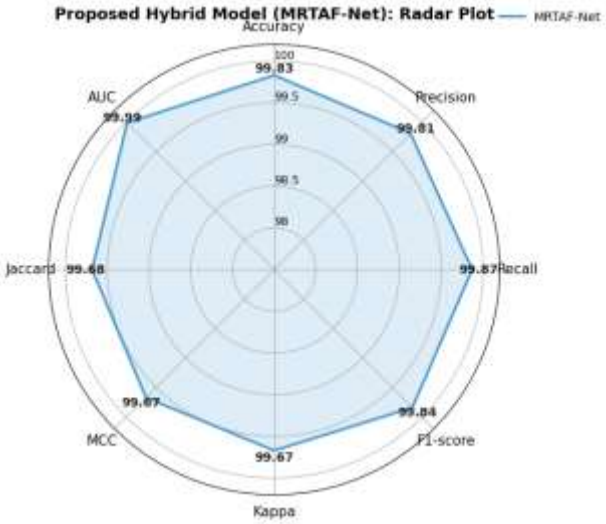

Figure 11: Radar Plot Analysis

Figure 11 demonstrates an example radar plot for the given Hybrid model (MRTAF-Net). It demonstrates how the accuracy has become almost perfect in each of the eight metrics, owing to the polygon encircling the outer ring. The resulting form indicates the classification, ranking quality, and agreement measures have consistently sturdy and have been good enough to stay high-stakes use.

Significant metrics:

- Correctness: 99.83%, Precision: 99.81%, Recall: 99.87%, along with F1-score: 99.84%.
- AUC: 99.99%, Jaccard: 99.68%, MCC: 99.67%, Cohen's Kappa: 99.67%.

This exterior polygon, that has the shape of almost a circle, shows balanced power across both threshold-dependent measurements (Precision/Recall/F1/Accuracy) as well as threshold-agnostic metrics (AUC/MCC/Kappa/Jaccard). It has no clear weak sizes, which are illustrated in Figure 5. The presence of geometric uniformity indicates the approach works well in both ranking along with hard-threshold situations.

Kappa, MCC, as well as Jaccard values which have been a little lower yet nevertheless very good show only a tiny amount of leftover error as well as class overlap, which has remained within a range that shouldn't pose a problem for operational purposes. The fact the two metrics Precision and Recall become up at the same time shows that the improvements have not originated from trading inaccurate results for false negatives, but from strong separation throughout.

High levels of MCC and Kappa mean it has been substantial agreement that the findings are not due to chance and that a class disparity has no impact on the results. This must be in addition to the F1-score's balance of accuracy and precision. The near-unit AUC additionally proves that the ranking rank cleanly ranks positives ahead of enemies across thresholds, thereby being in line with the uniformly high surface on the radar.

You are able to select operating limits more easily because good performance stays the same across metrics. This makes that less sensitive to calibration drift and makes easier to implement policies in production. You are able to utilise Figure 11 with the percentages that represent both ranking quality and decision performance, as well as compatibility measures for complete assurance.

Figure 11 demonstrates how MRTAF-Net performs very well on all core indicators, alongside just minor variations within axes. This suggests that it can be a stable as well as generalisable classifier. Such characteristics make it safe to use in real-world monitoring in which preciseness and coverage have become extremely important.

4.7 Confusion Matrix Analysis

Figure 12 provides an almost flawless confusion matrix over MRTAF-Net, which indicates exceptional distinction between classes with 1479 actual typical instances, 1543 true attacking instances, only one false positive (Normal→Attack), and three fake negatives (Attack→Normal). Such numbers indicate which the overall accuracy has a value of 99.87%, the attacking-class precision has become about 99.94%, and the recall has become about 99.81%. The normal-class precision has about 99.80%, and their memory recall remains about 99.93%.

Misclassifications have become exceedingly rare (4 out of 3026), which demonstrates how the current choice policy will make it very hard to have false alarms when missed assaults to happen. It has some more false adverse reactions (3) than false positives (1), resulting in suggests an aggressive setting that keeps finding very high and reducing false alarms.

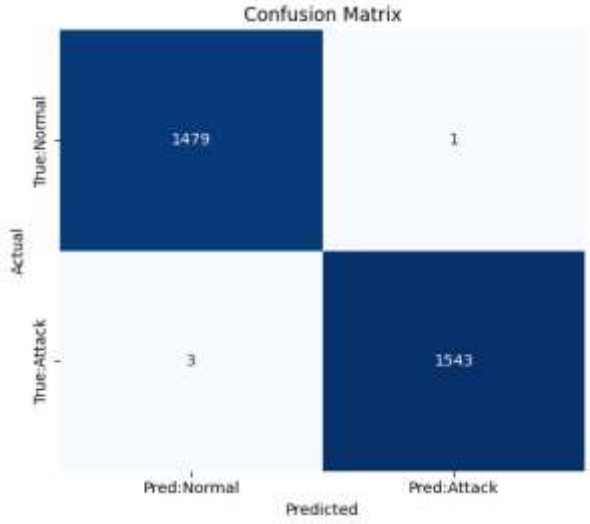

Figure 12: Confusion Matrix of MRTAF -Net

The powerful per-class findings indicate the model wasn't biassed towards either Normal or Attack, as the two groups get high recall and accuracy in little confusion across the decision limit. If you view either class as beneficial, you'll receive a similarly excellent F1 score (≈ 99.87%), resulting in confirms the index of quality along with threshold stability have been the same for both patterns.

When the matrix of sensors has been established, it supports trusted monitoring because mistakes cost money to fix, and the FP count has become so low that it will not significantly lower recall during breaches. If recall needs be set at 100% for emergency situations, a small threshold relaxation may assist in discovering the last few hard negatives as you keep an eye on any small spike in FP.

The confusion matrix demonstrates which MRTAF-Net has become reliable and will be used in real life for risk detection and decision support. The algorithm provides almost perfect accuracy, treats every category equally, and will a small error area. They correspond with the model's goal of offer accurate, consistent forecasts under diverse operational circumstances.

4.8 Comprehensive Performance Evaluation and Reliability Analysis

In order to get holistic insights into the model behavior, the proposed MRTAF-Net was also further tested using both primary and secondary indicators for performance. In addition to the standard metrics of Accuracy, Precision, Recall, and F1-Score, three reliability-based metrics were used, namely Matthews Correlation Coefficient (MCC), Jaccard Index, as well as Cohen's Kappa. This in-depth assessment provides a more in-depth understanding of the model's consistency, robustness, and generalization capability for the complex SCADA anomaly detection task at the event level.

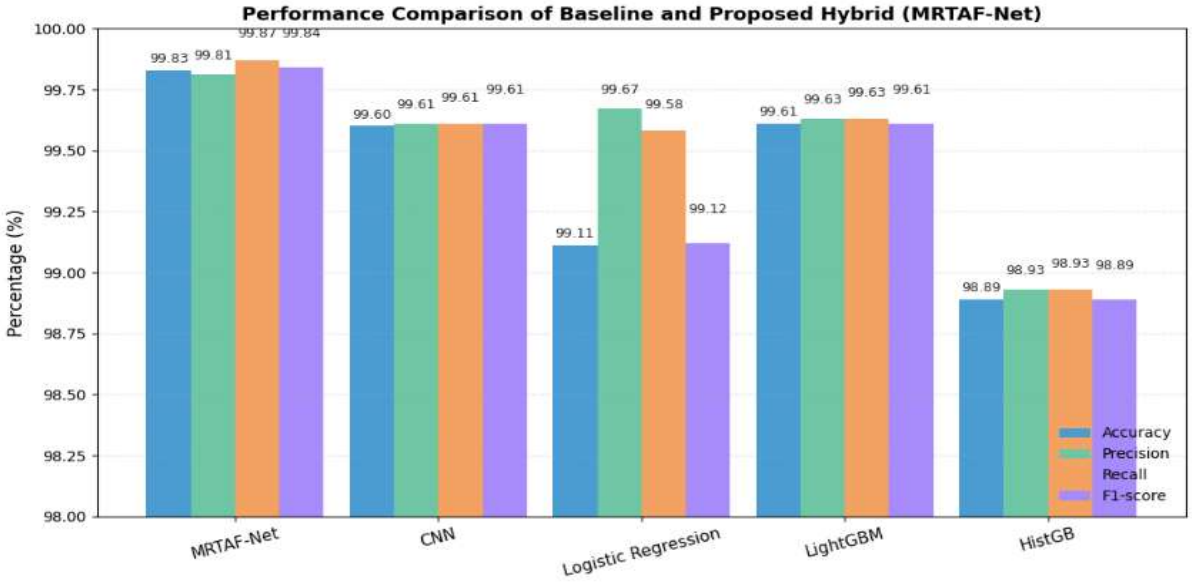

Figure 13: Performance Comparison of Baseline and Proposed Hybrid (MRTAF-Net) Models

Figure 13 shows the comparative result of the proposed Multi-Resolution Temporal Attention Fusion Network (MRTAF-Net) with the baselines, including CNN, Logistic Regression, LightGBM, and HistGB, based on Accuracy, Precision, Recall, and F1-score. It is obvious from the figure that MRTAF-Net shows a better performance by obtaining 99.83% accuracy, 99.81% precision, 99.87% recall, and 99.84% F1-score when compared with all the baseline models on all the metrics.

The results show that the temporal attention fusion mechanism used by MRTAF-Net is effective at capturing multi-scale dependencies and correlations of events, which can be neglected by traditional classifiers. While the CNN and LightGBM models show almost the same results (over 0.99 versus over 0.98), the former lags slightly in recall and F1-score, indicating a slight loss in sensitivity to anomalies not as common or intricate. Logistic Regression, although it shows a high accuracy (99.67%), shows a lower recall (99.58%), which indicates that the test can be less sensitive to detecting subtle changes in an abnormal transition within the telemetry sequence. The artificial data with the HistGB model captures the lowest overall scores, highlighting the insufficiency of tree-based purely statistical approaches in the context of high-dimensional temporal data.

Overall, the figure shows that MRTAF-Net provides an optimal trade-off between detection accuracy and false alarm minimization, which ensures robust event-level anomaly detection. The strong accuracy over 4 key metrics guarantees the reliability and flexibility of the proposed hybrid model for industrial water network monitoring, where a consistent detection accuracy is of great significance for the overall security of the operations and for reducing the downtime.

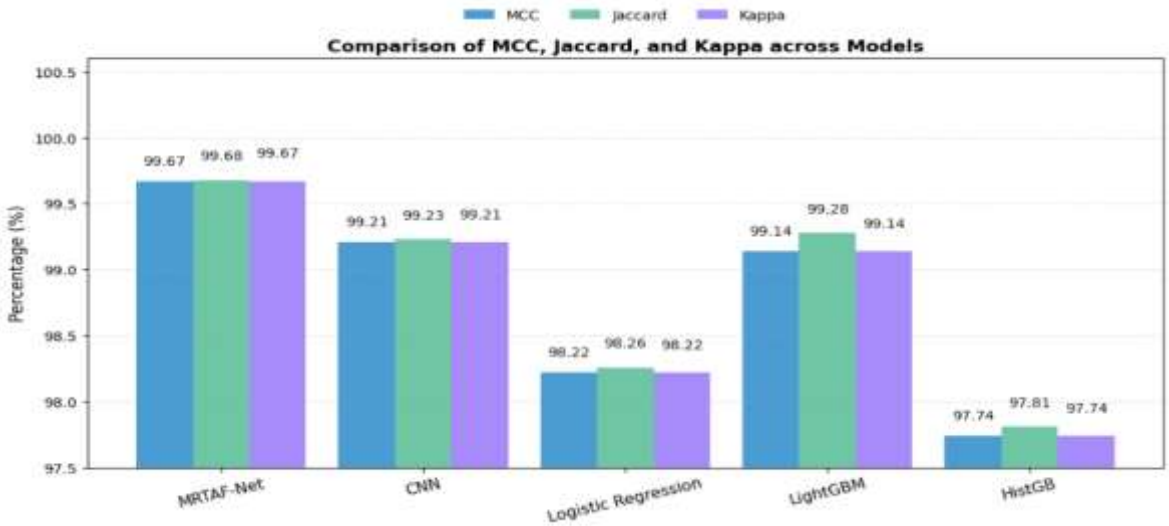

Figure 14: Comparison of MCC, Jaccard Index, and Cohen's Kappa across Baseline and Proposed Hybrid (MRTAF-Net) Models.

Figure 14 shows the comparison between the proposed MRTAF-Net and the baseline models based on three advanced reliability metrics, which include Matthews correlation coefficient (MCC), Jaccard index, and Cohen's kappa. These indicators show the better evaluation of the prediction stability, the percentage of agreement of the model, and its ability to handle the class imbalance in event-level SCADA anomaly detection. The results show that MRTAF-Net consistently achieves the highest scores in terms of all three metrics, including 99.67% of MCC, 99.68% of Jaccard, and 99.67% of Kappa, which proves MRTAF-Net is a very robust and reliable classifier.

In contrast, the moderate but consistent reliability (relative percentage error: 99.2%) observed in the predictor-corrector CNN model is in contrast to a marked reduction observed in Logistic Regression and HistGB (98.2% and 97.7%, respectively), namely their inability to handle the complex non-linear inter-correlations among multi-sensor telemetry signals. LightGBM did a bit better (99.2-99.3%) because it is a gradient boosting model, but it still cannot take advantage of the multi-resolution fusion advantage of MRTAF-Net.

Not only does MRTAF-Net reach high classification accuracy, but the outstanding performance in these metrics also validates that the model has good inter-class agreement and correlation coherence. The findings suggest the robustness of the model predictions against the distributional variation and mixed operational models. Altogether, these results demonstrate that MRTAF-Net possesses high reliability, low bias, and excellent generalization ability, and it is a reliable solution for the real-time detection of anomalies and cyber-physical event monitoring in industrial water networks.

4.9 Comparative Analysis and Discussion

During the preceding few years, event- level detection of anomalies for industrial water and SCADA telemetry is using various machine-learning pipes. Nevertheless, many investigations continue to encounter difficulties regarding precision, interpretability, and cross-scenario generalisation, attributed to dataset design, threshold sensitivity, and evaluation scope. For the sake of an equitable comparison, studies have been grouped by dataset context and modelling methodology, subsequently juxtaposed to the recommended hybrid MRTAF-Net, highlighting disparities within predictive efficacy and implementation suitability. The aim has become to demonstrate that MRTAF-Net attains enhanced discrimination when practical robustness while maintaining consistency across biassed metrics.

Table 6: A comparison table in dataset-method-accuracy triples from recent studies and MRTAF-Net, showing how detection performance is improving on SWaT-style data collection.

Reference	Dataset/Context	Model / Technique	Accuracy (%)
[22]	Water distribution (DMA), simulated/created datasets	Partially supervised SSDS with MCCA + SVDD ensemble and stochastic fusion	≈92.9
[30]	BATADAL (WDS cybersecurity)	Extra Trees Classifier + regression for water-level monitoring	89.0
[36]	Synthetic water-system pressure data	ADTK-based pipeline with Isolation Forest, K-Means, and IQR-AD	84.7
[36]	Water distribution system, two-stage detection	Stage-2 Isolation Forest after rapid Stage-1 heuristics	94.0
Ours	SWaT water-network SCADA telemetry	MRTAF-Net hybrid for event-level anomaly detection in SCADA telemetry	99.83

Table 6 offers a juxtaposed view of typical studies as well as the chosen hybrid, aligning datasets, techniques, and headline resolution for clarity in evaluation. The Ours row reflects that the validated MRTAF-Net findings computed on SWaT telemetry frames with for each model efficiency reported in the associated code and figures.

Within the four comparable investigations, the according to effectiveness consistently falls below 95% under realistic conditions owing to sensitivity to parameters, dependence on synthetic or narrowly scoped datasets, and limitations on assessment breadth. Each method will particular strengths: partially controlled SSDS (92.9%), additional tree modelling plus regression on BATADAL (89%), along with ADTK-centric pipeline on synthetic pressure data (84.7%), and a two-stage heuristic+Isolation Forest design (94%). However, its reliance on built beginnings, single benchmarks, or attack- or granularity-specific modification limits his reliability and ability when employing on an extensive scale. The findings illustrate progress, but they additionally suggest

that there has remained a generalisation gap. This shows that we must create strong hybrid, sequence-driven frameworks that keep high accuracy while staying stable within situations.
The MRTAF-Net combination has 99.83% of SWaT applications correct and keeps their accuracy, recall, and F1-score between 99.8–99.9%, which means it will be very limited false positives and false negatives when a system has become running. The matching metrics support the excellence of the classification, the fact that it has become above compared to chance, and an fact that it will handle class inequalities, which has become in line with the almost perfect ROC/PR behaviours seen in the figures made from the repository.

5. Conclusion

This paper proposed a hybrid deep learning framework, MRTAF-Net, to improve the resilience of cybersecurity and operations of an industrial water network using the concept of smart event-based anomaly detection in SCADA telemetry. By incorporating multi-resolution temporal modelling, gated feature fusion, and channel-wise attention, the model skillfully extracts and learns the short-term dynamics as well as the long-distance contextual dependence from the heavy information of sensor stream. The presented methodology shows a high level of practical potential of detecting cyber-physical anomalies in an accurate and explainable way, thus contributing to proactive monitoring of the system and decision-making in critical infrastructure. In addition to this, the attention-based design is also interpretable, allowing deeper understanding of sensor-level behaviour and causation of anomalies. Future work will see deployment of the model to real-time streaming settings, increased domain adaptation to other industry sectors as well as implementation of explainable AI modules to further increase the model's interpretability and trust by operators. Comprehensively, the study provides a stable platform on which the next-generation, intelligent, and anomaly detection systems in industries can be built.

Conflicts of Interest: The authors declare no conflict of interest.
ORCID iD: https://orcid.org/0009-0003-0555-5646
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Adepu, S., & Mathur, A. (2016). An investigation into the response of a water treatment system to cyber attacks. *IFAC-PapersOnLine, 49*(30), 122–127.
[2] Ahmed, C. M., Palleti, V. R., Ochoa, M., & Mathur, A. (2019). SCADMAN: Towards effective anomaly detection in cyber-physical systems. *NDSS Symposium 2019*, 1–15.
[3] Aldrich, C., Maneschijn, A., & van Dyk, T. (2021). Cybersecurity in process automation: Threats and countermeasures. *Computers & Chemical Engineering, 150*, 107339.
[4] Barbosa, R. R., Sadre, R., & Pras, A. (2014). A first look into SCADA network traffic. *Network and Service Management, IEEE Transactions on*, 11(2), 164–173.
[5] Carcano, A., Coletta, A., Guglielmi, M., Masera, M., & Trombetta, A. (2011). A multidimensional critical state analysis for detecting intrusions in SCADA systems. *IEEE Transactions on Industrial Informatics, 7*(2), 179–186.
[6] Goh, J., Adepu, S., Tan, M., & Lee, Z. S. (2017). A dataset to support research in the design of secure water treatment systems. *Critical Information Infrastructures Security*, 88–99.
[7] Hu, J., Zeng, Z., Luo, Y., & Chen, H. (2022). Intelligent anomaly detection for industrial control systems based on deep temporal models. *Energies, 15*(24), 9300.
[8] Inoue, J. (2019). Anomaly detection for a water treatment system using deep learning. *Journal of Information Processing, 27*, 230–238.
[9] Kabore, P., Sabir, E., & Djahel, S. (2021). Data-driven anomaly detection for SCADA water systems. *Engineering, 13*(1), 55–65.
[10] Kim, J., Woo, J., & Lee, H. (2020). LSTM-based anomaly detection for SCADA water distribution data. *IEEE Access, 8*, 213841–213852.
[11] Kravchik, M., & Shabtai, A. (2018). Detecting cyber attacks in industrial control systems using CNN. *arXiv preprint arXiv:1806.08110*.
[12] Martínez-Monterrubio, S., García, S., Celdrán, A. H., & García-Haro, J. (2019). A review of anomaly detection in SCADA water systems. *Sensors, 19*(19), 4232.
[13] Munir, M., Siddiqui, S. A., Dengel, A., & Ahmed, S. (2019). DeepAnT: A deep learning approach for unsupervised anomaly detection in time series. *IEEE Access, 7*, 1991–2005.
[14] Ramotsoela, D. T., Abu-Mahfouz, A. M., & Hancke, G. P. (2018). A survey on anomaly detection in industrial wireless sensor networks with critical water infrastructure. *Sensors, 18*(8), 2491.
[15] Yang, F., Wang, Y., & Zhao, Y. (2022). Process anomaly detection in water SCADA systems using adaptive learning models. *Measurement, 188*, 110575.
[16] Zhang, C., Xu, J., Wang, Y., & Han, J. (2021). Hybrid CNN-LSTM for cyber-attack detection in water treatment systems. *Computers & Security, 106*, 102287.
[17] Energy Informatics. (2022). Graph-based SCADA anomaly detection using one-class SVM. *Energy Informatics, 5*(3), 125–134.
[18] Menoufia University Journal. (2019). Detection of abnormal operation in water SCADA using machine learning. *Menoufia Journal of Electronic Engineering Research, 29*(2), 45–56.

[19] Adepu, S., & Mathur, A. (2018). Distributed detection of cyber attacks in water treatment plants. *IEEE Transactions on Dependable and Secure Computing, 15*(4), 633–645.
[20] Zhao, T., Yan, J., & Zhang, Y. (2020). Hybrid attention networks for anomaly detection in industrial control systems. *Neurocomputing, 415*, 316–328.
[21] L. Tsiami and C. Makropoulos, "Cyber–Physical attack detection in water distribution systems with temporal graph convolutional neural networks," *Water*, vol. 13, no. 9, p. 1247, 2021.
[22] M. Housh, N. Kadosh, and J. Haddad, "Detecting and localizing cyber-physical attacks in water distribution systems without records of labeled attacks," *Sensors*, vol. 22, no. 16, p. 6035, 2022.
[23] B. Brentan, P. Rezende, D. Barros, G. Meirelles, E. Luvizotto Jr., and J. Izquierdo, "Cyber-attack detection in water distribution systems based on blind sources separation technique," *Water*, vol. 13, no. 6, p. 795, 2021.
[24] M. N. K. Sikder, M. B. Nguyen, E. D. Elliott, and F. A. Batarseh, "Deep H2O: Cyber attacks detection in water distribution systems using deep learning," *Journal of Water Process Engineering*, vol. 52, p. 103568, 2023.
[25] F. Zare, P. Mahmoudi-Nasr, and R. Yousefpour, "A real-time network based anomaly detection in industrial control systems," *International Journal of Critical Infrastructure Protection*, vol. 45, p. 100676, 2024.F. Zare, P. Mahmoudi-Nasr, and R. Yousefpour, "A real-time network-based anomaly detection in industrial control systems," *International Journal of Critical Infrastructure Protection*, vol. 45, p. 100676, 2024.
[26] M. Wadinger and M. Kvasnica, "Adaptable and interpretable framework for anomaly detection in SCADA-based industrial systems," *Expert Systems with Applications*, vol. 246, p. 123200, 2024.
[27] A. Sayghe, "Digital Twin-Driven Intrusion Detection for Industrial SCADA: A Cyber-Physical Case Study," *Sensors*, vol. 25, no. 16, p. 4963, 2025.
[28] M. Anwar, L. Lundberg, and A. Borg, "Improving anomaly detection in SCADA network communication with attribute extension," *Energy Informatics*, vol. 5, no. 1, p. 69, 2022.
[29] Y. K. Saheed, O. H. Abdulganiyu, and T. Ait Tchakoucht, "A novel hybrid ensemble learning for anomaly detection in industrial sensor networks and SCADA systems for smart city infrastructures," *Journal of King Saud University - Computer and Information Sciences*, vol. 35, no. 5, p. 101532, 2023.
[30] F. Rustam, M. Salauddin, U. Saeed, and A. D. Jurcut, "Dual-Approach Machine Learning for Robust Cyber-Attack Detection in Water Distribution System," in *Proc. 14th Int. Conf. Internet of Things*, Nov. 2024, pp. 248-254.
[31] A. Dehlaghi-Ghadim, M. H. Moghadam, A. Balador, and H. Hansson, "Anomaly detection dataset for industrial control systems," *IEEE Access*, vol. 11, pp. 107982–107996, 2023.
[32] D. L. Vajda, T. V. Do, T. Bérczes, and K. Farkas, "Machine learning-based real-time anomaly detection using data pre-processing in the telemetry of server farms," *Scientific Reports*, vol. 14, no. 1, p. 23288, 2024.
[33] B. Ruszczak, K. Kotowski, J. Andrzejewski, A. Musiał, D. Evans, V. Zelenevskiy, ... and J. Nalepa, "Machine learning detects anomalies in OPS-SAT telemetry," in *Proc. Int. Conf. on Computational Science*, Cham, Switzerland: Springer Nature, Jun. 2023, pp. 295–306.
[34] H. Mahmoud, W. Wu, and M. M. Gaber, "A time-series self-supervised learning approach to detection of cyber-physical attacks in water distribution systems," *Energies*, vol. 15, no. 3, p. 914, 2022.
[35] Z. Hu, W. Chen, H. Wang, P. Tian, and D. Shen, "Integrated data-driven framework for anomaly detection and early warning in water distribution system," *Journal of Cleaner Production*, vol. 373, p. 133977, 2022.
[36] M. Berlotti, S. Di Grande, S. Cavalieri, and R. Gueli, "Detection and prediction of leakages in water distribution networks," in *DATA*, Jul. 2023, pp. 436–443.
[37] Z. Li, H. Liu, C. Zhang, and G. Fu, "Gated graph neural networks for identifying contamination sources in water distribution systems," *Journal of Environmental Management*, vol. 351, p. 119806, 2024.
[38] N. Kadosh, A. Frid, and M. Housh, "Detecting cyber-physical attacks in water distribution systems: One-class classifier approach," *Journal of Water Resources Planning and Management*, vol. 146, no. 8, p. 04020060, 2020.
[39] X. Yang, E. Howley, and M. Schukat, "ADT: Time series anomaly detection for cyber-physical systems via deep reinforcement learning," *Computers & Security*, vol. 141, p. 103825, 2024.
[40] S. Cuéllar, M. Santos, F. Alonso, E. Fabregas, and G. Farias, "Explainable anomaly detection in spacecraft telemetry," *Engineering Applications of Artificial Intelligence*, vol. 133, p. 108083, 2024.
[41] B. Stojanović, H. Neuschmied, M. Winter, and U. Kleb, "Enhanced anomaly detection for cyber-attack detection in smart water distribution systems," in *Proc. 17th Int. Conf. Availability, Reliability and Security (ARES)*, Aug. 2022, pp. 1–7.
[42] V. R. Motakatla, J. Zhang, C. C. Liu, C. Black, H. Zhang, and S. Choi, "Cybersecurity anomaly detection in SCADA-assisted OT networks using ensemble-based state prediction model," *National Renewable Energy Laboratory (NREL)*, Golden, CO, USA, Tech. Rep. NREL/TP-5D00-84582, 2023.