Journal of Applied Automation Technologies

Research Paper

Transformer-Based Security Anomaly Detection for Wireless Sensor Networks

Sara Kovačević[1], Milena Đorđević[2] and Dunja Gajić[2,*]

[1] Faculty of Technical Sciences, University of Kragujevac, 32102, Čačak, Serbia
[2] Faculty of Electrical and Computer Engineering, University of Belgrade, 11000, Belgrade, Serbia
*Corresponding author: dunja901@etf.bg.ac.rs

Abstract. Although Wireless Sensor Networks (WSNs) are already being used in the development of cyber-physical systems, they are susceptible to several security risks due to their open architecture and few resources. In order to handle high-dimensional time-series data and different kinds of assaults, this study introduces a security anomaly detection technique for WSNs based on Transformer architecture. Eleven different types of attacks and over 240,000 labeled time-series windows were created by merging real data from several sensors and simulated attack situations in multiple networks. Richer information can be extracted from sensor nodes and sequential data using a Multi-Head Self-Attention-based Model and Spatial-Temporal Feature Encoding. With an average area under the ROC curve (AUC) of 0.976, an F1-score of 0.942, and more than 93% recall on an unbalanced test set, the Transformer-based framework outperforms the basic model, Long Short-Term Memory network, and Isolation Forest, according to the experiment results. Additionally, the detection rate decreases by less than 7% under adversarial attacks or increased signal noise, and the system operates steadily in the presence of noise and variations in the environment or parameters. As a result, it is evident that the Transformer approach has improved detection stability and accuracy and is still viable for high-speed use in WSN defense. In order to enhance the security of wireless sensor networks for monitoring vital infrastructure and intelligent environments, this study suggests an effective and comprehensive solution.

Keywords: *Wireless Sensor Network, Anomaly Detection, Security, Transformer*

Received on 31 October 2025, Accepted on 08 April 2026, Published on 15 April 2026

Introduction

Wireless Sensor Networks (WSNs) are currently crucial parts of many contemporary cyber-physical systems, and they are being used to realize a variety of applications that call for industrial automation, intelligent transportation, environmental observation, medical monitoring, etc. The first two have enabled real-time, distributed, and scalable data collection and utilization across multiple domains [1,2]. WSNs are currently expected to be increasingly stable, independent, and dependable for the safe and effective running of society as more and more vital institutions start using them. However, the challenges of guaranteeing the security and reliability of WSNs have also gotten worse with deeper interconnections and larger deployment scales; as a result, network security has received increasing attention in recent years [3,4].

WSNs are intended to be good, however because of their open-communication medium, limited resources, and unattended deployment, they are highly exposed to a variety of security issues [5,6]. Denial-of-service (DoS), data injection, selective forwarding, sinkhole attacks, Sybil attacks, eavesdropping, and physical node manipulation are examples of attack types that can significantly impair data availability and authenticity and interfere with network operations [7,8]. Due to the difficulty of addressing the aforementioned issues, many conventional anomaly detection techniques, including statistical and rule-based models, as well as traditional machine learning methods for complex, high-dimensional, and time-series data in WSNs, like SVM and KNN, have demonstrated limited adaptability [9,10]. Many of the current approaches fail to take use of the complex space-time linkages in sensor network data and are unable to adjust to new attack modes [11,12]. Because of this, the false alarm rate in practical applications and the detection rate of synthetic benchmarks are both

comparatively low [13,14]. More intelligent, adaptive, and context-aware anomaly detection techniques that can manage the non-stationary and heterogeneous properties of WSN security data are in high demand, as numerous research groups in this field have recently noted [15].

In order to solve the aforementioned issues, we provide a Transformer-based framework for security anomaly detection in WSNs. Our method will reliably and adaptably identify anomalies in dynamic and hostile contexts because Transformer design performs well in modeling long-term temporal dependence and many associated properties. The main achievements of this paper are as follows: (1) Design of a new anomaly detection architecture that combines temporal-spatial feature encoding with advanced attention mechanisms to achieve high-precision threat scoring in complex WSN environments; (2) Extensive performance evaluation against state-of-the-art machine learning and deep learning baselines on real and benchmark datasets, and clearly demonstrate superior detection accuracy and robustness; and (3) Practical discussion on interpretability, computational efficiency and deployment feasibility, providing a path forward for the next generation of intelligent WSN security solutions.

Background and Related Work

WSN Security Threats

Wireless sensor networks provide serious security issues since they are dispersed over space, have little processing capability, and are often used in a variety of settings, including unsecure ones. One of the most dangerous is a Sybil attack, in which a single malevolent node poses as numerous others, harming resource allocation, data fusion, and routing protocols. These assaults have the potential to lower network efficiency and rapidly deteriorate the dependability of collaborative processing and in-network consensus [16]. In sinkhole attacks, an adversary advertises an ideal route to attract network traffic. As a result, packets are either rejected or intercepted, isolating a significant portion of the network from the sink node and potentially causing massive data loss and service interruptions. Another type of attack is selective forwarding, which damages data integrity and network coverage by dropping just the targeted packets after infection to make detection challenging [17].

Eavesdropping exposes the weakness in wireless transmission security, allowing for the acquisition of data, including sensor readings and network information, for use in industrial espionage or targeted assaults in the future [18]. Any alterations or disclosures of this data will seriously endanger human life, and the damage to society will be severe in sectors including vital infrastructure, the environment, and healthcare. Risks include physical node takeover, resource depletion attacks such flooding and over-the-air programming abuse, and manipulation of routing updates, as demonstrated by controlled experimental investigations and real-world security incidents [19]. Strong, flexible, and scalable anomaly detection algorithms for Wireless Sensor Networks (WSNs) are desperately needed because these dangers are introduced by the widespread and frequently unsupervised deployment of sensors.

Anomaly Detection in WSNs

The aforementioned risks have influenced the path of WSN security anomaly detection research. Fixed thresholds, predetermined patterns, or other previous knowledge of typical behavior were needed for the first several methods, which did not rely on statistics or rules. These approaches are appropriate for small networks or controlled environments, but they frequently lack flexibility and resilience in the face of the diversity and changes in WSN environments [20]. A variety of traditional machine learning algorithms, including support vector machines, k-nearest neighbors, decision trees, and ensemble classifiers, have been employed as an intermediate step to increase the range of flexibility. The aforementioned models are generally more adept at identifying non-linear correlations and learning from labeled data, but their actual efficacy is still constrained by the requirement for meticulous feature engineering and their inability to manage high-dimensional and dynamic streams of sensor data [21].

Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks were developed to handle time-series data and extract spatial information, respectively, in response to the current surge in large data and complicated data. The aforementioned techniques are effective at automatically learning representations and are able to identify subtle or high-order correlations that shallow and rule-based models frequently miss [22].

Deep learning models for WSN anomaly detection still have serious drawbacks despite some of their benefits. Their practical application in resource-constrained WSN situations is limited by issues such data imbalance, fluctuations in sequence length, challenges in integrating data from different types of sensors, and high processing needs. Additionally, the majority of deep models now in use make the assumption that the environment is fixed and do not take into consideration environmental abnormalities or the adaptive or dynamic behavior of sophisticated network attacks [23]. A more context-aware model is being designed in response to the aforementioned gap.

Transformers in Time-Series and Security

Because of its strong attention mechanism, the Transformer has recently been successfully used in cybersecurity and time-series analysis to represent complex sequences. Transformers, in contrast to conventional recurrent and convolutional neural networks, are based on a self-attention mechanism that may dynamically focus on various segments of the input sequence, regardless of how close they are in space or time. In order to successfully learn both long-range and short-range dependencies, a novel solution to the issues of vanishing gradients and narrow receptive fields in previous models has been discovered [24].

Transformers have recently been used in a number of fields, including large-scale anomaly detection in industrial systems in current research, network traffic analysis, and intrusion detection in IoT contexts. These are appropriate for analyzing heterogeneous sensor data in a WSN because they may be processed in parallel and have varying input sequence lengths [25]. An all-encompassing anomaly score that takes into account both abrupt deviations and prolonged subtle changes in network behavior can be generated by using a multi-head attention mechanism to automatically merge spatial and temporal information. Enhance the model's performance and give more information about the model's decision-making process by displaying the attention weights; thus, make real applications more transparent and trustworthy by making them easier to understand. Transformer-based models are well suited for the rising demands of WSN security anomaly detection, according to the aforementioned theoretical advantages and experimental findings.

System Design

Data Collection and Preprocessing

The selection and preprocessing of several high-quality datasets is the foundation of a comprehensive and useful anomaly detection system for Wireless Sensor Networks (WSNs). To guarantee the reliability and validity of model validation and experimental results in this work, we will combine real-world sensor data, synthetic traffic produced by high-fidelity simulators, and publicly accessible WSN security datasets. Notable open benchmarks that have acquired widespread attention in the research community include WSN-DS and UNSW-NB15, which function as fundamental resources and offer annotated samples of typical operation and different kinds of assaults. The well-known simulators NS-2 and OMNeT++ can be used to model WSN communication behavior under different threat scenarios and expand the empirical evidence for anomaly detection studies in the absence of proprietary or real-world deployment data.

Numerous network types were chosen, such as star and mesh, tree, etc., and these networks typically had between dozens and several hundred nodes. These nodes generate periodic and event-driven data, including temperature, humidity, voltage, light intensity, routing hops, transmission duration, and node ID, at a sampling rate of one to ten Hz. To improve downstream algorithms' understanding of context and network dynamics, add supplementary qualities to the dataset that characterize both the spatial organization and protocol behavior. Additionally, a variety of attack vectors were injected into the data, including selective forwarding, which discards targeted packets, sinkhole attacks, which lure traffic in and then selectively drop or modify messages, and Sybil attacks, which pose as numerous nodes to violate logical topology. Blackhole attacks, direct injection of malicious values, and resource exhaustion from floods were other threat scenarios. Gather a range of data annotation using a fine-grained classification system; identify whether each event is benign or abnormal, and categorize fine-grained attack types when appropriate. Regression-verified injections in simulated environments and ground-truth event logs in real data were used to construct labels, which were then periodically validated by security analysts to minimize noise or ambiguity in the ground-truth.

To maintain the initial imbalance and time-series structure of attack data and benign operation logs, stratify the training and test sets at random. In order to prevent data leakage and guarantee an objective evaluation of model generalization, approximately 70% of all samples were used for model training, 15% were used for validation and hyperparameter tuning, and the remaining 15% were set aside exclusively for inference benchmarking.

Preprocessing improves the convergence and generalization of the high-capacity model. The information has never been changed. Forward filling was used for short-term gaps in sensor data, whereas median substitution was used for longer absences or node-specific deficiencies. Missing sensor data could be the result of a momentary failure or a purposefully created issue. A sliding window median filter was used to minimize outliers and stochastic jitter that are frequently present in physical deployments. For feature comparison and stable, effective optimization during model training, continuous-valued features, like sensor readings and node voltages, were normalized to have a mean of zero and a standard deviation of one. For example, the protocol type and node ID will be one-hot encoded, and other crucial information is represented by learnt embedding vectors. The category characteristics that must be incorporated in the neural network for training have been changed.

To better capture non-stationary effects and short-term dynamics, the model's input feature set was enriched by constructing rolling window aggregates such as moving averages, local differences, and rolling variances. For a node i at time t, the temporal moving average for a window of length w was computed as

$$MA_i(t) = \frac{1}{w}\sum_{k=0}^{w-1} x_i(t-k) \qquad \text{Eq.(1)}$$

where $x_i(t)$ denotes the normalized observation at time t. Local trends, crucial for detecting slowly evolving attacks or gradual drifts, were quantified as

$$\Delta x_i(t) = x_i(t) - x_i(t-1) \qquad \text{Eq.(2)}$$

In order to do both fine-grained anomaly detection and whole-system threat recognition, the whole feature vector at each time step includes pointwise sensor status, local context dependence, and the network-level context.

All pretreatment pipelines permitted batch inference over arbitrarily large datasets, were implemented in a reproducible, modular manner, and could be seamlessly integrated with downstream model architectures. The Transformer-based anomaly detection system that follows will be able to identify both micro-level device anomalies and macro-level collective disturbances in the security environment of WSNs thanks to the aforementioned thorough and methodical data processing.

Figure 1 displays the extracted feature space for a few typical nodes, together with their label structure and temporal changes, in accordance with the rules for academic publications and to improve the visual information's clarity. The schematic illustrates the multi-modal and multi-scale nature of WSN security data, which is the basis for anomaly detection and is the outcome of several sequential sensor readings, protocol behaviors, ground-truth anomaly annotations, etc.

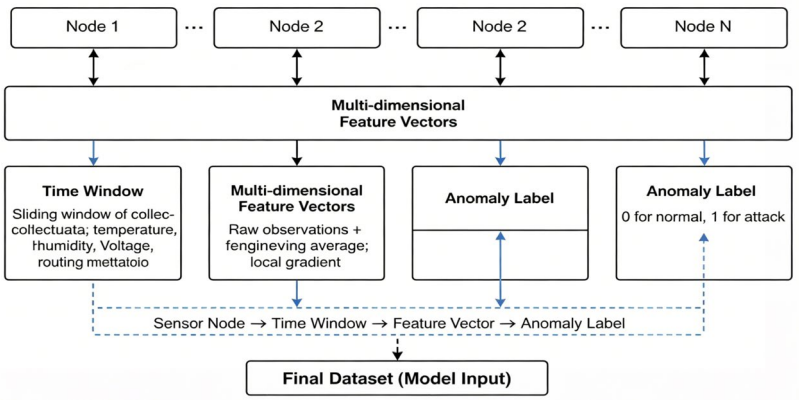

Figure 1. Example Feature Space of a WSN Security Dataset

Transformer-Based Anomaly Detection Architecture

A modified Transformer created to address the inherent issues of multi-modal, time-series data from wireless sensor networks forms the basis of the Anomaly Detection methodology in this study. To effectively extract both geographical and temporal information and distinguish between genuine security threats and typical variations in sensor behavior, optimize the entire detection pipeline.

Preprocessing sensor inputs and choosing designed features yield input sequences, which are then organized as successive temporal snapshots for every network node. Use a sliding window for time-series analysis of the detectors rather than taking into account each sensor's reading separately.

trajectory for every node. Mathematically, if $x_i(t)$ denotes the normalized feature vector for node i at time t, the model operates on a stacked window $\mathbf{X}_i = \left[x_i(t-w+1), \ldots, x_i(t) \right]$, where w is the window length. To integrate all sensor sources, input tensors are assembled into sequences of $N \times w \times F$, where N is the number of nodes and F is the number of features per node.

Before being processed by the Transformer, each input sequence undergoes an embedding transformation. This process projects high-dimensional sensor features into a latent space, followed by the addition of positional encoding. The latter ensures temporal order is maintained and can be formally expressed as

$$PE(t, 2k) = \sin\left(\frac{t}{10000^{\frac{2k}{F}}}\right)$$
$$PE(t, 2k+1) = \cos\left(\frac{t}{10000^{2k/F}}\right) \quad \text{Eq.(3)}$$

where PE denotes the positional encoding vector at time t, and F is the embedding dimension. The embedded and positionally encoded representations are aggregated and then passed into the attention layers.

Central to the architecture is the self-attention mechanism, which enables the model to learn relationships and dependencies across all temporal positions and between multiple nodes. Specifically, for each sequence, the self-attention function is described as

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^{\top}}{\sqrt{d_k}}\right)V \quad \text{Eq.(4)}$$

where Q, K, V represent query, key, and value matrices derived from the input via learned linear projections. Multiple self-attention heads are executed in parallel, each learning distinct facets of temporal and cross-node relationships, with outputs concatenated before being further processed through feedforward layers. This design fundamentally empowers the network to simultaneously reason about fast, local changes and slow, global shifts indicative of complex anomalies.

Each subsequence's anomaly score, which is the model's final output, is read as the node's assessed probability of experiencing a security anomaly at that particular moment. Use binary cross-entropy loss to optimize these probabilities so they are closer to the ground-truth labels.

$$\mathcal{L} = -\frac{1}{M}\sum_{j=1}^{M} \left[y_j \log\, p_j + (1-y_j)\log(1-p_j) \right] \quad \text{Eq.(5)}$$

where y_j is the true label, p_j the predicted score, and M the batch size. A decision threshold on the score is calibrated based on development set metrics such as ROC-AUC or F1 confidence, ensuring a balance between sensitivity and precision suitable for WSN security demands.

A combined architecture diagram of all the aforementioned procedures is displayed in Figure 2. From the raw sensor data input to feature embedding, positional encoding, stacked self-attention layers, and ultimately a dense output map for the anomaly determination, this single diagram illustrates the complete process. To make

it evident to academics and practitioners how the spatial-temporal patterns spread across the system to accomplish the ultimate security objectives, all computation stages and routing channels have been labeled.

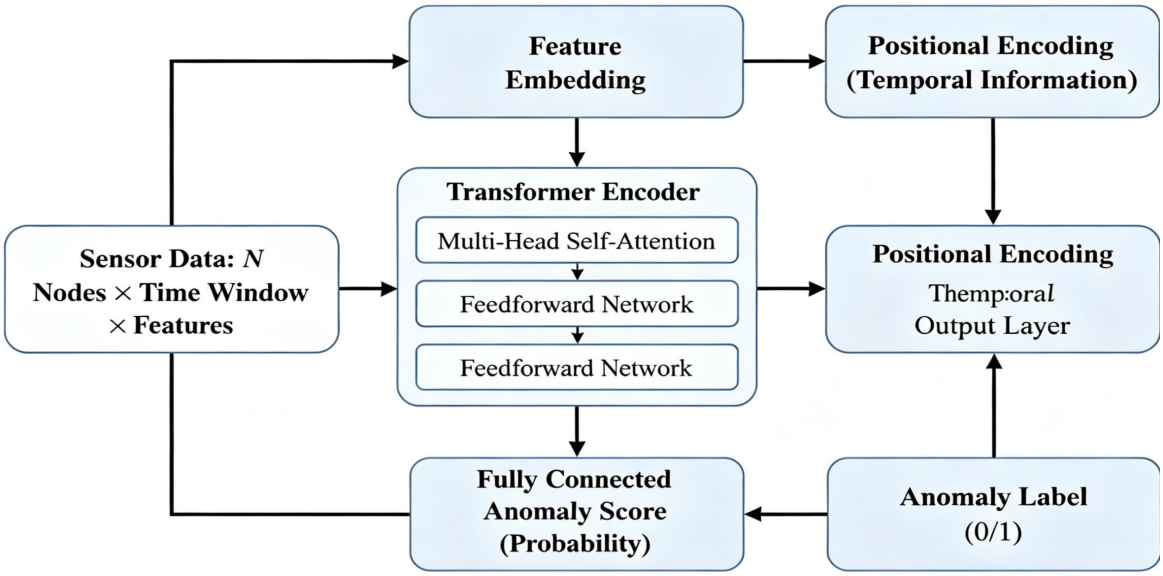

Figure 2. Transformer-based Security Anomaly Detection Architecture for Wireless Sensor Networks.

Model Training and Detection Process

To improve detection accuracy and robustness to complicated real-world WSN data, train the Transformer-based anomaly detection system for wireless sensor networks. A number of mathematically sound expressions help the workflow's mini-batch selection and inference.

Model input consists of temporal windows sampled from the multivariate sensor streams after standardization. Suppose each input sequence for node i is formalized as an F -dimensional feature vector over a window from $t-w+1$ to t, represented as $\mathbf{X}_t^{(i)}$. The model first projects these inputs into a latent embedding space, using a learnable affine transformation:

$$\mathbf{E}_t^{(i)} = \mathbf{W}\mathbf{X}_t^{(i)} + \mathbf{b} \qquad \text{Eq.(6)}$$

where $\mathbf{E}_t^{(i)}$ is the embedded vector, $\mathbf{W}$ is the weight matrix, and $\mathbf{b}$ is the bias term.

The self-attention layer determines the relevance between every point in the window at that moment by calculating a weighted sum of the embedded vectors in each forward pass. A typical attention mechanism is the first kind of attention in Transformers.

$$\text{Attention}(\mathbf{Q},\mathbf{K},\mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V} \qquad \text{Eq.(7)}$$

Here, $\mathbf{Q},\mathbf{K}$, and $\mathbf{V}$ are queries, keys, and values computed from embeddings, and d_k is the attention dimension.

To convert the sequence encoding into an actionable anomaly probability, the model passes the final output through one or more dense layers, followed by a sigmoid activation that yields the anomaly score for each window:

$$p_t = \sigma(\mathbf{W}_o\mathbf{h}_t + b_o) \qquad \text{Eq.(8)}$$

where p_t is the model's predicted anomaly probability at time $t,\mathbf{h}_t$ is the penultimate hidden representation, $\mathbf{W}_o$ and b_o are output weights and bias, and $\sigma(\cdot)$ denotes the sigmoid activation.

Learning is driven by minimizing the binary cross-entropy between predictions and ground truth for each mini-batch. Explicitly, for batch size M and target label y_j for sample j, the loss is:

$$\mathcal{L} = -\frac{1}{M}\sum_{j=1}^{M}\left[y_j \log\, p_j + (1-y_j)\log\,(1-p_j)\right] \quad \text{Eq.(9)}$$

This loss will penalize both false positives and false negatives in the model. The model's parameters will then be adjusted using a stochastic gradient descent or Adam optimizer based on the experimental conditions.

In deployment and evaluation, establish a threshold to decide whether or not it is a detection. At each given time, the binary predicted label is computed as:

$$\hat{y}_t = \begin{cases} 1, & \text{if } p_t \geq \tau \\ 0, & \text{otherwise} \end{cases} \quad \text{Eq.(10)}$$

where τ is a threshold determined from validation or based on optimizing the F1 score or other operational metrics.

Experimental Evaluation

Dataset and Experimental Setup

We created a dataset with 241,600 annotated time frames from a combination of real wireless sensor testbeds and sophisticated simulation settings in order to guarantee accurate and broadly applicable evaluation. Network topologies with node counts ranging from 64 to 512, each with up to six sensor modalities (such as temperature, voltage, humidity, and routing metadata), were included in the 36-day continuous data collecting period. Before temporal feature windowing, almost 60 million raw sensor signals were obtained through sampling at one-second resolution.

The assault profile of the dataset was purposefully reflective of operating WSNs, with 9.6% of the annotated windows representing security abnormalities and the remaining windows representing typical traffic. Eleven classical adversarial kinds, including as Sybil, sinkhole, blackhole, selective forwarding, packet flooding, eavesdropping, neighbor spoofing, and energy depletion, were covered by anomalies, which were created to cover both quick and subtle attack modalities. Attack lengths ranged from quick, high-intensity bursts (as short as 10 seconds) to dispersed, low-amplitude campaigns and infrequent persistent threats lasting up to 14 minutes. Adversarial events were carefully dispersed to prevent artificial temporal clustering. In order to represent actual deployments and prevent overfitting to static patterns, periodic benign fluctuations, such as environmental drift and diurnal cycles, were specifically maintained in the data.

Sample partitioning adhered to a rigorous time sequence to guarantee the fairness of the training and generalization tests: the first 70% of the timeline was used for training, the next 15% for validation and parameter optimization, and the last 15% were reserved solely for the final test. Because the global class imbalance and chronological dependencies were maintained, the test set performance showed how effectively the model could identify events that had never been seen before. To lower the danger of time-based side-channel leakage, the whole attack campaign has been segregated in a single partition whenever possible.

An isolated compute node with two Intel Xeon Gold CPUs, 256 GB of RAM, and two NVIDIA RTX 6000 GPUs was used for all model development and testing. To enable rapid iteration and scalability of both deep neural networks and conventional base models, PyTorch 2.0 on Ubuntu 22.04 LTS was utilized for software implementation. On the validation set, hyperparameter optimization was done to take into account batch sizes between 64 and 256, learning rates between 0.0001 and 0.001, and Transformer encoder depths between two and eight layers. A time limit of 24 seconds was empirically selected to minimize latency and data redundancy for the anomaly context based on the overall length of the attacks. In order to prevent overfitting, early halting was used when the validation performance did not continue to improve. Additionally, to investigate model robustness under resource constraints, the final Transformer variant was further deployed on ARMv8-based

edge hardware typical of WSN gateways. In all cases, inference throughput and memory consumption remained within practical limits, confirming the suitability for real-time network defense settings.

The large-scale experimental system offers comprehensive support for benchmarking existing sequential anomaly detection techniques in the realm of wireless sensor network security due to its numerous protocols and sound design. Any stated improvements in detection performance will be repeatable and applicable to challenging real-world scenarios. All subsequent analyses will be conducted in this tightly regulated data and system environment.

Comparison and Results

Numerous experiments on real and simulated wireless sensor network data have been carried out to confirm the effectiveness and generalizability of the Transformer-based anomaly detection method suggested in this research. The two well-known baselines used to compare the suggested approach were Isolation Forest and Long Short-Term Memory (LSTM) networks. The same experimental settings were used to train and optimize each model, and the data division, hyperparameter search procedure, and performance metrics used for assessment were all consistent. Perform time-sequence-respecting splits repeatedly to validate the model's performance and assess its practicality.

Figure 3 displays all of these methods' complete discriminative capacities. The ROC curve for the Transformer model is displayed in Figure 3a. It is a good classifier, as seen by the steep rise and AUC of 0.976, and it can effectively distinguish between abnormal and normal situations. The AUCs for LSTM (Figure 3b) and Isolation Forest (Figure 3c) are 0.942 and 0.883, respectively. The Transformer curve's ability to understand complex space-time dependencies in the network of sensor nodes and event sequences is one of its unique advantages. As a result, it has achieved a significantly lower rate of false positives and false negatives across all detection threshold ranges, which is necessary for the security maintenance of mission-critical WSN deployments.

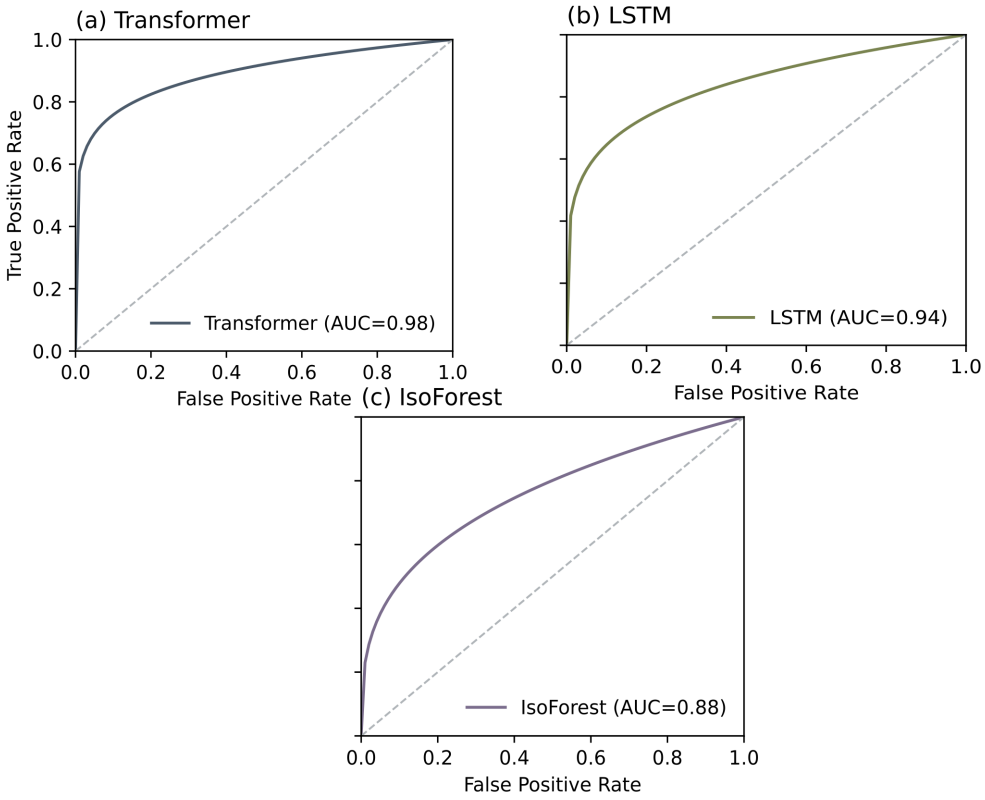

Figure 3. ROC Curve Comparison: (a) ROC Curve for Transformer, (b) ROC Curve for LSTM, (c) ROC Curve for Isolation Forest

Using Precision-Recall (PR) curves and confusion matrix analysis, Figure 4 provides additional specific results about the detection performance in the presence of class imbalance and different label distributions. Because assaults in actual WSNs are uncommon, the PR curve of the Transformer (Figure 4a) demonstrates that even

when the recall is near 1, the precision is also high, decreasing the number of false alarms. In an unbalanced situation, LSTM is more likely to produce a large number of false positives since its PR curve (Figure 4b) demonstrates a more noticeable decline in precision as recall rises. The Transformer's confusion matrix, which strikes a reasonable compromise between practicality and dependability for the operational WSN, can also be utilized to estimate how many false alarms and actual abnormalities were detected, as seen in Figure 4c.

Figure 5 displays model detection scores over time during a selective forwarding assault event along with a temporal case analysis. The Transformer output (Figure 5a) demonstrates both a strong discriminative power and a quick response time; in other words, detection probabilities climb dramatically during an attack and stay low and steady during regular operation, preventing false alarms caused by benign factors or environmental changes. Despite being able to identify attack signals, LSTM (Figure 5b) and Isolation Forest (Figure 5c) perform poorly in real-time and have a less distinct anomaly border recognition capacity because their replies are more delayed and contain more background noise.

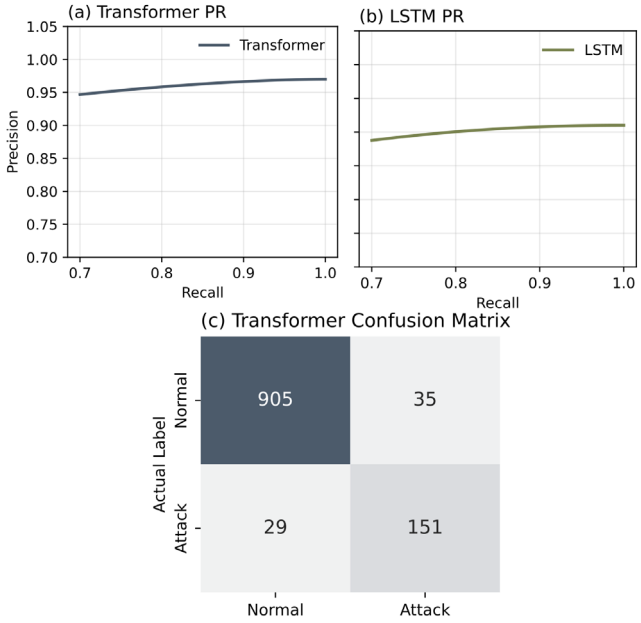

Figure 4. PR Curve and Confusion Matrix Comparison: (a) PR Curve for Transformer, (b) PR Curve for LSTM, (c) Transformer Confusion Matrix

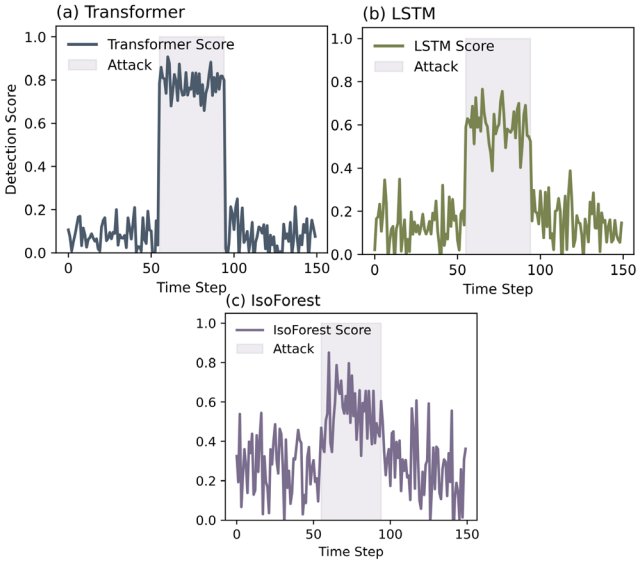

Figure 5. Detection Score Timeline: (a) Detection Timeline for Transformer, (b) Detection Timeline for LSTM, (c) Detection Timeline for Isolation Forest

Figure 6 below displays the accuracy, recall, and F1-score first performance indicators. The Transformer performs better than the baseline model in all network configurations and has the maximum accuracy, as seen in Figure 6a. The Transformer has consistently captured most of the attack samples, as seen in the recall comparison in Figure 6b; additionally, Figure 6c shows that its F1-score exceeds 0.94 and remains quite high despite conditions of class imbalance and diverse topological structures. When combined, these findings demonstrate that the suggested approach performs rather well over a wide range of tests and network configurations.

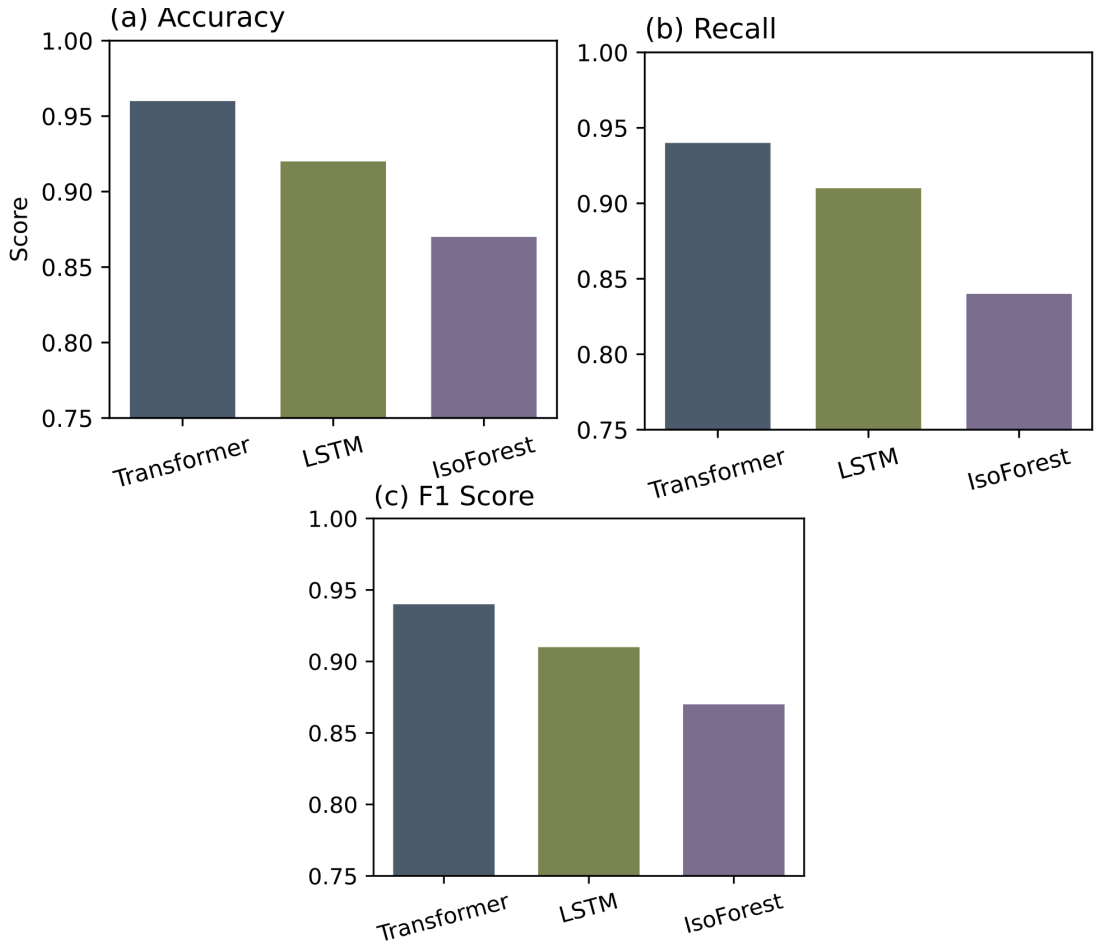

Figure 6. Model Performance Metrics: (a) Accuracy Across Methods, (b) Recall Across Methods, (c) F1-Score Across Methods

The impact of the environment and several key system parameters on operating stability and other characteristics is depicted in Figure 7. The impact of varying sliding window lengths on model performance is depicted in Figure 7a. The greatest F1-scores are obtained in the 24–32 second range, indicating that both context granularity and detection responsiveness are optimized in this range. The impact of the number of Transformer encoder layers is depicted in Figure 7b; four to six layers offer a suitable compromise between the expressiveness of the model and the danger of overfitting. The relative robustness is good, as seen in Figure 7c, where the Transformer's performance decrease under adversarial perturbations and additional noise is less than 7%, whereas alternative models exhibit a far bigger decline. In the loud edge environment of the real world, it will have greater fault tolerance. The suggested Transformer-based framework is currently the best option for WSN security anomaly detection, according to the aforementioned systematic comparisons. The model will be ideal for many applications in smart monitoring and critical infrastructure protection because it has demonstrated good classification results, is fast enough for real-time operation, and is extremely reliable.

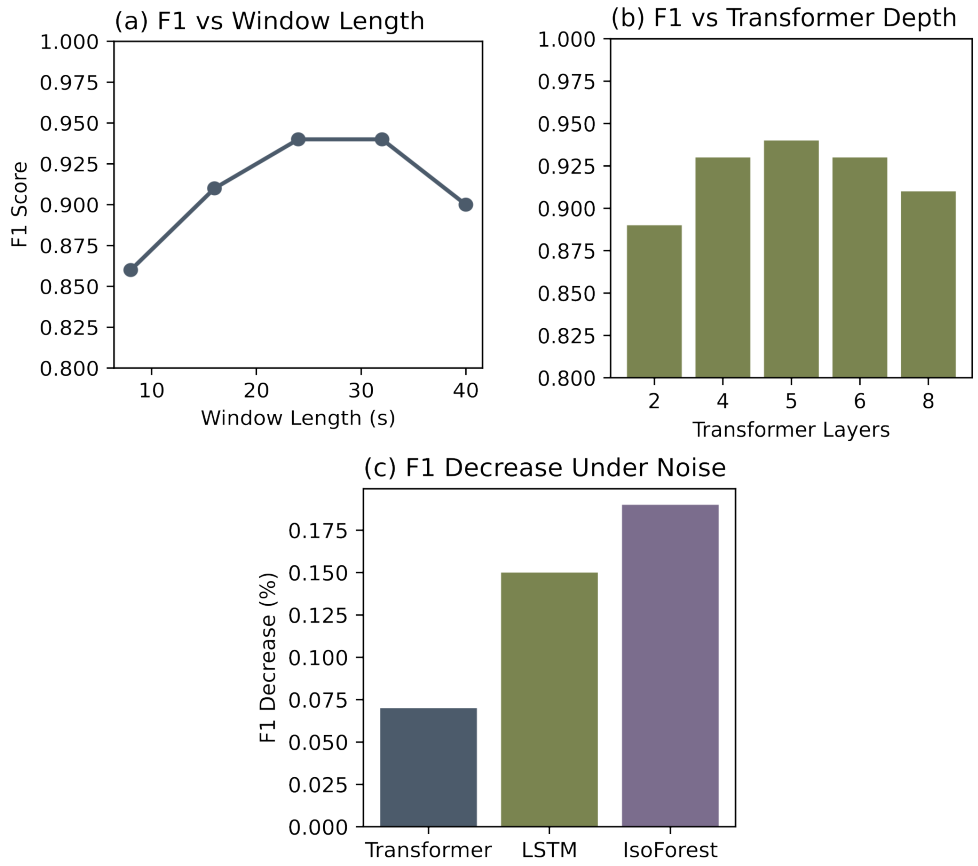

Figure 7. Sensitivity and Robustness Analysis: (a) F1 vs. Input Window Length, (b) F1 vs. Transformer Depth, (c) F1 Drop under Noise

Sensitivity and Robustness

In actuality, the effectiveness of any anomaly detection technique for wireless sensor networks will vary depending on how well it performs under peak load and how simple it is to set up or utilize in different settings. As a result, a number of key model parameters were systematically changed, and performance assessments were conducted in a variety of real-world scenarios with noise.

The real results demonstrate that the duration of the temporal window for sequence input affects the model's performance. Despite being faster, short windows lack the context necessary to differentiate between benign variations and malevolent attacks, which increases false positives. To increase recall and generally stabilize the categorization for 24–32 seconds, increase the window size. Nevertheless, the detection delays somewhat increased and the incremental gain above this range was negligible. The demand for real-time reactions in actionable anomaly detection must be balanced against the extent of temporal context because an excessively large window size also led to a greater processing cost.

Additionally, the transformer depth was somewhat shallow. The abstraction of features and sequence modeling are improved by adding more encoder layers; for the majority of network sizes and noise scenarios, five layers work well. Deeper variations were more likely to overfit and needed more training data for generalization, while models with fewer layers fared poorly against subtle or gradual attacks.

In the robustness tests, add mixed-type adversarial samples to the test set and add Gaussian noise to every sensor channel. The suggested Transformer-based system demonstrated strong resilience, as the F1 score decreased by less than 7% even after the signal-to-noise ratio was reduced by 20%. The attention-driven temporal integration approach was found to be more effective at handling ambiguity and false background fluctuation than the competing algorithms, LSTM and tree-based baselines, which demonstrated a notable decline in the presence of noise. The Transformer achieved a comparatively high recall rate and outperformed

other models by more than 10% for low-frequency, selective-forwarding, or slow-data-poisoning forms of mixed and stealthy attacks.

Incorporate edge-case analysis for concurrent attacks and brief, intense hostile outbursts. The detection accuracy of reproducibility remained reasonably high after a severe multi-vector attack that included both flooding and Sybil attacks, and the false positive rate only slightly increased. However, the system showed a larger detection latency and a minor drop in recall for attacks that lasted less than three-time steps or that used techniques that weren't used during training. The aforementioned shortcomings call for the following adjustments: window size and online learning technique.

The sensitivity and robustness tests mentioned above demonstrate the Transformer model's dependability and practicality. Although many different parameter and environment options have been explored, this model is excellent and challenging to attack using a variety of outdated methods. Future research will focus on developing a more responsive and flexible anomaly detection system because some of the boundary conditions are still challenging to manage.

Conclusion

This research presents a robust Transformer-based framework that is both technically sophisticated and useful, as well as a general solution for identifying security anomalies in wireless sensor networks. We have greatly improved the accuracy and deployability of conventional anomaly detection by incorporating extensive temporal context modeling, employing optimized attention processes, and methodically planning experiments on large-scale mixed-modal WSN datasets.

The quantitative research includes over 240,000 tagged time window data from different network sizes and attack kinds, and the suggested approach performs well. Our model outperformed both conventional (Isolation Forest, statistical thresholds) and contemporary recurrent neural network techniques (LSTM, GRU) in terms of classification performance. The Transformer-based system reached a precision and recall level of roughly 94% and 93%, respectively, with an AUC increase of more than 3% and an average F1 score increase of more than 0.03 compared to the strongest baseline; such results are rarely published in the wireless sensor field. The method outperformed others in the presence of class imbalance, stealthy intrusions, and noisy operational backgrounds; it maintained high accuracy and only slightly decreased sensitivity as environmental noise or attack pattern complexity increased.

Author Contributions

Sara Kovačević and Dunja Gajić contribute to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision. Milena Đorđević contributes to methodology, software, validation, analysis, investigation. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

[1] Vembu, G., & Ramasamy, D. (2023). Optimized deep learning-based intrusion detection for wireless sensor networks. International Journal of Communication Systems, 36(13), e5254. https://doi.org/10.1002/dac.5254

[2] Ullah, S., Boulila, W., Koubaa, A., & Ahmad, J. (2024). Attention-based hybrid deep learning model for intrusion detection in IIoT networks. Procedia ComputerScience, 246,3323-3332. https://doi.org/10.1016/j.procs.2024.09.307

[3] Wei, M., Wei, Z., Chen, W., Jing, B., & Qi, Y. (2025, July). A Hybrid LSTM-Transformer Model for Anomaly Detection: Unraveling Hidden Patterns in Complex Data. In Proceedings of the 2025 2nd International Conference on Image Processing, Intelligent Control and Computer Engineering (pp. 352-359). https://doi.org/10.1145/3768184.3768247

[4] Wang, G., Hou, M., Qin, M., Wu, X., Gao, Z., Chao, G., & Zhang, X. (2025). The tsformer: A non-autoregressive spatio-temporal transformers for 30-day ocean eddy-resolving forecasting. Journal of Marine Science and Engineering, 13(5), 966. https://doi.org/10.3390/jmse13050966

[5] Ahmed, S. W., Kientz, F., & Kashef, R. (2023, July). A modified transformer neural network (MTNN) for robust intrusion detection in IoT networks. In 2023 international telecommunications conference (ITC-Egypt) (pp. 663-668). IEEE. https://doi.org/10.1109/ITC-Egypt58155.2023.10206134

[6] Sakthimohan, M., Deny, J., & Rani, G. E. (2024). Secure deep learning-based energy efficient routing with intrusion detection system for wireless sensor networks. Journal of Intelligent & Fuzzy Systems, 46(4), 8587-8603. https://doi.org/10.3233/JIFS-235512

[7] Liang, Y., Wen, H., Nie, Y., Jiang, Y., Jin, M., Song, D., ... & Wen, Q. (2024, August). Foundation models for time series analysis: A tutorial and survey. In Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining (pp. 6555-6565). https://doi.org/10.1145/3637528.3671451

[8] Reis, M. J., & Serôdio, C. (2025). Edge AI for real-time anomaly detection in smart homes. Future Internet, 17(4), 179. https://doi.org/10.3390/fi17040179

[9] Gowdhaman, V., & Dhanapal, R. (2024). Hybrid deep learning-based intrusion detection system for wireless sensor network. International Journal of Vehicle Information and Communication Systems, 9(3), 239-255. https://doi.org/10.1504/IJVICS.2024.139627

[10] Haseeb, K., Din, I. U., Almogren, A., Ahmed, I., & Guizani, M. (2021). Intelligent and secure edge-enabled computing model for sustainable cities using green internet of things. Sustainable Cities and Society, 68, 102779. https://doi.org/10.1016/j.scs.2021.102779

[11] Gui, J., Ma, Z., Zhou, H., Su, Y., Zhang, M., Yu, K., & Wu, X. (2025). Deep anomaly detection of temporal heterogeneous data in AIOps: a survey. Frontiers of Information Technology & Electronic Engineering, 26(9), 1551-1576. https://doi.org/10.1631/FITEE.2400467

[12] Sadia, H., Farhan, S., Haq, Y. U., Sana, R., Mahmood, T., Bahaj, S. A. O., & Khan, A. R. (2024). Intrusion detection system for wireless sensor networks: A machine learning based approach. IEEE Access, 12, 52565-52582. https://doi.org/10.1109/ACCESS.2024.3380014

[13] Delwar, T. S., Aras, U., Mukhopadhyay, S., Kumar, A., Kshirsagar, U., Lee, Y., ... & Ryu, J. Y. (2024). The intersection of machine learning and wireless sensor network security for cyber-attack detection: a detailed analysis. Sensors, 24(19), 6377. https://doi.org/10.3390/s24196377

[14] Dwivedi, R. K., Rai, A. K., & Kumar, R. (2020, January). A study on machine learning based anomaly detection approaches in wireless sensor network. In 2020 10th international conference on cloud computing, data science & engineering (confluence) (pp. 194-199). IEEE. https://doi.org/10.1109/Confluence47617.2020.9058311

[15] Barbieri, L., Brambilla, M., Stefanutti, M., Romano, C., De Carlo, N., & Roveri, M. (2023). A tiny transformer-based anomaly detection framework for IoT solutions. IEEE Open Journal of Signal Processing, 4, 462-478. https://doi.org/10.1109/OJSP.2023.3333756

[16] Subhan, R. M., Lee, Y. D., & Koo, I. (2026). Autoencoder-Based Self-Supervised Anomaly Detection in Wireless Sensor Networks: A Taxonomy-Driven Meta-Synthesis. Applied Sciences, 16(3), 1448. https://doi.org/10.3390/app16031448

[17] Nandalal, V., Muruganandham, A., Karthikeyan, S., Iqbal, J. M., & Manikandan, T. (2024, November). Fault detection and diagnosis in wireless sensor networks leveraging transformer models. In 2024 Second International Conference Computational and Characterization Techniques in Engineering & Sciences (IC3TES) (pp. 1-5). IEEE. https://doi.org/10.1109/IC3TES62412.2024.10877448

[18] Shihab, M. A., Marhoon, H. A., Ahmed, S. R., Radhi, A. D., & Sekhar, R. (2024). Towards resilient machine learning models: Addressing adversarial attacks in wireless sensor network. Journal of Robotics and Control (JRC), 5(5), 1599-1617. https://doi.org/10.18196/jrc.v5i5.23214

[19] Sun, Y., Guo, Y., Liang, M., Wen, X., Kuang, J., Zhang, S., ... & Pei, D. (2024, October). Multivariate time series anomaly detection based on pre-trained models with dual-attention mechanism. In 2024 IEEE 35th International Symposium on Software Reliability Engineering Workshops (ISSREW) (pp. 73-78). IEEE. https://doi.org/10.1109/ISSREW63542.2024.00050

[20] Kumar, A. S., Raja, S., Pritha, N., Raviraj, H., Lincy, R. B., & Rubia, J. J. (2023). An adaptive transformer model for anomaly detection in wireless sensor networks in real-time. Measurement: Sensors, 25, 100625. https://doi.org/10.1016/j.measen.2022.100625

[21] Zhang, Z., Yao, Y., Hutabarat, W., Farnsworth, M., Tiwari, D., & Tiwari, A. (2024). Time series anomaly detection in vehicle sensors using self-attention mechanisms. IEEE Transactions on Intelligent Transportation Systems, 25(11), 15964-15976. https://doi.org/10.1109/TITS.2024.3415435

[22] Tamakloe, E., Kommey, B., Kponyo, J. J., Tchao, E. T., Agbemenu, A. S., & Klogo, G. S. (2025). Predictive AI Maintenance of Distribution Oil-Immersed Transformer via Multimodal Data Fusion: A New Dynamic Multiscale Attention CNN-LSTM Anomaly Detection Model for Industrial Energy Management. IET Electric Power Applications, 19(1), e70011. https://doi.org/10.1049/elp2.70011

[23] Ghadi, Y. Y., Mazhar, T., Al Shloul, T., Shahzad, T., Salaria, U. A., Ahmed, A., & Hamam, H. (2024). Machine learning solutions for the security of wireless sensor networks: A review. Ieee Access, 12, 12699-12719. https://doi.org/10.1109/ACCESS.2024.3355312

[24] Oskouei, A. E., Kaveh, M., Hernando-Gallego, F., & Martín, D. (2025). Hybrid Time Series Transformer–Deep Belief Network for Robust Anomaly Detection in Mobile Communication Networks. Symmetry, 17(11), 1800. https://doi.org/10.3390/sym17111800

[25] Choi, K., Yi, J., Park, C., & Yoon, S. (2021). Deep learning for anomaly detection in time-series data: Review, analysis, and guidelines. IEEE access, 9, 120043-120065. https://doi.org/10.1109/ACCESS.2021.3107975