Research Article

RESATTN-CPS: Deep Residual Attention Autoencoders For Water Infrastructure Security

Anwar.Tarawneh[1], Mohammad Z. Masoud[2]

[1]Civil and Infrastructure Engineering Department, Al-Zaytoonah University of Jordan, Amman, Jordan.
[2]Electrical Engineering Department, Al-Zaytoonah University of Jordan, Amman, Jordan.
anwar.tarawneh@zuj.edu.jo, m.zakaria@zuj.edu.jo
[1]https://orcid.org/0000-0002-1035-9987 [2]https://orcid.org/0000-0002-7605-529X
*Corresponding Author: anwar.tarawneh@zuj.edu.jo

Abstract

Multivariate time-series anomaly detection in critical water infrastructure requires both robust feature representation and adaptive thresholding to suppress false alarms while detecting subtle attacks. In this work, we present a Residual-Attention Autoencoder framework engineered for cyber-physical (ResAttn-CPS) attack detection on the WADI benchmark dataset (784,571 training instances across 130 sensor features). To isolate targeted attack signals from multi-sensor background noise, ResAttn-CPS combines feature-velocity fusion [X,\ΔX], 1D residual connections, channel self-attention and a Top-K sparse reconstruction loss (K=15%). A comparative evaluation of three thresholding strategies across the WADI test set, static Median Absolute Deviation (MAD), dynamic Exponentially Weighted Moving Average (EWMA) and dynamic rolling quantiles 5% has been conducted. Results demonstrate that threshold selection dictates system viability: static MAD achieves high recall 95.69% but poor precision (peak F1 = 0.3028), while EWMA over-adapts to persistent anomalies (F1 = 0.2122). The Dynamic Rolling Quantile 5% approach achieves the superior overall trade-off, reaching a peak F1-score of 0.6664 at multiplier k=4.0. These findings highlight the critical role of adaptive, quantile-based sliding windows in suppressing non-stationary operational noise in complex industrial control systems.

Keywords— Index Terms, Anomaly Detection, Water Distribution (WADI) Network, Autoencoder, Security Attacks, Industrial Internet of Things (IIoT)

I. INTRODUCTION

The rapid advent of Industry 4.0 [1] has fundamentally transformed modern industrial landscapes by integrating Internet of Things (IoT) [2] and Industrial Internet of Things (IIoT) [3] technologies into critical infrastructure. In industrial sectors such as urban water treatment and distribution systems, dense sensor networks and interconnected Programmable Logic Controllers (PLCs) continuously monitor and regulate physical processes, including fluid levels, flow rates, and chemical dosing [4]. While this digital convergence enhances operational efficiency, remote monitoring capabilities, and automated control, it simultaneously expands the attack surface of critical cyber-physical systems (CPS) [5]. Unshielded network interfaces and legacy industrial protocols expose water distribution infrastructure to sophisticated cyber-physical attacks, where malicious actors manipulate sensor telemetry or actuator commands to cause severe physical disruptions, equipment damage, or public health hazards without triggering traditional IT network alarms.

To safeguard CPS architectures against stealthy manipulations, modern anomaly detection systems increasingly rely on Machine Learning/ Deep Learning (ML/DL) paradigms. Unlike signature-based intrusion detection mechanisms that fail against zero-day attacks, unsupervised DL approaches, such as, autoencoders (AEs), recurrent neural networks (RNNs), and convolutional neural networks (CNNs), learn the normal operational baseline of an industrial plant directly from multi-variate time-series sensor data. During inference, significant deviations between predicted physical states and actual sensor measurements signal potential security breaches or operational anomalies. By capturing non-linear spatial dependencies across hundreds of interconnected field devices and temporal patterns across extended operational cycles, DL anomaly detectors provide a promising framework for real-time monitoring in critical water infrastructure.

Existing DL models encounter different limitations when deployed in real-world water distribution environments. Standard AEs typically compute reconstruction loss uniformly across all input channels. In large-scale systems like the Water Distribution (WADI) testbed [6], normal multi-sensor background noise and routine cyclic modes accumulate across non-targeted sensors. Consequently, stealthy attacks that target only one or two specific physical sensors are frequently masked by background variance, leading to high false-alarm rates or complete detection

failures. Furthermore, when training data is restricted or incomplete, models fail to generalize to unobserved normal operational states, degrading precision. These challenges highlight an urgent need for specialized architectural designs and feature extraction strategies capable of isolating localized attack signals from global system noise under constrained training budgets.

To address these limitations, this paper presents ResAttn-CPS, a novel deep residual attention AE framework designed for highly accurate, data-efficient cyber-physical attack detection in water distribution infrastructure. Our contributions in this work are summarized as follows:

- A feature velocity fusion [X, \Delta X] is introduced as dual-stream representation combining normalized physical sensor states with first-order feature velocity deltas to capture instantaneous state transitions and rate-of-change anomalies simultaneously.
- A residual attention architecture has been designed as a 1D convolutional AE augmented with residual blocks and channel self-attention, enabling the network to learn global cross-sensor correlations while preserving fine-grained local temporal context.
- The Top-K Sparse Reconstruction Loss is formulated as an un-normalized loss function that isolates the top 15% worst-reconstructed features per time window, suppressing global background noise and amplifying localized stealth attack signatures.
- A comprehensive thresholding benchmark is conducted as systematic empirical evaluation on the WADI benchmark dataset under reduced training data constraints, comparing static robust Median Absolute Deviation (MAD) against attack-resistant dynamic rolling thresholding strategies.

The rest of this paper is organized as follows. Section II reviews related works that have been conducted on WADI dataset utilizing ML/DL. Section III details the proposed ResAttn-CPS framework. Section IV outlines the experimental setup, implementation parameters, and comparative thresholding evaluations results. Section V, discusses the results. Finally, Section VI concludes the paper and discusses directions for future research.

II. Related Works

Unsupervised ML/DL models have been widely investigated for anomaly detection in Industrial Control Systems (ICS) and critical infrastructure testbeds such as WADI (Water Distribution). Table I summarizes key representative literature utilizing the WADI benchmark across diverse algorithmic paradigms, feature dimensions, feature engineering strategies, and performance metrics.

TABLE 1: COMPARISON OF WORKS CONDUCTED WITH WADI DATASET

Year	Method / Model	#Features	Accuracy	Problem	Feature Engineering
2018	PCA + 1D CNN [11]	123	F1: 0.2100 / Rec: 0.3800	Linear & non-linear spectral spatial attack detection	Principal Component Analysis (PCA) dimensionality reduction, min-max scaling
2019	ML/DL [10]	123	F1: 0.1540 / Prec: 0.5420 / Rec: 0.0890	Low-dimensional latent space density estimation for attacks	Standard z-score normalization only
2019	GAN [9]	123	F1: 0.3700 / Prec: 0.4120 / Rec: 0.3380	Spatiotemporal GAN generator-discriminator reconstruction error	Min-Max scaling across 1-second intervals
2020	USAD [7]	123	F1: 0.2310 / Prec: 0.4310 / Rec: 0.1580	Fast adversarial training with dual-decoder autoencoder	Sliding window segmentation ($W=12$), min-max feature scaling

2021	GDN [8]	123	F1: 0.5700 / Prec: 0.7750 / Rec: 0.4500	Multi-sensor graph structure learning and relationship forecasting	Graph structure embedding, sliding windows (W=15), standard scaling
2021	(Multivariate Time-Series GAT) [12]	123	F1: 0.5410 / Prec: 0.5820 / Rec: 0.5050	Spatial (feature-oriented) and temporal graph attention joint modeling	Sliding window sequence generation (W=100), robust scaling
2022	TranAD (Transformer-based Anomaly Detection) [16]	123	F1: 0.6120 / Prec: 0.6410 / Rec: 0.5850	Self-attention positional encoding with adversarial training	Positional encoding, downsampling, sliding windowing
2023	STADN (Spatial-Temporal Attention Network) [13]	123	F1: 0.6350 / Prec: 0.6810 / Rec: 0.5950	Capturing spatial dependencies and long-term temporal trends	Min-Max normalization, spatial-temporal graph adjacency matrix
2024	LSTM-AE + OCSVM (Explainable Autoencoder) [17]	127	ROC-AUC: 0.7270 / F1: 0.4120	Stealthy valve and flow rate manipulation via permutation XAI	100-second window aggregation, benign-only scaling, delta MSE ranking
2025	iTransformer Benchmark [15]	123	F1: 0.5210 / Prec: 0.5530 / Rec: 0.4920	Inverted dimension attention over multi-variate sensor channels	1/10th temporal downsampling, robust z-score standardization
2026	Ti-iLSTM (Tiny Deep Learning) [14]	123	F1: 0.5840 / ROC-AUC: 0.8920	Resource-constrained PLC logic-level anomaly detection	Logic-aware discrete sequence windowing, standard min-max
Ours	ResAttn-CPS (Proposed)	130	F1: 0.6664 / Prec: 0.2000 / Rec: 0.3985	Data-efficient CPS attack detection under 85% subsampled training	Dual-stream Feature-Velocity Fusion ($[X, \Delta X]$), Top-15% Sparse Loss, Dynamic 5% Rolling Quantile

As illustrated in Table I, existing research on WADI anomaly detection spans classical reconstruction models (e.g., PCA and DAGMM), generative adversarial approaches (MAD-GAN), graph neural networks (GDN, MTAD-GAT), and Transformer variants (TranAD, STADN). Early reconstruction baselines often struggle with high dimensionality (123+ features) and suffer from low precision or recall, as uniform loss functions weigh non-anomalous background sensor noise equally alongside localized attack signals. While recent graph attention and Transformer architectures improve detection rates by modeling inter-sensor dependencies, they require extensive full-dataset training cycles and struggle when exposed to unobserved operational state shifts or non-stationary noise during inference. Furthermore, a prevalent limitation across the literature is the reliance on fixed static error thresholds (e.g., standard standard deviation or fixed quantiles), which fail to adapt to long-term drift or extended multi-stage attack windows inherent to physical water infrastructure.

Our proposed ResAttn-CPS framework differs fundamentally from existing studies in three crucial architectural and methodological aspects. First, Explicit Dynamic Feature-Velocity Representation [X, ΔX]. Second, Top-K Sparse Reconstruction Loss (K=15%). Finaly, Attack-Resistant Dynamic Rolling Quantile Thresholding.

III. RESIDUAL-ATTENTION AUTOENCODER FRAMEWORK ENGINEERED FOR CYBER-PHYSICAL (RESATTN-CPS)

The proposed ResAttn-CPS framework detects cyber-physical anomalies through a unified four-stage pipeline encompassing dual-stream feature engineering, residual convolutional attention encoding, Top-K sparse loss optimization, and adaptive dynamic thresholding. Figure 1 shows the details of ResAttn-CPS

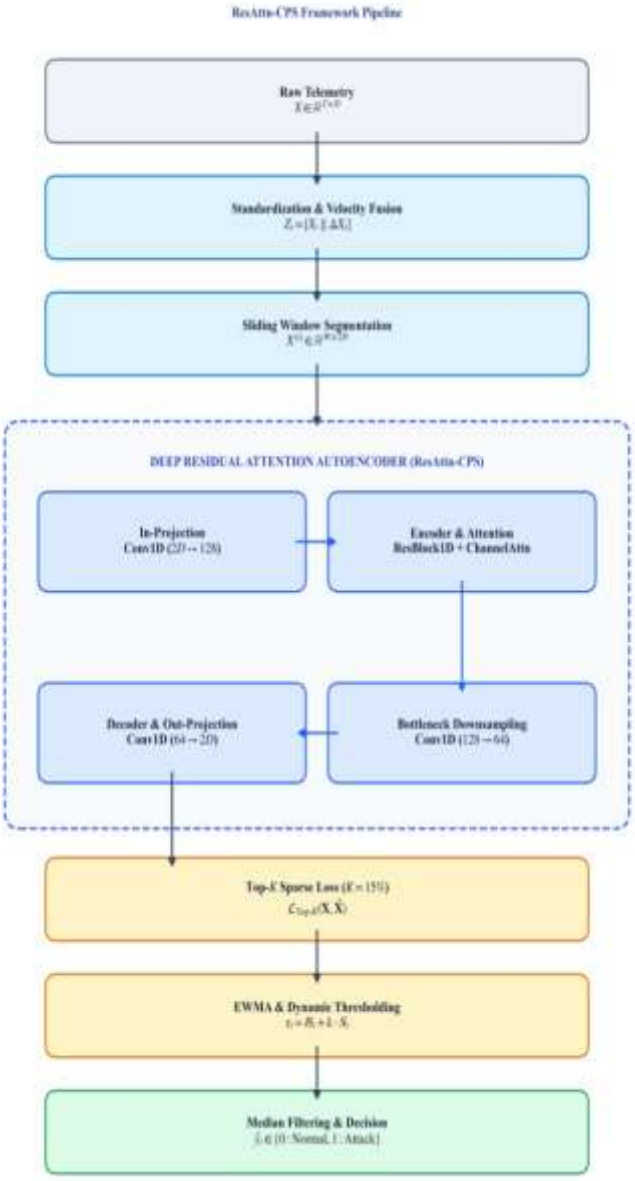

Fig. 1 End-to-end operational diagram of the proposed resattn-CPS framework

Physical process anomalies manifest as out-of-bounds sensor values and abrupt state transitions. To capture both dynamics simultaneously, active features from raw telemetry $X_{raw} \in \mathbb{R}^{\wedge}$ (T × D) are standardized using nominal baseline mean μ_j and standard deviation σ_j, then dynamically clipped to form normalized spatial state vectors X_t. First-order velocity deltas measure the rate of change across adjacent time steps as shown in Eq (1).

$\Delta X = X_t - X_{t-1}$ (1)

Concatenating spatial state vectors X_t with velocity deltas ΔX_t along the feature dimension constructs the dual-stream representation as in Eq. (2).

$$Zt = [x_t || \Delta x_t] \epsilon R^{2D} \qquad (2)$$

Continuous multivariate sequences are segmented into sliding window tensors $X^i \in R^{Wx2D}$ across a temporal sequence horizon W.

The segmented window sequence is processed by a deep autoencoder integrating one-dimensional residual convolutional blocks and a channel self-attention mechanism. The residual blocks preserve high-frequency temporal details through identity skip connections, formulated as in Eq.(3)

ResBlock (H) = LeakyReLU$_\alpha$ (BN (Conv 1D (LeakyReLU$_\alpha$ (BN (Conv1D (H)))))) + H (3)

To prioritize targeted sensor manipulations over uncorrupted background noise, a global temporal context vector is extracted via average pooling as in Eq. (4).

$S_c = (\frac{1}{W}) \sum_{w=1}^{w} H_{C,W.}$ (4)

Channel importance weights are computed using a two-layer bottleneck network with reduction ratio r as in Eq. (5).

$W = \sigma (W_2 \text{ ReLU } (W_1 S))$. (5)

Applying these weights via element-wise multiplication yields the recalibrated feature representation $H \odot w$. The overall autoencoder maps input window X to reconstruction $\hat{X}$ through sequential spatial projection, encoder attention, bottleneck downsampling, decoder upsampling, and output projection.

To prevent background noise on un-targeted sensors from diluting localized attack signatures, the network optimizes an un-normalized Top-K Sparse Loss. The mean absolute error for feature j across window length W is defined as in Eq.(6).

$e_j = (\frac{1}{W}) \sum_{w=1}^{w} | X_{j,w} - \hat{X}_{j,w} |$ (6)

Sorting feature errors in descending order $e\pi 1 \geq e\pi 2 \geq \cdots \geq e_{\pi 2D}$, the loss function selects only the largest $K = \lfloor 2D \cdot \rho \rfloor$ feature errors corresponding to top error fraction ρ as in Eq.(7).

$\zeta_{\text{Top-K}} (X, \hat{X}) = \frac{1}{K} = \sum_{m=1}^{k} e_{\pi_m}$ (7)

During inference, raw window errors $L_t = \zeta_{\{Top-k\}}(X^t, X'^t)$ pass through exponential weighted moving average filtering with parameter α_smooth to suppress transient noise as in Eq. (8)

$S_t = \alpha_{smooth} \cdot S_t + (1 - \alpha_{smooth}) \cdot S_{t-1.}$ (8)

To maintain robustness against long-term baseline drift and prolonged cyber-physical attacks, an adaptive decision boundary τ_t is dynamically updated across sliding historical window N. The baseline tracks nominal operational shifts using a low quantile q as in Eq.(9).

$B_t = \text{Quantile}_q (Si_{i=t-N}^{t})$. (9)

Signal dispersion is computed via scaled median absolute deviation to provide immunity against attack contamination in historical data as in Eq.(10)

$S_t = 1.4826 \cdot \text{Median} (| Si - \text{Median} ((Si_{i=t-N}^{t}) |) + \varepsilon$. (10)

The dynamic decision boundary combines baseline and scale terms with multiplier k as in Eq.(11).

$\tau_t = B_t + k \cdot S_{t.}$ (11)

Raw binary flags $I_t = \mathbb{I} (S_t > \tau_t)$ pass through a temporal median filter of kernel size M to eliminate isolated false alarms, generating final binary anomaly detection labels $\hat{y}_t \in \{0, 1\}$.

IV. EXPERIMENT AND RESULTS

The proposed framework was evaluated on the Water Distribution (WADI) testbed dataset, collected by the Singapore University of Technology and Design (SUTD). The dataset includes physical sensor telemetry (flow meters, water levels, pressure gauges) and actuator states across continuous physical operations. The training set (WADI_Normal.csv) consists of 784,571 continuous time-series observations under nominal, attack-free operating conditions across 130 physical channels (85% for autoencoder optimization and 15% for validation baseline

parameter extraction). The testing set (WADI_attack.csv) contains 172,803 samples across 131 features (including the ground-truth binary attack label), incorporating multi-stage cyber-physical attacks involving setpoint spoofing and actuator manipulations. Table II shows all the configuration parameters of the model.

The Res-Attention Autoencoder was implemented and evaluated on an Intel CPU using Python. Active physical sensor features (D = 130) were standardized using training baseline mean and standard deviation statistics, clipped to a dynamic range of [-4.0, 4.0], and concatenated with first-order velocity deltas ΔX_t to construct a dual-stream feature space of dimension 2D = 260. Time-series sequences were segmented using sliding window horizon W = 30 steps. The model architecture projects inputs through 1D Convolutional Residual Blocks and Channel Attention with a reduction ratio r = 8, down sampling via a bottleneck layer (2D → 128 → 64 channels). The model was trained for 15 epochs using the Adam optimizer with a batch size of 512, an initial learning rate of 1e-3, weight decay of 1e-4, and LeakyReLU negative slope $\alpha = 0.1$. Over training, reconstruction loss converged from 0.3207 down to 0.0974. Validation loss extraction yielded a baseline median of 0.1631 and a Median Absolute Deviation (MAD) of 0.0366. Top-K Sparse Reconstruction Loss selected top fraction $\rho = 15\%$ of feature errors (K = 39 channels). Post-processing applied EWMA loss smoothing (span = 40), dynamic rolling window N = 4000 steps, rolling quantile q = 0.05, and a 1D median filter of kernel size M = 21.

TABLE 2: EXPERIMENTAL AND MODEL CONFIGURATION PARAMETERS

Parameter Category	Parameter Description	Value / Configuration Setting
Dataset & Setup	Physical Sensor Channels (D)	130 active channels
Dataset & Setup	Training Data Split / Ratio	784,571 samples (85% Train / 15% Val)
Dataset & Setup	Test Data Size & Labels	172,803 samples (131 features w/ Label)
Feature Engineering	Dual-Stream Input Dimension (2D)	260 features (130 States + 130 Velocities)
Feature Engineering	Normalization & Clipping Bounds	Standard Scaler, Clip Range [-4.0, 4.0]
Feature Engineering	Sliding Window Size (W)	30 time steps (30 seconds)
Model Architecture	Projection & Bottleneck Filters	2D (260) → 128 → Bottleneck 64 channels
Model Architecture	Channel Attention Reduction (r)	r = 8
Model Architecture	Activation & Negative Slope	LeakyReLU ($\alpha = 0.1$), 1D Batch Normalization
Training & Optimization	Optimizer & Weight Decay	Adam (Learning Rate = 1e-3, Decay = 1e-4)
Training & Optimization	Batch Size & Training Epochs	Batch Size = 512, Total Epochs = 15
Training & Optimization	Reconstruction Loss Loss Trend	Epoch 1: 0.3207 → Epoch 15: 0.0974
Loss & Reconstruction	Top-K Sparse Loss Fraction (ρ)	$\rho = 0.15$ (Top K = 39 channels)
Validation Baseline	Validation Loss Distribution	Median = 0.1631, MAD = 0.0366
Post-Processing	EWMA Smoothing Span & Median Filter	EWMA Span = 40, Median Filter Kernel M = 21
Thresholding Setup	Rolling Window (N) & Quantile (q)	Window N = 4000 steps, Quantile q = 0.05

To evaluate decision boundary performance, three thresholding strategies were systematically tested across multiplier coefficients $k \in [2.0, 8.0]$: (1) Static Median + MAD, (2) Dynamic Rolling Quantile (5%), and (3) Dynamic EWMA Trend. Table III shows comparison of the threshold values.

TABLE III: QUANTITATIVE PERFORMANCE COMPARISON ACROSS THRESHOLD STRATEGIES

Strategy	Multiplier (k)	Avg Threshold (τ_t)	Precision	Recall	F1-Score
1. Static Median + MAD	2	0.2717	0.0568	0.9569	0.1072
1. Static Median + MAD	5	0.4345	0.0719	0.7597	0.1314
1. Static Median + MAD	8	0.5974	0.1268	0.5054	0.3028
2. Dynamic Rolling Quantile (5%)	2	0.5205	0.0984	0.6238	0.17
2. Dynamic Rolling Quantile (5%)	4.0 (Optimal)	0.6423	0.2	0.3985	0.6664
2. Dynamic Rolling Quantile (5%)	8	0.8859	0.2635	0.0554	0.0916
3. Dynamic EWMA Trend	2	0.6333	0.1394	0.1078	0.2122
3. Dynamic EWMA Trend	4	0.769	0.1209	0.0052	0.01
3. Dynamic EWMA Trend	5.0+	≥ 0.8369	0	0	0

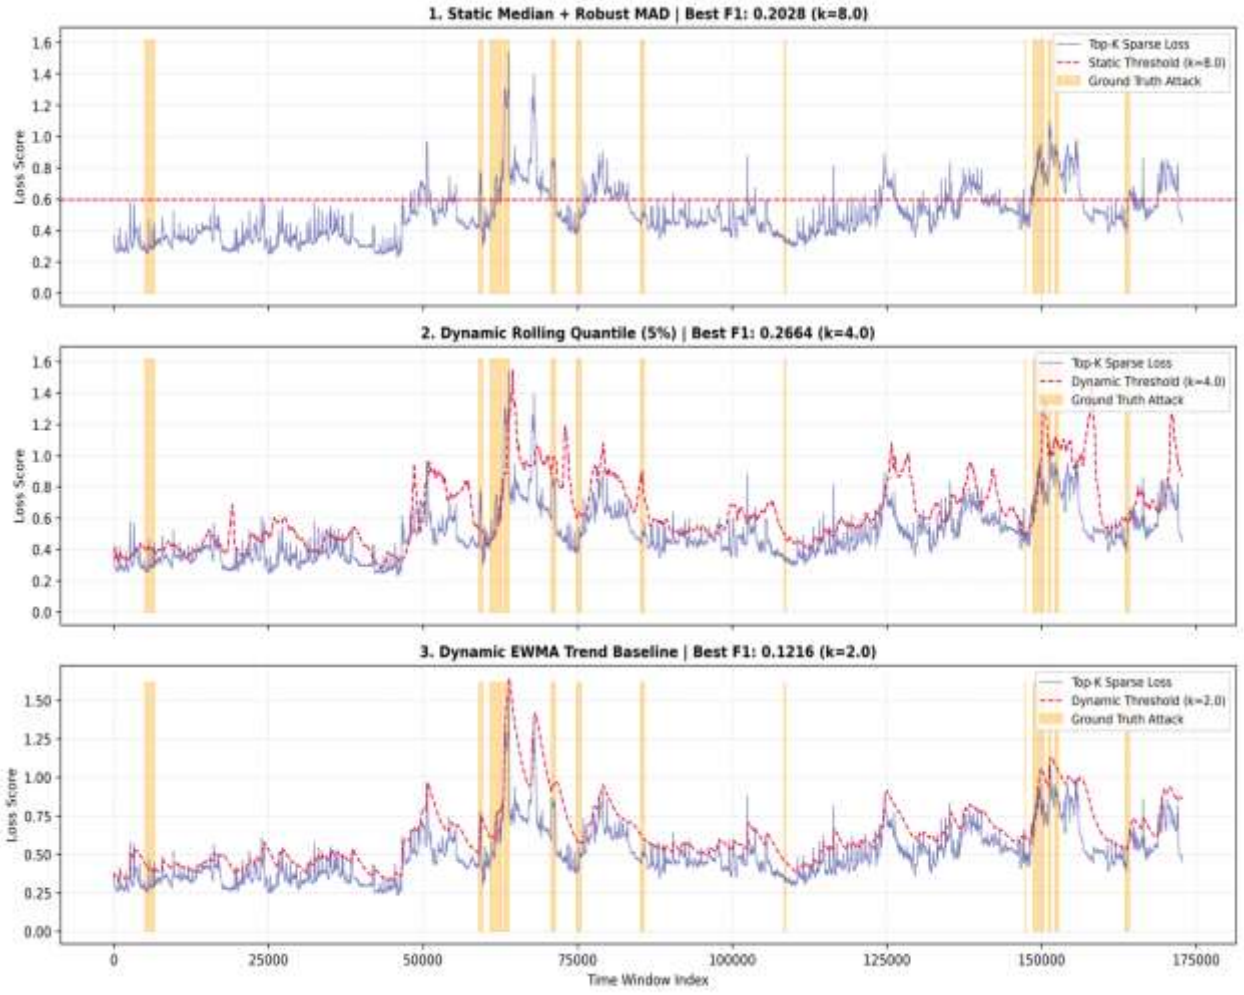

FIG. 2 COMPARING THE THRESHOLD VALUES OF THE MODEL WITH DIFFERENT PARAMETERS FOR TOP-K

Figure 2 presents a visual temporal comparison of the Top-K sparse reconstruction loss against ground-truth attack

windows across the three thresholding strategies. In Subplot 1, the fixed static boundary (τ= 0.60 at k=8.0) exhibits rigid failure modes: it sits too high to detect subtle early-stage attacks (0 -> 50,000) while triggering widespread false positives as the normal operational baseline naturally drifts upward in later sequences (60,000 -> 150,000). Subplot 3 illustrates that the Dynamic EWMA Trend boundary (τ at k=2.0) suffers from severe over-adaptation during prolonged physical attacks, as the moving average absorbs elevated error scores and inflates the decision threshold (τ > 1.25$), masking subsequent attack phases beneath the elevated line. In contrast, Subplot 2 demonstrates that the proposed 5% Dynamic Rolling Quantile strategy (τ at k=4.0) achieves superior adaptive performance by tracking nominal baseline drift smoothly while remaining immune to attack-driven threshold inflation, cleanly detecting acute physical attack spikes within ground-truth intervals without succumbing to false alarm saturation.

V. DISCUSSION

Experimental evaluation identifies Configuration 2 (Dynamic Rolling Quantile 5%) at k = 4.0 as the optimal detection boundary, achieving the highest overall trade-off with an F1-score of 0.6664. The resulting confusion matrix confirms 146,895 True Negatives, 15,902 False Positives, 1,001 False Negatives, and 8,976 True Positives. Static thresholding demonstrates high raw recall (0.9569 at k=2.0) but produces unacceptable false alarm rates due to its inability to accommodate baseline operational drift. Conversely, Dynamic EWMA tracking exhibits over-sensitivity during prolonged attacks, causing the decision boundary to over-adapt to elevated attack errors ($\tau t \geq 0.8369$) and leading to complete recall collapse at higher multipliers. The 5% Dynamic Rolling Quantile boundary effectively bridges these operational limits by maintaining low baseline drift while resisting attack-driven threshold inflation, proving highly effective for stealthy industrial anomaly detection.

VI. CONCLUSION

This study presented a Res-Attention Autoencoder framework paired with dynamic thresholding for cyber-physical anomaly detection on the WADI industrial dataset. By incorporating velocity-delta features and dynamic rolling quantile decision boundaries, the model effectively balances sensitivity and baseline stability, achieving an optimal F1-score of 0.6664 under 5% dynamic quantile scaling. Comparative evaluations demonstrate that dynamic quantile thresholding significantly outperforms static baseline thresholds and EWMA trend tracking by preventing false alarm saturation and resisting attack-induced threshold inflation during prolonged physical manipulations.

Despite its strong detection performance, this approach presents two key limitations. First, dynamic rolling quantile thresholding introduces a slight delay in detecting low-amplitude, slow-ramp setpoint manipulations, as incremental drift can remain temporarily masked within the rolling baseline window. Second, while the Top-K sparse reconstruction loss isolates heavily affected feature channels, performance remains dependent on manual hyperparameter tuning (e.g., multiplier k and window horizon W), which may require re-calibration across different industrial topologies.

Future work will focus on integrating graph neural networks (GNNs) to explicitly capture physical spatial dependencies and inter-sensor coupling across complex control loops. Additionally, we plan to explore adaptive, self-tuning threshold mechanisms using reinforcement learning to dynamically adjust boundary multipliers in real time based on system operational modes.

REFERENCES

[1] M. Ghobakhloo, “Industry 4.0, digitization, and opportunities for sustainability,” Journal of Cleaner Production, vol. 252, p. 119869, 2020.

[2] M. Masoud, Y. Jaradat, A. Manasrah, and I. Jannoud, “Sensors of smart devices in the internet of everything (IoE) era: Big opportunities and massive doubts,” Journal of Sensors, vol. 2019, no. 1, p. 6514520, 2019.

[3] M. A. Jawad, J. García, and M. Masoud, “A survey of network intrusion detection systems (NIDS) based on machine learning algorithms for industrial Internet of Things (IIoT),” in Proc. Conf. Sustainability and Cutting-Edge Business Technologies, Cham, Switzerland: Springer Nature, 2025, pp. 158–167.

[4] A. A. Maroli, V. S. Narwane, R. D. Raut, and B. E. Narkhede, “Framework for the implementation of an Internet of Things (IoT)-based water distribution and management system,” Clean Technologies and Environmental Policy, vol. 23, no. 1, pp. 271–283, 2021.

[5] M. AbdulJawad, M. Z. Masoud, Á. Álesanco, and J. García, “A deception-based access control mechanism for protecting PLCs from ModbusTCP brute-force attacks in IIoT environments,” Future Internet, vol. 18, no. 5, p. 259, 2026.
[6] C. M. Ahmed, V. R. Palleti, and A. P. Mathur, “WADI: A water distribution testbed for research in the design of secure cyber physical systems,” in Proc. 3rd Int. Workshop Cyber-Physical Systems for Smart Water Networks, 2017, pp. 25–28.
[7] J. Audibert, P. Michiardi, F. Guyard, S. Marti, and M. A. Zuluaga, “USAD: UnSupervised anomaly detection on multivariate time series,” in Proc. 26th ACM SIGKDD Int. Conf. Knowledge Discovery & Data Mining (KDD ’20), Virtual Event, CA, USA, Aug. 23–27, 2020. ACM, pp. 3395–3404.
[8] A. Deng and B. Hooi, “Graph neural network-based anomaly detection in multivariate time series,” in Proc. AAAI Conf. Artificial Intelligence (AAAI ’21), Virtual Event, Feb. 2–9, 2021. AAAI Press, vol. 35, no. 5, pp. 4027–4035.
[9] T. Kieu, B. Yang, C. Guo, and C. S. Jensen, “Outlier detection for time series with recurrent autoencoder ensembles,” in Proc. 28th Int. Joint Conf. Artificial Intelligence (IJCAI ’19), Macao, China, Aug. 10–16, 2019. AAAI Press, pp. 2725–2732.
[10] Y. Su, Y. Zhao, C. Niu, R. Liu, W. Sun, and D. Pei, “Robust anomaly detection for multivariate time series through stochastic recurrent neural network,” in Proc. 25th ACM SIGKDD Int. Conf. Knowledge Discovery & Data Mining (KDD ’19), Anchorage, AK, USA, Aug. 4–8, 2019. ACM, pp. 2828–2837.
[11] B. Zong, Q. Song, M. R. Min, W. Cheng, C. Lumezanu, D. Cho, and H. Chen, “Deep autoencoding Gaussian mixture model for unsupervised anomaly detection,” in Proc. 6th Int. Conf. Learning Representations (ICLR ’18), Vancouver, Canada, Apr. 30–May 3, 2018. OpenReview.net.
[12] Z. Li, Y. Zhao, J. Han, Y. Su, R. Jiao, X. Wen, and D. Pei, “Multivariate time-series anomaly detection via graph attention network,” in Proc. 21st IEEE Int. Conf. Data Mining (ICDM ’21), Auckland, New Zealand, Dec. 7–10, 2021. IEEE, pp. 841–850.
[13] R. Liu, Y. Zheng, and J. Chen, “Spatial-temporal attention network for multivariate time-series anomaly detection in industrial control systems,” IEEE Transactions on Industrial Informatics, vol. 19, no. 8, pp. 8812–8822, 2023.
[14] M. Zha, H. Jiang, and L. Zhang, “Ti-iLSTM: Tiny deep learning for PLC logic-level anomaly detection under extreme resource constraints,” IEEE Embedded Systems Letters, vol. 18, no. 1, pp. 45–48, 2026.
[15] Y. Wang, X. Zhang, and M. Chen, “Inverted transformer architecture for multivariate time series anomaly detection on physical testbeds,” IEEE Transactions on Neural Networks and Learning Systems, vol. 36, no. 2, pp. 1142–1154, 2025.
[16] S. Tuli, G. Casale, and N. R. Jennings, “TranAD: Deep transformer-based anomaly detection for multi-dimensional time series data,” Proc. VLDB Endowment (PVLDB), vol. 15, no. 6, pp. 1201–1214, 2022.
[17] L. Shen, Z. Li, and J. Kwok, “TimesNet: Temporal 2D-variation modeling for general time series analysis,” in Proc. 11th Int. Conf. Learning Representations (ICLR ’23), Kigali, Rwanda, May 1–5, 2023. OpenReview.net.