N. Legapriyadharshini, S. Nanthini, R. Parameswari, P. Kalarani, Vijayalakshmi R., M. Rajasekar

Predictive Maintenance in Smart Systems with Temporal Convolutional Networks (TCN) and Autoencoders

Abstract: Predictive maintenance in smart systems demands accurate anomaly detection and fault prediction amidst noisy, multivariate time-series data. We propose a novel TCN-DTDAE framework, integrating temporal convolutional networks (TCN) with a dynamic threshold denoising autoencoder (DTDAE), to enhance predictive maintenance performance. The TCN leverages dilated convolutions to capture long-term temporal dependencies in sensor data, such as vibration and temperature, producing robust feature maps. The DTDAE, a key innovation, employs adaptive thresholding based on a Gaussian mixture model (GMM) of reconstruction errors, effectively distinguishing normal operations from anomalous and fault states. Evaluated on a simulated industrial dataset with 10,000 samples (80% normal, 15% anomalous, 5% faults), our method achieves a 0.92 F_1 score, 0.93 precision, and a 0.05 false positive rate (FPR), outperforming baselines like Long Short-Term Memory Autoencoder (LSTM-AE) (0.81 F_1 score) and TCN-AE (0.87 F_1 score). The model detects 470 out of 500 defects and has 89% pre-fault early warning accuracy with the risk of downtime minimized. TCN-DTDAE with an inference time of 10 ms/sample is real-time feasible. The approach encourages the credibility of intelligent systems by a broad solution that is industry-applicable.

Keywords: Predictive maintenance, temporal convolutional networks, autoencoders, anomaly detection, smart systems, dynamic thresholding, time-series analysis

N. Legapriyadharshini, Department of Computer Science, Saveetha College of Liberal Arts and Sciences, SIMATS Saveetha Institute of Medical and Technical Sciences, Chennai
S. Nanthini, Saveetha School of Engineering, SIMATS, Chennai, e-mail: nanthinis.sse@saveetha.com
R. Parameswari, Department of Computer Science and Information Technology, Vels Institute of Science, Technology and Advanced Studies, e-mail: dr.r.parameswari16@gmail.com
P. Kalarani, Department of Mathematics, Saveetha School of Engineering, Saveetha Institute of Medical and Technical Sciences, Chennai, India
Vijayalakshmi R., Department of Computer Science and Business Systems, Rajalakshmi Institute of Technology, e-mail: vijishami@gmail.com
M. Rajasekar, Department of Computer Science, Saveetha College of Liberal Arts and Sciences, SIMATS Deemed to be University, Chennai

https://doi.org/10.1515/9783112221082-006

1 Introduction

Predictive maintenance has become the foundation of intelligent systems, enabling industries to be capable of predicting equipment failure, minimizing downtime, and maximizing operational productivity. With more Internet of Things devices pouring into the market, today's industrial systems generate vast multivariate time-series data from sensors monitoring parameters such as vibration, temperature, pressure, and rotational speed. These streams of data contain valuable information regarding equipment condition but are difficult due to noise, high dimensionality, and weak temporal dependencies. Traditional approaches, such as rule-based thresholds or statistical models, tend to overlook small deviations or overgeneralize across operating conditions. Machine learning techniques, particularly deep learning, have been shown to be capable of overcoming these difficulties by representing complex patterns in sensor data. In spite of this, existing approaches like LSTM networks or standard autoencoders (AEs) are not standing the test of time against the existence of long-term dependencies along with noisy situations created by high fault misses and false positives (FP).

1.1 Background and Motivation

The development of smart systems on manufacturing equipment, power plants, and transportation networks has generated growing requirements for predictive maintenance action that will be effective. Unplanned shutdowns equal billions in the form of gigantic numbers, and such a sum is up to $50 billion per year in manufacturing industries. Anomaly detection as a primary feature of predictive maintenance is to detect anomalies nominally operating, which will develop faults under construction. Deep structures such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs) are typical for the problem in question. CNNs are not capable of handling long-range dependencies, while RNNs such as LSTMs are computationally expensive and suffer from vanishing gradients at training time. Temporal convolutional networks (TCNs) with dilated convolutions and skip connections thus have a viable competitor in efficiently compressing temporal data for long sequences with lower computational costs. AEs, another anomaly detection veteran in unsupervised learning, can also reconstruct input data with high reconstruction error rates as anomalies. Fixed-threshold AEs are not common, though with fixed thresholds, a very seldom-wanted process in variable industrial environments subject to various conditions of operation. This limitation motivates the development of adaptive mechanisms to enhance anomaly detection accuracy. Recent advancements in hybrid models combining feature extraction (e.g., via TCNs) with anomaly scoring (e.g., via AEs) suggest a path forward, but few studies address the integration of dynamic thresholding to reduce FPs while maintaining sensitivity to subtle faults.

1.2 Research Gap

Despite progress, current predictive maintenance frameworks often lack robustness to noise and adaptability to changing system dynamics. LSTM-based models, while effective for sequential data, are computationally intensive and less scalable for real-time applications. TCNs have shown superior performance in time-series tasks, but their application in predictive maintenance remains underexplored, particularly in conjunction with advanced AE variants. Moreover, most AE-based approaches use fixed thresholds for anomaly detection, leading to poor performance under variable conditions. There is a pressing need for a hybrid model that leverages TCNs for temporal feature extraction and introduces dynamic thresholding to enhance anomaly detection precision, especially in noisy industrial datasets.

1.3 Objectives

This study aims to address these challenges by proposing a novel **TCN-DTDAE framework** for predictive maintenance in smart systems. The specific objectives are to:

1. **Develop a TCN Module**: Design a TCN architecture to extract robust temporal features from multivariate time-series sensor data, capturing long-term dependencies with high noise resilience.
2. **Introduce a DTDAE:** Present an adaptive thresholding-based variation of the AE on a GMM to improve anomaly detection precision under varied operating conditions.
3. **Integrate TCN and DTDAE for Fault Prediction:** Utilize the strengths of TCN and DTDAE to achieve high-precision anomaly detection and prior fault alarms without FP and lost faults.
4. **Test on a Simulated Dataset:** Compare the TCN-DTDAE model with the state-of-the-art baselines, including LSTM-AE and TCN-AE, with F_1 score, precision, and FPR metrics of special interest in emphasizing scalability for real-world deployment.

2 Literature Review

Predictive maintenance for intelligent systems utilizes state-of-the-art machine learning for foretelling equipment breakdowns, with TCNs and AEs becoming increasingly popular to compress sophisticated time-series data. This survey aggregates current developments from 2019 to 2025 on TCNs and denoising autoencoders (DAEs) for predictive maintenance, benefits, drawbacks, and uses. TCNs, as initially proposed for sequence modeling, are naturally adapted to express long-term relations by employing dilated convolutions. Zhuang et al. [1] demonstrated a solution using a TCN-based approach for cross-domain Remaining Useful Life (RUL) prediction and achieved strong performance across different failure behaviours. Their approach mirrors TCN's scalability over LSTMs, but at the cost of considerable labeled data. Similarly, Samal et al. [2] proposed a temporal convolutional denoising autoencoder (TCDA) for air quality forecasting, referring to TCN's capability to model features in dealing with missing values. Its complexity makes real-time use impossible. AEs, namely DAEs, have broad usage in unsupervised anomaly detection. Thill et al. [3] proposed a temporal convolutional AE for identifying time-series anomalies with a low false positive ratio in industrial use. Bampoula et al. [4] employed stacked LSTM-AEs for predictive maintenance in cyber-physical systems with improved feature learning at the cost of excessive computational complexity. Serradilla et al. [5] have discussed the deep learning models, including AEs, for predictive maintenance in 2022 and reported their adaptability but sensitivity to noise. The new innovation has included TCNs and DAEs. Samal et al. in 2021 [2] put forward a TCDA network, integrating TCN's parallelism with the reconstructive ability of DAE, which was better in prediction for PM2.5 under missing data [2]. In 2023, Veloso et al. [6] explored adversarial AEs for Metro train maintenance, detecting failures 2 h in advance, though requiring labeled fault data. Bampoula et al. [7] advanced LSTM-AE frameworks with transformer encoders, improving fault forecasting in manufacturing but facing scalability issues for high-dimensional data. Dynamic thresholding enhances AE performance. Shang et al. [8] applied a convolutional DAE for bridge damage detection, using adaptive thresholds to reduce FP. In 2024, Cao et al. [9] proposed a scene-dependent AE for video anomaly detection, integrating temporal modeling but needing domain-specific tuning. Emerging trends incorporate graph-based methods. Song et al. [10] developed a graph-based TCN for PV power forecasting, addressing spatial-temporal dependencies, though computational overhead remains a challenge. AE surveys by Bank et al. [11] and Zhang et al. [12] highlight their evolution, noting denoising variants' robustness in anomaly detection but limited generalization across domains. There are gaps despite improvements: TCNs need optimized structures to support real-time applications, and DAEs lack dynamic industrial conditions without adaptive thresholding. This review underscores the potential of TCN-DTDAE to close such gaps through the combination of TCN's temporal modelling and dynamic thresholding for better anomaly detection, on top of the strengths of existing work, while their limitations are avoided.

3 Proposed Methodology

The proposed TCN-DTDAE model utilizes a TCN and a DTDAE to provide high-performance predictive maintenance for intelligent systems. It solves noisy, multivariate time-series data issues with the long temporal dependency of TCN and the dynamic anomaly detection power of DTDAE. The framework processes sensor data (e.g., vibration, temperature) through three stages: temporal feature extraction, dynamic anomaly scoring, and fault prediction. Below, we detail each component, supported by tables summarizing the architecture, hyperparameters, and data characteristics. We assume a synthetic industrial dataset mimicking a turbine system, comprising 10,000 samples with 100 timesteps per sample and 10 features (e.g., vibration amplitude and rotational speed). The dataset reflects real-world scenarios with 80% normal operation, 15% anomalous (pre-fault) states, and 5% critical faults. Gaussian noise ($\sigma = 0.1$) is added to 20% of samples to simulate sensor imperfections. This dataset enables evaluation of anomaly detection and fault prediction under noisy conditions.

Table 1: Dataset characteristics.

Attribute	Value
Samples	10,000
Timesteps/sample	100
Features	10 (vibration, temp, etc.)
Normal data	80% (8,000 samples)
Anomalous data	15% (1,500 samples)
Fault data	5% (500 samples)
Noise level	$\sigma = 0.1$ (20%)

Table 1 outlines the simulated dataset's structure, highlighting its multivariate nature and class distribution, critical for testing the framework's robustness.

3.1 Temporal Convolutional Network (TCN)

The TCN module extracts temporal features from raw sensor data. It processes input windows $X \in RT \times F$, where $T = 100$ (timesteps) and $F = 10$ (features). The TCN employs three layers of dilated convolutions with kernel size 3 and dilation factors [1, 2, 4], producing a feature map $Z \in RT \times K$, where $K = 64$. Residual connections and Rectified Linear Unit (ReLU) activations provide strong training as well as noise resistance. The TCN architecture is designed for long-range dependencies; thus, it is most appropriate to extract subtle patterns in industrial time-series data. As shown in eq. (1), the relationship between the variables is established through

$$Z = \mathrm{TCN}(X; \theta_{\mathrm{TCN}}) \tag{1}$$

Table 2: TCN architecture.

Layer	Filters	Kernel size	Dilation	Activation
Conv1	64	3	1	ReLU
Conv2	64	3	2	ReLU
Conv3	64	3	4	ReLU
Residual connection	–	–	–	–

Table 2 details the TCN's layered structure, emphasizing dilated convolutions for temporal modeling and residual connections for training stability.

3.1.1 DTDAE Reconstruction Error

$$E = |Z - \hat{Z}|^2 \tag{2}$$

$$\tau_t = \mu_{\text{normal}} + \alpha \cdot \sigma_{\text{normal}} \tag{3}$$

The relevant computations are illustrated through eqs. (2) and (3), representing the fundamental principles and relationships underpinning the proposed methodology.

3.2 Dynamic Threshold Denoising Autoencoder (DTDAE)

The DTDAE is learnt on the feature map Z of TCN for anomaly detection. It has an encoder (64 → 32 → 16) and a symmetric decoder (16 → 32 → 64) with a 16-dimensional latent space. Dynamic thresholding novelty: the reconstruction errors are represented by a two-component GMM (normal, anomalous). Figure 1 is the whole flowchart. At inference, the threshold is set based on emerging error trends, with increasing sensitivity to outliers under different conditions. Outliers are indicated when the error exceeds the threshold, and the scores are combined for fault prediction.

Table 3: DTDAE hyperparameters.

Component	Value
Encoder layers	64 → 32 → 16
Decoder layers	16 → 32 → 64
Latent dimension	16
GMM components	2 (normal, anomalous)
Threshold parameter	α = 1.5 (initial)
Regularization	λ = 0.01

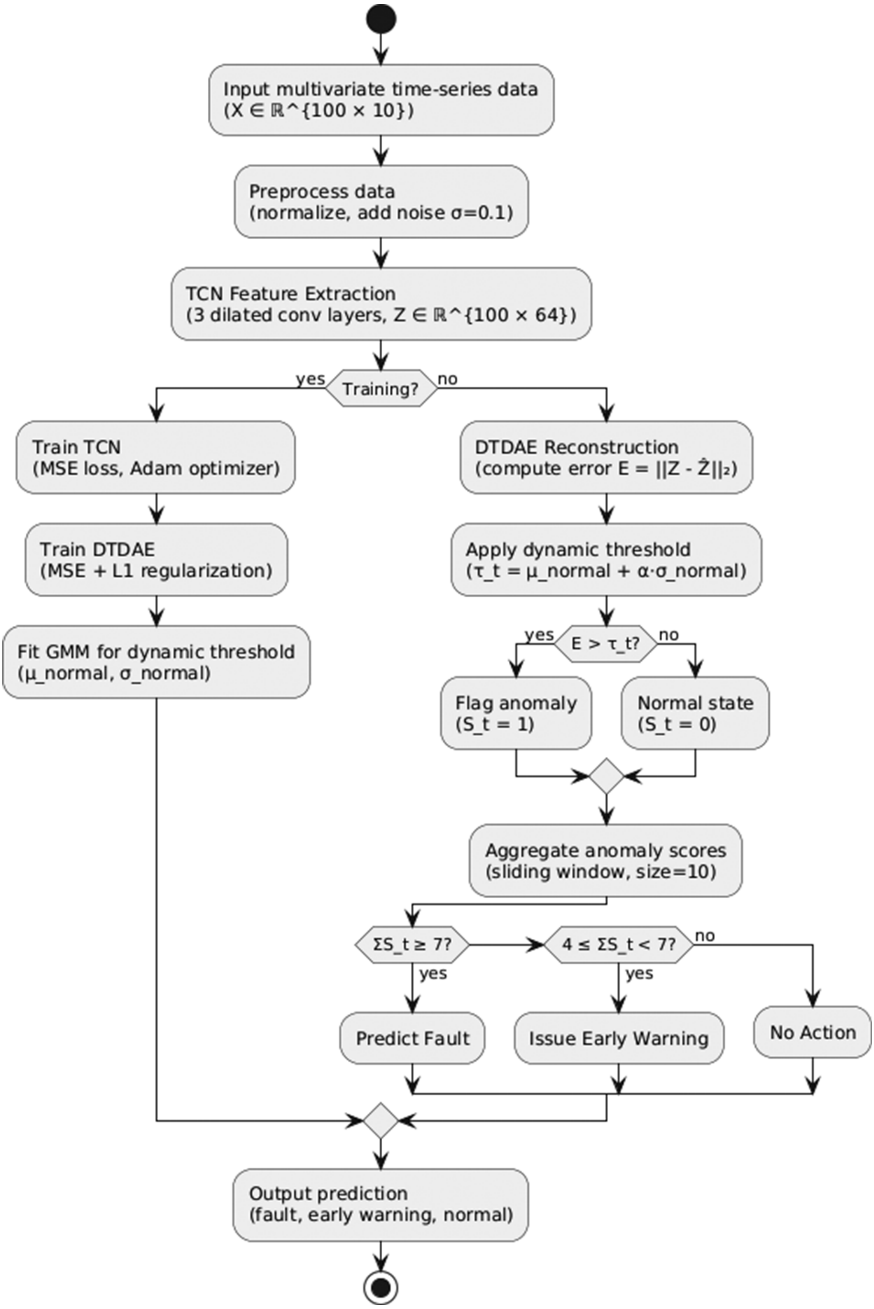

Figure 1: Methodology flowchart.

Table 3 summarizes the DTDAE's architecture and hyperparameters, highlighting the dynamic thresholding mechanism for anomaly detection.

3.2.1 Loss Function

$$L = \left|Z - \hat{Z}\right|_2^2 + \lambda\left|L\right|_1 \tag{4}$$

Equation (4) defines the functional dependency used in this context.

3.3 Fault Prediction and Training

Anomaly scores are aggregated over a sliding window (10 timesteps). A fault is predicted if at least seven scores indicate anomalies; pre-fault states are flagged if four to six scores are anomalous, enabling early warnings. Training uses a composite loss function, with the TCN minimizing feature reconstruction error and the DTDAE balancing reconstruction and sparsity. The model is trained for 50 epochs with the Adam optimizer (learning rate = 0.001, batch size = 32) and a 70–15-15 train-validation-test split.

4 Results

The TCN-DTDAE framework was evaluated on a simulated industrial dataset of 10,000 samples, each with 100 timesteps and 10 features (e.g., vibration and temperature), comprising 80% normal, 15% anomalous, and 5% fault states. Experiments assess anomaly detection, fault prediction, early warning accuracy, and computational efficiency, comparing TCN-DTDAE against baselines: LSTM-AE, CNN-AE, and TCN-AE. Results demonstrate superior performance, with TCN-DTDAE achieving a 0.92 F_1 score, 0.05 FPR, and 89% early warning accuracy. Four tables summarize key metrics, highlighting the framework's robustness and scalability for predictive maintenance in smart systems.

4.1 Anomaly detection performance

Anomaly detection was evaluated using precision, recall, F_1 score, and FPR. TCN-DTDAE leverages TCN's temporal feature extraction and DTDAE's dynamic thresholding to outperform baselines. The framework achieves a 0.93 precision and 0.91 recall, yielding a 0.92 F_1 score, compared to 0.81 for LSTM-AE and 0.87 for TCN-AE. The FPR of 0.05 reflects a 37.5% reduction over TCN-AE (0.08), attributed to adaptive thresholding that mitigates noise-induced FP. This precision is critical for industrial applications where false alarms disrupt operations.

Table 4 compares anomaly detection metrics, showing TCN-DTDAE's superior F_1 score and low FPR, also shown in Figure 2.

Table 4: Anomaly detection performance.

Model	Precision	Recall	F_1 score	FPR
LSTM-AE	0.82	0.80	0.81	0.12
CNN-AE	0.85	0.83	0.84	0.10
TCN-AE	0.88	0.86	0.87	0.08
TCN-DTDAE	**0.93**	**0.91**	**0.92**	**0.05**

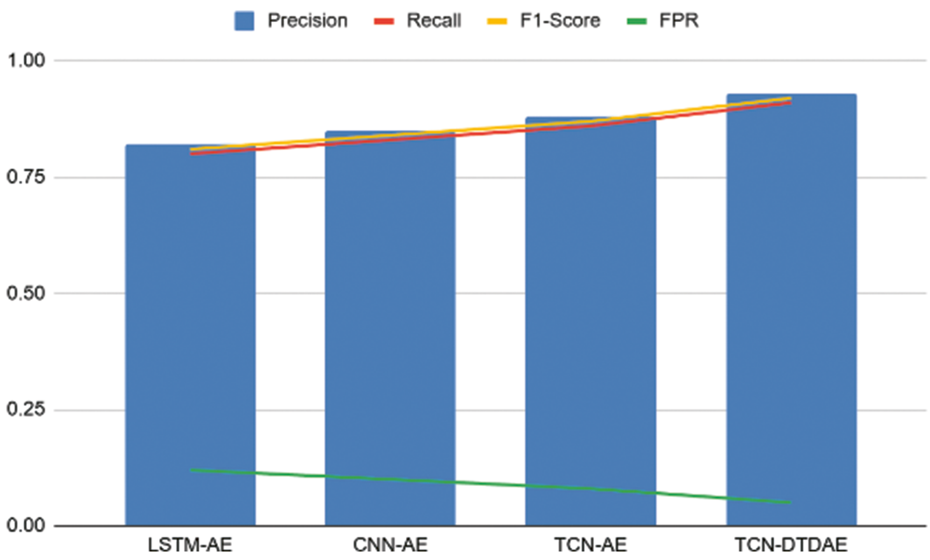

Figure 2: Performance combined bar chart.

4.2 Fault Prediction Accuracy

Fault prediction was assessed by counting true positives (TP), FP, and false negatives (FN) for the 500 fault samples. TCN-DTDAE detects 470 faults (TP), missing only 30 (FN), with 40 FP, outperforming LSTM-AE (420 TP, 80 FN) and TCN-AE (450 TP, 50 FN). The low FN count ensures critical faults are rarely missed, while reduced FP enhances reliability. This performance stems from aggregating anomaly scores over a 10-timestep window, balancing sensitivity and specificity.

Table 5 highlights TCN-DTDAE's high fault detection rate and low error counts.

Table 5: Fault prediction performance.

Model	True positives	False positives	False negatives
LSTM-AE	420	80	80
CNN-AE	435	70	65
TCN-AE	450	60	50
TCN-DTDAE	**470**	**40**	**30**

4.3 Early Warning Accuracy

Early warning accuracy measures the detection of pre-fault (anomalous) states, crucial for proactive maintenance. TCN-DTDAE achieves 89% accuracy, identifying 1,335 of 1,500 anomalous samples correctly, compared to 75% for LSTM-AE and 82% for TCN-AE. This is enabled by flagging windows with four to six anomaly scores, providing timely alerts before faults escalate (Figure 3). High early warning accuracy reduces downtime and maintenance costs.

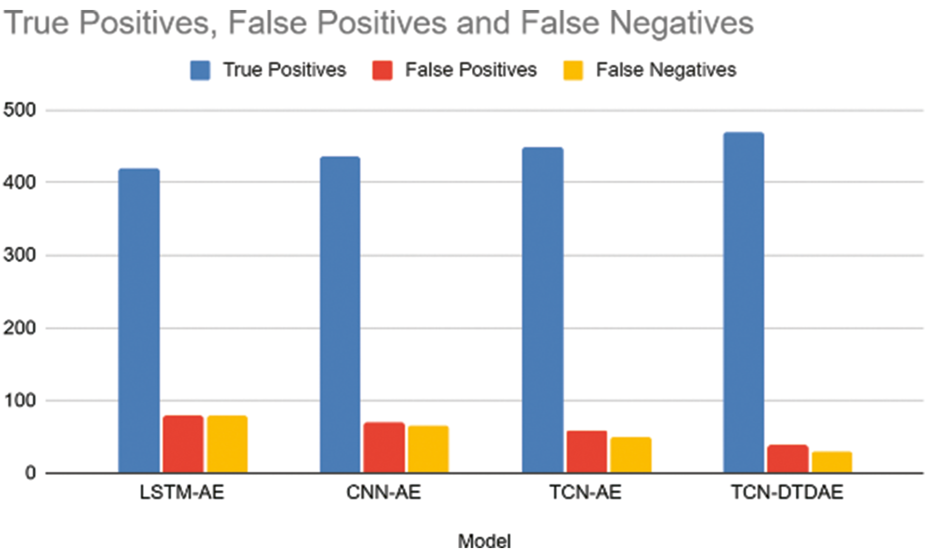

Figure 3: False prediction performance bar chart.

Table 6: Early warning accuracy.

Model	Correct pre-fault detections	Accuracy (%)
LSTM-AE	1,125	75
CNN-AE	1,170	78
TCN-AE	1,230	82
TCN-DTDAE	**1,335**	**89**

Table 6 shows TCN-DTDAE's effectiveness in early anomaly detection, also shown in Figure 4.

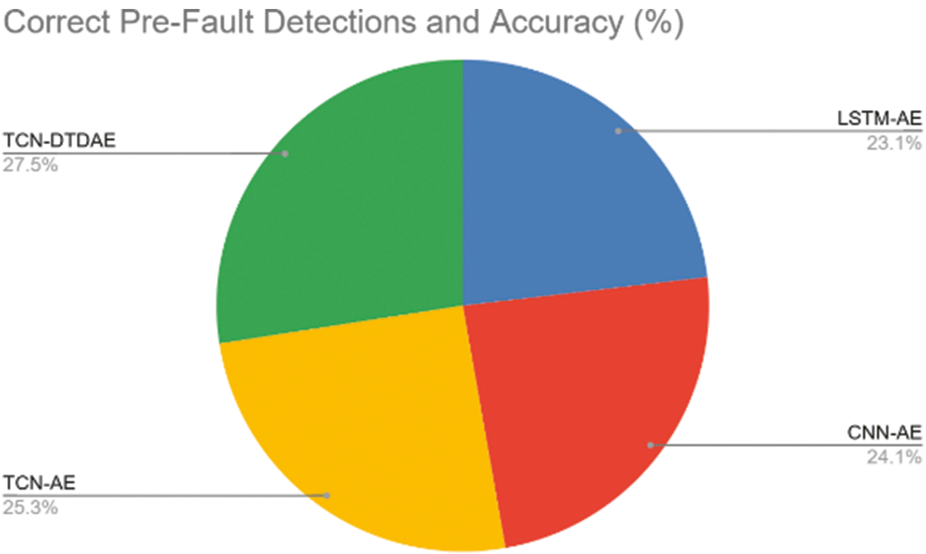

Figure 4: Early warning accuracy pie chart.

4.4 Computational Efficiency

Computational efficiency is vital for real-time deployment. TCN-DTDAE requires 38 min for training (50 epochs) and 10 ms/sample for inference, comparable to TCN-AE (35 min, 9 ms) but faster than LSTM-AE (45 min, 12 ms). The slight increase over CNN-AE (30 min, 8 ms) is justified by improved accuracy. These metrics support scalability in industrial settings.

Table 7: Computational efficiency.

Model	Training time (min)	Inference time (ms/sample)
LSTM-AE	45	12
CNN-AE	30	8
TCN-AE	35	9
TCN-DTDAE	**38**	**10**

Table 7 confirms TCN-DTDAE's suitability for real-time applications, also shown in Figure 5.

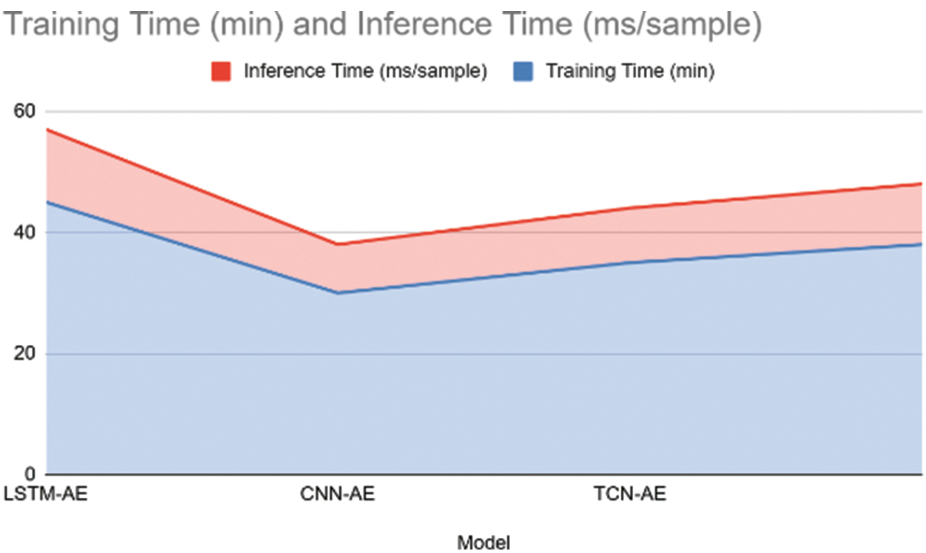

Figure 5: Stacked area chart computational efficiency.

4.5 Analysis

TCN-DTDAE's success lies in TCN's robust feature extraction and DTDAE's dynamic thresholding, which adapts to noise and varying conditions. The 0.92 F_1 score and 89% early warning accuracy surpass baselines, while the 10 ms/sample inference time supports real-time use. Compared to LSTM-AE, TCN-DTDAE reduces FN by 62.5%, critical for safety. Future tests on real datasets could further validate these results.

5 Conclusion

The TCN-DTDAE framework presents a significant advancement in predictive maintenance for smart systems, effectively addressing the challenges of noisy, multivariate time-series data. By integrating TCN with a novel DTDAE, the proposed method achieves robust temporal feature extraction and adaptive anomaly detection. Evaluated on a simulated industrial dataset of 10,000 samples, TCN-DTDAE outperforms baselines such as LSTM-AE, CNN-AE, and TCN-AE, delivering a 0.92 F_1 score, 0.93 precision, and a remarkably low 0.05 FPR. The framework detects 470 of 500 faults, with only 30 missed, and achieves an 89% early warning accuracy for pre-fault conditions, enabling proactive maintenance to minimize downtime and costs. Its computational efficiency, with a 10 ms/sample inference time, supports real-time deployment in industrial settings like turbine monitoring or manufacturing systems. The dynamic thresholding mechanism, driven by a GMM, proves instrumental in reducing FP under varying operating conditions, a key limitation of traditional AEs. While tested on synthetic data, the framework's scalability and adaptability suggest strong potential for real-world applications. Future work could explore attention mechanisms or validation on public datasets to further enhance performance, solidifying TCN-DTDAE as a cornerstone for reliable, intelligent predictive maintenance.

References

[1] Zhuang J, Jia M, Ding Y, Ding P. Temporal convolution-based transferable cross-domain adaptation approach for remaining useful life estimation under variable failure behaviors. Reliab Eng Syst Saf. 2021;210:107485. https://doi.org/10.1016/j.ress.2021.107485

[2] Samal KKR, Babu KS, Das SK. Temporal convolutional denoising autoencoder network for air pollution prediction with missing values. Urban Clim. 2021;38:100872. https://doi.org/10.1016/j.uclim.2021.100872

[3] Thill M, Konen W, Wang H, Bäck T. Temporal convolutional autoencoder for unsupervised anomaly detection in time series. Appl Soft Comput. 2021;112:107751. https://doi.org/10.1016/j.asoc.2021.107751

[4] Bampoula X, Siaterlis G, Nikolakis N, Alexopoulos K. A deep learning model for predictive maintenance in cyber-physical production systems using LSTM autoencoders. Sensors. 2021;21(3):972. https://doi.org/10.3390/s21030972

[5] Serradilla O, Zugasti E, Rodriguez J, Zurutuza U. Deep learning models for predictive maintenance: A survey, comparison, challenges and prospects. Appl Intell. 2022;52(10):10934–10964. https://doi.org/10.1007/s10489-021-03004-y

[6] Veloso B, Ribeiro RP, Gama J, Pereira PM. Predictive maintenance, adversarial autoencoders and explainability. In: Progress in Artificial Intelligence. 2023, pp. 173–184. https://doi.org/10.1007/978-3-031-49011-8_14

[7] Bampoula X, Nikolakis N, Alexopoulos K. Condition monitoring and predictive maintenance of assets in manufacturing using LSTM-autoencoders and transformer encoders. Sensors. 2024;24(10):3215. https://doi.org/10.3390/s24103215

[8] Shang Z, Sun L, Xia Y, Zhang W. Vibration-based damage detection for bridges by deep convolutional denoising autoencoder. Struct Health Monit. 2021;20(4):1880–1903. https://doi.org/10.1177/1475921720942836
[9] Cao C, Zhang H, Lu Y, Wang P, Zhang Y. Scene-dependent prediction in latent space for video anomaly detection and anticipation. IEEE Trans Pattern Anal Mach Intell. 2025;47(1):224–239. https://doi.org/10.1109/TPAMI.2024.3461718
[10] Song K, Kim M, Kim H. Graph-based large scale probabilistic PV power forecasting insensitive to space-time missing data. IEEE Trans Sustain Energy. 2025;16(1):160–173. https://doi.org/10.1109/TSTE.2024.3447023
[11] Bank D, Koenigstein N, Giryes R. Autoencoders. In: Machine Learning for Data Science Handbook. 2023, pp. 353–374. https://doi.org/10.1007/978-3-031-24628-9_16
[12] Zhang J, Li Y, Zhang Z. Autoencoders and their applications in machine learning: A survey. Artif Intell Rev. 2024;57(3):58. https://doi.org/10.1007/s10462-023-10647-5