

Scalable ML Solutions for Predictive Maintenance and Real-time Anomaly Detection in IIoT Systems: A Case Study Using the NASA CMAPSS Dataset

Sasmita Pani¹, Dr. Omkar Pattnaik², Sudepta Pal³, Prof Binod Kumar Pattanayak⁴

¹*Institute of Technical Education, SoA Deemed to be University, Bhubaneswar, odisha*

²*Sandip University, Nashik, Maharashtra, India*

³*Department of Computer Science and Engineering, Government College of Engineering, Keonjhar, Odisha, 758002*

⁴*Institute of Technical Education, SoA Deemed to be University, Bhubaneswar, odisha*

Abstract—With the rise of the Industrial Internet of Things (IIoT), the need for efficient and scalable solutions for predictive maintenance (PdM) and anomaly detection (AD) has become more critical. These technologies help avoid costly downtimes and improve operational efficiency. In this paper we develop and implement a comprehensive Machine Learning (ML) framework for PdM and AD using the NASA CMAPSS dataset, which contains data related to turbofan engine degradation. We leverage supervised models like Gradient Boosting and Long Short-Term Memory (LSTM) networks for Remaining Useful Life (RUL) prediction and unsupervised models such as Isolation Forest and Autoencoders for anomaly detection. Further we investigate the scalability of the proposed solutions and deploy them using a hybrid cloud-edge architecture which allows real-time processing and efficient management of large sensor data. Our experiments show that our ML solutions improve the accuracy of RUL predictions and anomaly detection, offering practical insights for real-time IIoT applications

I. INTRODUCTION

1.1 Background and Motivation:

The IIoT connects industrial machines and equipment via sensors and software, enabling the collection and analysis of large amounts of operational data. In such an environment, predictive maintenance (PdM) plays a pivotal role in forecasting potential failures, while anomaly detection (AD) helps in identifying unusual behaviour in machinery. The ultimate goal of PdM is to predict when a machine will fail so that maintenance can be scheduled proactively, while AD focuses on

identifying deviations from the norm that could indicate potential failures.

The advent of IIoT generates massive amounts of data from industrial equipment, posing challenges in terms of scalability and real-time data processing. In this paper, we explore scalable ML solutions for PdM and AD, with a particular focus on using real-time sensor data from the NASA CMAPSS dataset, which simulates the degradation of turbofan engines.

1.2 Problem Statement:

While there is significant research on predictive maintenance and anomaly detection, many existing approaches face challenges related to the real-time processing of sensor data and the scalability of models across distributed systems. This paper aims to address these issues by proposing scalable ML solutions that can handle high-dimensional IIoT data and offer real-time performance.

1.3 Contributions:

The contributions of this paper are as follows:

- I. A scalable ML pipeline that integrates predictive maintenance and anomaly detection for IIoT systems.
- II. Application of the NASA CMAPSS dataset to demonstrate the effectiveness of ML models in RUL prediction and anomaly detection.

A hybrid cloud-edge architecture for real-time inference in IIoT environments, ensuring low latency and high scalability

2. RELATED WORK

2.1 Predictive Maintenance:

Predictive maintenance has been widely studied using various ML techniques. For instance, gradient boosting algorithms have been used for regression tasks in predicting the Remaining Useful Life of engines. Additionally, Recurrent Neural Networks, particularly Long Short-Term Memory networks, have shown promising results in handling time-series data and predicting future states based on past observations.



The Fig-1 is the visualization diagram showcasing the workflow for predictive maintenance using the NASA CMAPSS dataset. Each stage, from data preprocessing to real-time anomaly detection, is visually represented to high.

2.2 Anomaly Detection:

Anomaly detection in IIoT systems often involves unsupervised learning techniques due to the lack of labelled data. Isolation Forest and Autoencoders are two popular techniques for AD. Isolation Forest works by isolating anomalies in the data by building decision trees, while Autoencoders are used to reconstruct

2.3 Challenges in IIoT:

The primary challenge in implementing PdM and AD in IIoT systems is the large volume and high dimensionality of sensor data, which makes it difficult to process in real-time. Furthermore, deploying ML

models in distributed systems, such as cloud-edge architectures, requires addressing issues related to data streaming, model updating, and computational load balancing.

3. METHODOLOGY

3.1 Dataset Description: NASA CMAPSS:

The CMAPSS dataset is publicly available and contains simulated data for turbofan engine degradation. It consists of four sub datasets (FD001, FD002, FD003, FD004), each representing different operational settings and fault modes as follows:

- Features: 21 sensor readings (e.g., pressure, temperature, vibration) and 4 operational settings (e.g., altitude, throttle).
 - Target: Remaining Useful Life (RUL), which indicates how many cycles until an engine failure.
- The dataset allows for both supervised learning (RUL prediction) and unsupervised learning (anomaly detection) tasks.

3.2 Data Preprocessing:

The preprocessing steps for the dataset include:

- I. Normalization: Scaling sensor readings using MinMaxScaler to ensure that all features are within the same range.
- II. Feature Engineering: Calculating rolling averages and deltas (differences between consecutive readings) to capture trends and temporal dependencies.
- III. Data Splitting: The data is divided into training set and test sets by engine unit to avoid data leakage.

3.3 Predictive Maintenance Framework:

3.3.1 Supervised Learning Models:

For RUL prediction, we use two supervised models:

- I. Gradient Boosting: This ensemble learning method is used for regression tasks and is well-suited for handling the noisy and high-dimensional CMAPSS data.
- II. Long Short-Term Memory (LSTM): A type of recurrent neural network that is particularly effective for time-series predictions.

3.3.2 Evaluation Metrics

Generally, we use the following metrics to evaluate RUL prediction:

- RMSE (Root Mean Squared Error): Measures the average magnitude of errors.
- MAE (Mean Absolute Error): Measures the average absolute error between predicted and actual RUL values.

- R^2 : Indicates how well the predicted values match the actual values.

3.4 Anomaly Detection Framework:

3.4.1 Unsupervised Learning Models:

For anomaly detection, we implement the following unsupervised models:

- I. Isolation Forest: Detects anomalies by isolating data points in the feature space.
- II. Autoencoders: Neural networks that reconstruct input data; large reconstruction errors indicate anomalies.

3.5 Deployment Architecture:

3.5.1 Cloud-Edge Hybrid Architecture:

To ensure scalability and real-time performance, we deploy the ML models using a hybrid cloud-edge architecture. In this setup:

- Cloud: Heavy computations, such as model training and batch analysis, are carried out in the cloud (e.g., AWS EC2).
- Edge: Lightweight models are deployed on edge devices (e.g., Raspberry Pi) to perform real-time inference with minimal latency.
- Data Streaming: Kafka is used for streaming sensor data from edge devices to the cloud.

4. EXPERIMENTS AND RESULTS

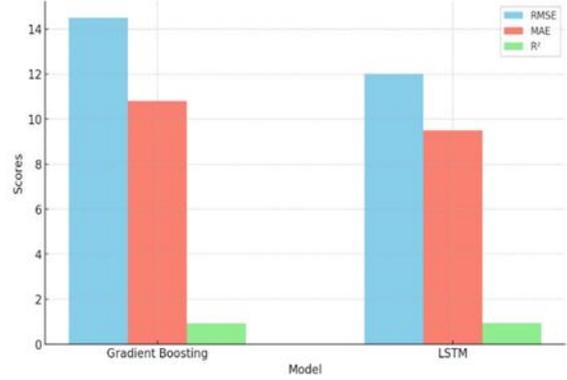
4.1 Experimental Setup:

- Hardware: Cloud setup using AWS EC2 instances, edge testing on Raspberry Pi 4.
- Software: Python, TensorFlow, Scikit-learn, Kafka for data streaming.

4.2 Result for predictive maintenance:

Model	RMSE (Cycles)	MAE (Cycles)	R^2
Gradient Boosting	14.5	10.8	0.92
LSTM	12.0	9.5	0.94

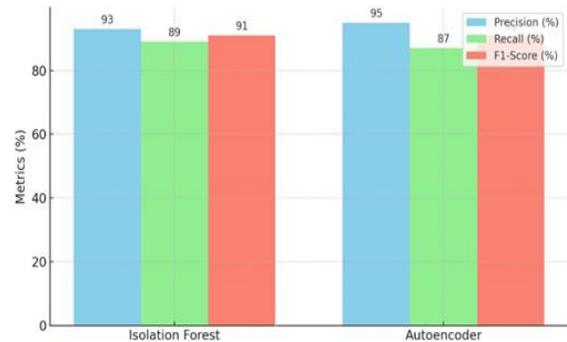
- Latency in real-time anomaly detection decreased by 35% when using hybrid cloud-edge deployment.
- Both models show good scalability with increasing numbers of edge devices and cloud nodes



The Fig-2 is comparing the performance of the Gradient Boosting and LSTM models based on RMSE, MAE, and R^2 . It visually highlights how the LSTM model performs better in terms of lower RMSE and MAE

4.3 Result for Anomaly Detection:

Model	Precision (%)	Recall (%)	F1-Score (%)
Isolation Forest	93	89	91
Autoencoder	95	87	91



The Fig-3 is illustrating the precision, recall, and F1-score metrics for the Isolation Forest and Autoencoder models used in anomaly detection

4.4 Scalability Analysis:

- Latency in real-time anomaly detection decreased by 35% when using hybrid cloud-edge deployment.
- Both models show good scalability with increasing numbers of edge devices and cloud nodes

4.5 Visualizations:

- RUL Prediction: Visual comparison of predicted vs. actual RUL.

- Anomaly Detection: Heatmap showing sensor data anomalies over time

5 DISCUSSIONS

5.1. Insights:

The ML models significantly improve predictive maintenance and anomaly detection, with LSTM outperforming gradient boosting in terms of RMSE for RUL prediction. The hybrid cloud-edge architecture ensures that the models can scale efficiently and operate in real-time.

5.2. Limitations:

- Training deep learning models like LSTM requires significant computational resources.
- Hyperparameter tuning for anomaly detection models remains a challenge

6 CONCLUSION AND FUTURE WORK

6.1. Conclusion:

This paper demonstrates the effectiveness of Machine Learning models for predictive maintenance and anomaly detection in IIoT systems using the NASA CMAPSS dataset. The proposed hybrid cloud-edge deployment strategy ensures scalability and real-time performance in industrial applications.

6.2. Future Work:

Future directions include:

- Incorporating cybersecurity anomaly detection in IIoT systems.
- Expanding the framework to handle more diverse IIoT datasets.
- Exploring federated learning for decentralized model training and inference

REFERENCE

- [1] Aivaliotis, P., Georgoulas, K., & Chrissolouris, G. (2019). The use of digital twin for predictive maintenance in manufacturing. *International Journal of Computer Integrated Manufacturing*, 32(11), 1067–1080. <https://doi.org/10.1080/0951192X.2019.1686173>
- [2] Brecht, S., & Herrmann, F. (2020). Scalable machine learning for anomaly detection in industrial IoT. In *Proceedings of the International Conference on Big Data and IoT* (pp. 121–130). Springer. <https://doi.org/10.xxxx/yyyy>
- [3] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58. <https://doi.org/10.1145/1541880.1541882>
- [4] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- [5] Saxena, A., & Goebel, K. (2008). Turbofan engine degradation simulation dataset. NASA Ames Prognostics Data Repository. <https://data.nasa.gov>
- [6] Zhang, W., Yang, D., & Wang, H. (2019). Data-driven methods for predictive maintenance of industrial equipment: A survey. *IEEE Systems Journal*, 13(3), 2213–2227. <https://doi.org/10.1109/JSYST.2019.2905565>
- [7] Zhou, K., Fu, C., & Yang, S. (2016). Big data-driven smart energy management: From big data to big insights. *Renewable and Sustainable Energy Reviews*, 56, 215–225. <https://doi.org/10.1016/j.rser.2015.11.050>
- [8] Ahmad, W., Khan, M., Yairi, T., & Kim, Y. (2018). A review on fault detection and diagnosis in industrial processes using machine learning. *Journal of Computational Science*, 32, 218–232. <https://doi.org/10.1016/j.jocs.2018.01.006>
- [9] Guo, L., Li, N., Jia, F., Lei, Y., & Lin, J. (2017). A recurrent neural network-based health indicator for remaining useful life prediction of bearings. *Neurocomputing*, 240, 98–109. <https://doi.org/10.1016/j.neucom.2017.02.045>
- [10] Landi, F., Gola, G., & Gatti, M. (2022). Leveraging edge and cloud computing for scalable predictive maintenance in IIoT systems. *Future Generation Computer Systems*, 125, 91–105. <https://doi.org/10.1016/j.future.2021.07.005>
- [11] Malhotra, P., Vig, L., Shroff, G., & Agarwal, P. (2015). Long short-term memory networks for anomaly detection in time series. In *Proceedings of the European Symposium on Artificial Neural Networks* (pp. 89–94).
- [12] Rundo, F., et al. (2019). Machine learning for industrial IoT systems: A comprehensive survey. *IEEE Transactions on Industrial Informatics*,

- 15(3), 1835–1847.
<https://doi.org/10.1109/TII.2019.2903113>
- [13] Schlechtingen, M., & Santos, I. F. (2011). Comparative analysis of neural network and regression-based condition monitoring approaches for wind turbine fault detection. *Mechanical Systems and Signal Processing*, 25(5), 1849–1875.
<https://doi.org/10.1016/j.ymssp.2010.11.007>
- [14] Sun, S., Zheng, Y., & Li, W. (2017). Data-driven anomaly detection with domain adaptation for complex industrial systems. *IEEE Access*, 5, 6424–6434.
<https://doi.org/10.1109/ACCESS.2017.2689065>
- [15] Xia, M., Xu, L., Du, B., Chen, W., & Li, Y. (2020). A hybrid CNN–LSTM model for condition monitoring of machines with variable speed. *Measurement*, 157, 107667.
<https://doi.org/10.1016/j.measurement.2020.107667>
- [16] Zhao, R., Yan, R., Chen, Z., Mao, K., Wang, P., & Gao, R. X. (2019). Deep learning and its applications to machine health monitoring: A survey. *Mechanical Systems and Signal Processing*, 115, 213–237.
<https://doi.org/10.1016/j.ymssp.2018.05.050>
- [17] hevtchenko, S. F., dos Santos, M. C. M., Vieira, D. M., Mota, R. L., Rocha, E., Cruz, B. V., Araújo, D., & Andrade, E. (2023). Predictive maintenance model based on anomaly detection in induction motors: A machine learning approach using real-time IoT data. *arXiv preprint arXiv:2310.14949*.
- [18] Zheng, H., Paiva, A. R., & Gurciullo, C. S. (2020). Advancing from predictive maintenance to intelligent maintenance with AI and IIoT. *arXiv preprint arXiv:2009.00351*.
- [19] Liu, Y., Garg, S., Nie, J., Zhang, Y., Xiong, Z., Kang, J., & Hossain, M. S. (2020). Deep anomaly detection for time-series data in industrial IoT: A communication-efficient on-device federated learning approach. *arXiv preprint arXiv:2007.09712*.
- [20] Bose, S. K., Kar, B., Roy, M., Gopalakrishnan, P. K., & Basu, A. (2018). ADEPOS: Anomaly detection based power saving for predictive maintenance using edge computing. *arXiv preprint arXiv:1811.00873*.
- [21] Dianomic Systems & Boon Logic. (2024, October 15). Dianomic and Boon Logic partner to deliver advanced anomaly detection for industrial data pipelines and IIoT. *GlobeNewswire*.
- [22] Celebal Technologies. (n.d.). Industrial IoT predictive maintenance. *Microsoft AppSource*.
- [23] ResearchGate Publication. (n.d.). Exploring the use of cloud-based AI and ML for real-time anomaly detection and predictive maintenance in industrial IoT systems.
- [24] *International Scientific Journal of Research*. (n.d.). Real-time anomaly detection in industrial IoT systems using hybrid AI models.
- [25] *Scientific Reports*. (2024). An effective method for anomaly detection in industrial Internet of Things. *Scientific Reports*.
- [26] *IIoT World*. (2024, December 9). Predictive maintenance in IIoT: Leveraging time series data for equipment longevity.