*Proceedings of the International Conference on Industrial Engineering and Operations Management*

Publisher: IEOM Society International, USA DOI: 10.46254/BA08.20250470
Published: December 20, 2025

# A Predictive Maintenance Framework for Elevators: Detecting Mechanical Anomalies Using Unsupervised Deep Learning and Outlier Models

**Md. Efatuzzaman Efat**
Institute of Information and Communication Technology
Bangladesh University of Engineering and Technology
Dhaka, Bangladesh
efatuzzaman@gmail.com

**Sayed Md. Fazle Rabbi**
Department of Electrical and Electronic Engineering
Ahsanullah University of Science and Technology
Dhaka, Bangladesh
sayedfazlerabbi45@gmail.com

**Sumiya Afrose**
Department of Robotics and Mechatronics Engineering
University of Dhaka
Dhaka, Bangladesh
s.afrose098@gmail.com

**Shaikh Farhana Supti**
Institute of Information and Communication Technology
Bangladesh University of Engineering and Technology
Dhaka, Bangladesh
shaikhfarhanasupti@gmail.com

**Md. Saiful Islam**
Professor
Institute of Information and Communication Technology
Bangladesh University of Engineering and Technology
Dhaka, Bangladesh
mdsaifulislam@iict.buet.ac.bd

## Abstract

The reliability and safety of high-traffic urban infrastructure elevator systems heavily rely on predictive maintenance. This paper proposes a method to identify mechanical failures in elevator door systems based on unsupervised learning methods. Based on a publicly available dataset of the elevator predictive maintenance project by Huawei, we model the time-series data of vibration, humidity, and door bearing sensors sampled at a rate of 4 Hz. The proposed pipeline combines classical models of outlier detection, such as Isolation Forest, Local Outlier Factor, and One Class SVM, with a deep learning-based LSTM autoencoder. The anomalies are identified through calculating reconstruction errors and comparing these values with a dynamic threshold that is formed based on training error distribution. The anomalies identified correspond with high-vibration oscillations that can be signs of mechanical anomalies, e.g. door jitter or control instability. The model obtains the capability to isolate abnormal events without fault labels at the start through visualization of both full and zoomed sensor timelines. The hybrid model is a powerful and understandable model that may be deployed in the real-time in the elevator repair systems, integrating the advantages of both classical and deep learning frameworks. The paper shows the importance of AI in shifting towards proactive instead of reactive maintenance, minimizing down time, and increasing equipment life in electromechanical systems. The methodology can be widely used on other mechanical platforms that can monitor sensors.

## Keywords

Predictive Maintenance (PdM), Time-Series Fault Detection, Machine Learning for Fault Diagnosis, Smart Maintenance Systems, Anomaly Detection.

## 1. Introduction

Traditionally, industries have relied on reactive maintenance (fixing parts after failure) or preventive maintenance (periodic servicing) to keep equipment running. However, reactive "run-to-fail" approaches lead to extended downtime and costly repairs, while strictly schedule-based maintenance can be inefficient – performing unnecessary service and still failing to catch unforeseen faults (Jardine et al. 2006; Zonta et al. 2020). For mission-critical equipment such as elevators, these limitations highlight the necessity for smarter maintenance strategies that minimize downtime and ensure passenger safety. Predictive maintenance (also known as condition-based maintenance) has emerged to address these inefficiencies. Rather than servicing machinery at set intervals or only after breakdowns, predictive maintenance uses real-time condition monitoring and data analysis to anticipate failures in advance. This practice allows maintenance to be done on a just-in-time basis – components are repaired or replaced immediately before they are likely to fail, thus preventing unexpected outages and reducing downtime (Jardine et al. 2006). Research has demonstrated that properly executed predictive maintenance significantly reduces unplanned downtime compared to purely time-based schedules (Sampaio et al. 2019). With the advent of affordable sensors and the Industrial Internet of Things (IIoT), it has become possible to continuously monitor elevator parameters (vibration, speed, motor current, etc.) and identify early indicators of mechanical faults. Using such data streams, operators can shift to data-driven maintenance decision-making (Wang et al. 2016). This shift is especially valuable for elevator systems, where high availability is required and any entrapment or failure can be very dangerous. Leading elevator manufacturers have begun investing in smart maintenance software that processes sensor data in real time to make service calls only when needed, improving reliability and customer satisfaction (Wang et al. 2016). Recent advances in artificial intelligence (AI) and machine learning (ML) are further driving the adoption of predictive maintenance in mechanical systems. High-tech ML algorithms can automatically detect complex patterns in maintenance history and sensor measurements that might be overlooked by human operators, enabling more accurate and earlier failure predictions (Zhang et al. 2019). In particular, AI's capability to process large volumes of data and intricate trends aligns with the demands of modern elevator maintenance, where each unit produces abundant operational data daily. In the context of Industry 4.0, machines are expected to be self-aware and self-predictive (Lee et al. 2014) — capable of monitoring their own status, protecting themselves, and requesting care proactively. AI driven predictive maintenance systems move toward this vision by constantly evaluating the elevator's condition and informing technicians about looming problems. Altogether, predictive maintenance is of great significance for elevator systems because it can make them safer, minimize downtime, and optimize maintenance resources (Zonta et al. 2020).

However, implementing predictive maintenance for elevators still presents challenges. Early fault detection (especially for elevator door mechanisms) is difficult due to limited fault examples, and models must generalize across varying usage patterns. There is also a need for interpretable and computationally efficient solutions that engineers can trust in real time. This paper proposes an approach to address these 3 needs by leveraging the latest

AI/ML techniques in particular, a deep learning model to detect elevator mechanical anomalies before they escalate. The rest of the paper is structured as follows. Section 2 provides a review of the relevant literature, including maintenance strategies, the role of machine learning (with emphasis on techniques like LSTM networks and autoencoders), prior research on elevator maintenance, and the key research gaps. Section 3 describes the proposed methodology and hybrid anomaly detection framework. Section 4 presents the results and discussion of our approach. Finally, Section 5 concludes the paper and outlines directions for future work.

### 1.1 Objectives

- Designing a data-driven anomaly detection system for elevator doors using multivariate time-series sensor data.
- Implementing a hybrid pipeline combining LSTM autoencoder and classical outlier models for unsupervised fault detection.
- Mapping and visualizing mechanical anomalies in real-time to support predictive maintenance actions
- Ensuring the entire system is lightweight, interpretable, and suitable for deployment in real-world elevator maintenance operations.

## 2. Literature Review

**Strategies of Maintenance: Reactive to Predictive.**
Maintenance practices have evolved from purely reactive fixes or time-based preventive schedules toward condition-based predictive strategies. Reactive maintenance takes no action until a failure occurs, resulting in expensive downtime and damage (Jardine et al. 2006). Preventive maintenance can mitigate risk, but its rigid schedule may lead to premature servicing or failure to catch unforeseen faults (Nunes et al. 2023; Zonta et al. 2020). Such inefficiencies are critical concerns in elevator systems that require high uptime and safety. Predictive maintenance (PdM) offers a data-driven alternative by continuously monitoring sensor data to predict failures. This approach enables timely interventions to prevent breakdowns and minimize downtime while extending the life of components (Sampaio et al. 2019). PdM systems employ sensors to measure parameters like motor vibration and current so that elevators are serviced according to actual wear rather than preset intervals. Studies have found that PdM saves costs and improves reliability (Alsyouf 2009), representing a significant improvement over conventional maintenance practice.

**Application of AI and Machine Learning in Predictive Maintenance.**
Beyond basic maintenance strategies, recent advances in AI and ML have greatly expanded PdM capabilities by enabling the learning the complex fault patterns from historical sensor data. SVMs, decision trees, and ANNs, which are the classical ML models, have been employed to identify anomalies and forecast failures (Zhang et al. 2019; Jardine et al. 2006). It was demonstrated that ANNs are capable of predicting failures of electric motors with high accuracy (Sampaio et al., 2019), and Martinez-Rego et al. (2011) showed that SVMs could be used to detect faults in vibration. Deep learning, in particular, LSTM networks, further improve PdM by modelling time-varying behaviours and extracting features from raw signals. LSTMs have been used in different systems to predict health conditions and probability of failure (Zhang et al. 2019). Iuculano and Babar (2024) applied LSTMs to sensor logs of elevators to anticipate safety-chain failures with a high score (F1 of about 85%). These sequence-based models are especially effective with the elevators since they have a sequential operational pattern. Lack of labeled fault data has made unsupervised learning crucial. Autoencoders, and in particular LSTM-based ones, detect anomalies based on the reconstruction errors against normal operation baselines (Dou et al. 2022). The advantage of such models is that they are capable of detecting faults early without fault labels, which is important in systems such as elevators where failures are seldom seen. Autoencoders combined with classifiers have also demonstrated high accuracy and low false alarm rates (Mishra et al. 2019).

**Previous Elevator Predictive Maintenance Research.**
Applying these concepts to elevator systems, a number of studies have employed machine learning for predictive maintenance in elevator maintenance. Yan and Lee (2004) focused on door fault classification using logistic regression, while Wang et al. (2009) focused on brake faults diagnosis with the help of fuzzy neural networks and wavelet analysis. These initial investigations indicated that ML could be used to identify faults. Due to the development of IoT and cloud computing, more holistic solutions have been developed. Wang et al. (2016) suggested Intelligent Predictive Maintenance (IPdM) system, which integrates CPS (Cyber-Physical Systems), sensors, and analytics in health monitoring of elevators. Mishra et al. (2019) applied deep autoencoders and random

forests to identify elevator faults with close to perfect accuracy and almost no false positives. The study by Iuculano and Babar (2024) involved the use of the LSTM models to process the event logs and forecast elevator breakdowns before they occur. Their article resolved the issue of unbalanced data and presented real-life examples of using the LSTM-based PdM for elevators. Such studies affirm that ML techniques, particularly deep learning, can be useful in identifying early signs of faults and scheduling maintenance.

### Gaps in Existing Research

Nevertheless, several challenges remain. Most models are not validated on a large scale, and may struggle to generalizing to elevators with different usage patterns (Nunes et al. 2023). There is little labeled data available which constrains supervised learning methods. Although unsupervised techniques can help such is alleviate this, their effectiveness in a variety of fault modes remains unexplored. The fusion of data across several sensors is also an unsolved problem with the majority of literature concentrating on a single modality. The other issue is interpretability, as black-box deep models make predictions difficult to trust by engineers (Linardatos et al. 2021). Lastly, there are computational limits of elevator control systems that require efficient and real-time-capable models. This paper attempts to fill these gaps by combining the deep learning (LSTM autoencoders) with traditional ML methods to identify faults using real elevator operation data. Our unsupervised model does not require fault labels and focuses on scalable interpretable maintenance analytics. To address the above challenges, the proposed hybrid approach is detailed in Section 3, which integrates an LSTM autoencoder with classical outlier detection models for elevator fault detection.

## 3. Methodology

### 3.1 Dataset Description and Preprocessing

This study employs the *Predictive Maintenance Dataset for Elevator Systems* developed by the Huawei German Research Center. The dataset was collected under real-world operating conditions to support the advancement of predictive maintenance methodologies in elevator systems (Huawei German Research Center 2020). It comprises multivariate time-series recordings captured from elevator door subsystems during high-traffic evening hours (16:30–23:30), ensuring the inclusion of representative mechanical activity across diverse usage scenarios. Figure 1 presents the readings of all three sensors over the full observation period (~8 hours).

Sensor data were sampled at a consistent frequency of 4 Hz and encompass three distinct channels:

- **Ball Bearing Sensor**, representing electromechanical movement and potential component degradation;
- **Vibration Sensor**, capturing dynamic oscillations, impact signatures, and anomalous motion patterns;
- **Humidity Sensor**, reflecting environmental conditions that may influence mechanical reliability or signal interpretation.

The dataset is of significant academic value due to its high temporal resolution, clean sampling structure, and its provenance from an industrial research environment. Its design enables robust benchmarking of time-series analysis techniques, anomaly detection algorithms, and deep learning models applied to mechanical fault prediction.

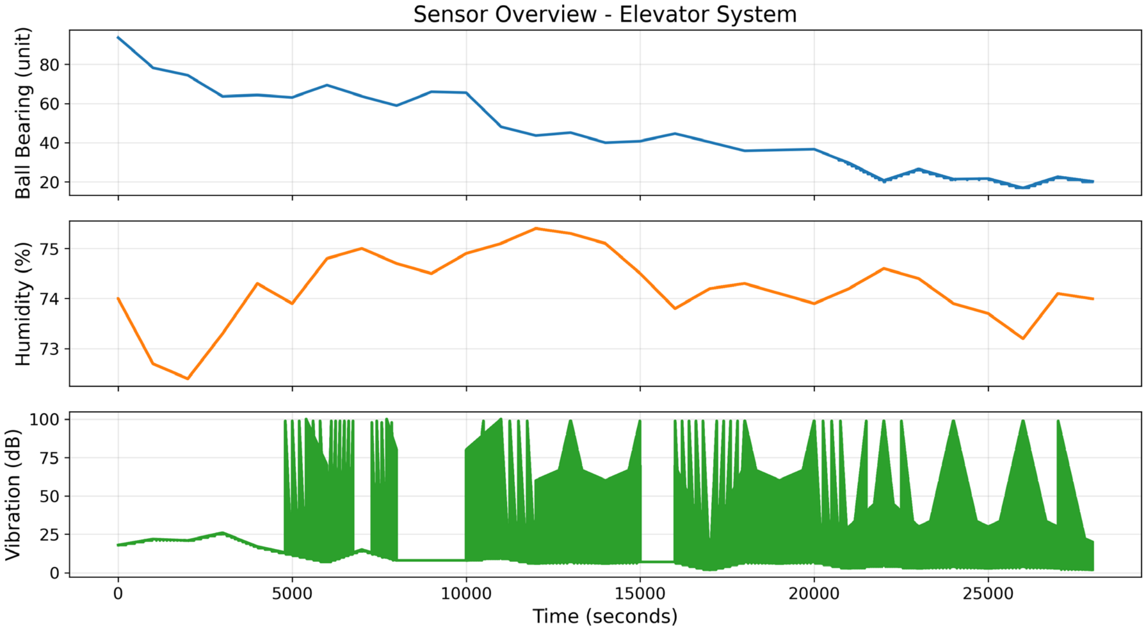

Figure 1. Time series of three key elevator door sensors over the full observation period.

Data were recorded over approximately eight hours at a sampling interval of 0.25 seconds, resulting in over 112,000 timestamped samples for each sensor. Basic preprocessing steps included cleaning any obvious outliers or missing readings (none were present in the dataset) and normalizing sensor scales where appropriate. The vibration readings (in dB) and the ball-bearing sensor 2 values were each scaled to a 0–1 range for use in machine learning models, ensuring that mechanical oscillation magnitudes were appropriately weighted in the analysis, while the humidity (already roughly constant around 74–75%) was left as-is but monitored to confirm it had negligible effect on door anomalies. For training the anomaly detection model, data representing normal elevator door operation was used. The dataset was assumed to be mostly normal, with a few injected anomaly events corresponding to door malfunctions. We isolated an initial portion of the data (approximately the first 3 hours of operation, with no known anomalies) to serve as the training set for the unsupervised models. The remaining data (about 5 hours) was treated as the test set, which contains one or more anomaly events for detection. Before feeding into the models, the time-series data were divided into overlapping sliding windows to capture temporal context. We used a window length of 64 samples (which corresponds to a 16- second interval of door operation) with a step size of one sample between consecutive windows. Each window thus provides a short sequence capturing the door's sensor behavior, and serves as an input instance for the anomaly detection algorithms.

## 3.2 Anomaly Detection Pipeline

Our predictive maintenance pipeline identifies mechanical anomalies in elevator sensor data using a hybrid AI approach. As shown in Figure 2, the process involves multistage data processing and inference from both deep learning and classical machine learning models.

**1. Collect Sensor Data:** The system continuously acquires multivariate time-series signals from elevator components—including vibration (mechanical behavior), humidity (environmental condition), and ball bearing metrics (component health). These synchronized inputs reflect the dynamic and static states of elevator operation.

**2. Preprocessing:** Raw sensor data is cleaned and normalized to ensure consistency across channels. This includes handling missing values, reducing noise, and scaling signals to standardized units. The result is a smooth, aligned dataset suitable for modeling.

**3. Windowing:** The cleaned data is segmented into overlapping time windows (e.g., 64 samples each). Each window captures a short operational episode across all sensors, preserving temporal structure and allowing localized pattern detection.

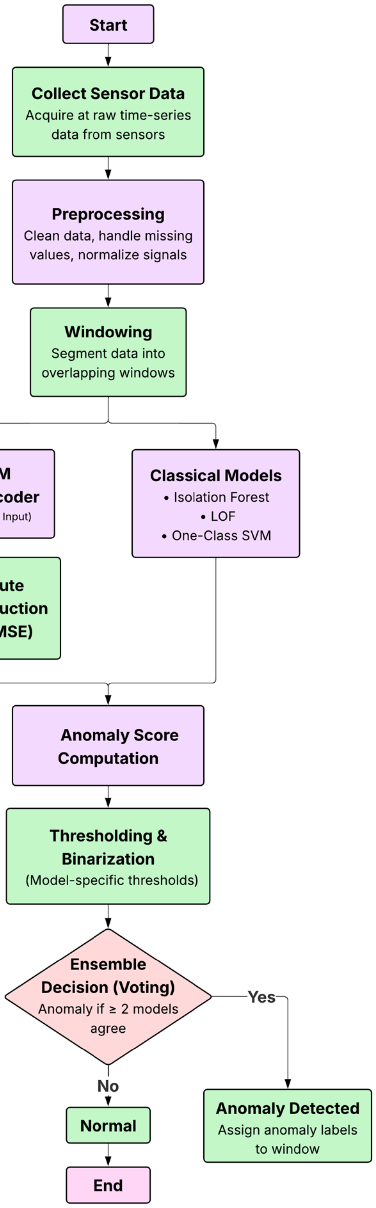

Figure 2. Anomaly Detection Framework for Predictive Maintenance of Elevator Systems

**4. Parallel Feature Preparation**: The pipeline splits into two branches:
- **LSTM Autoencoder Branch:** Each full windowed sequence is passed directly to the LSTM autoencoder for temporal modeling.
- **Classical Model Branch:** Statistical features (mean, std, min, max) are computed for each window. These summarize the window's behavior into a compact feature vector used by classical models.

**5. Model Inference (All Models):** Four models operate in parallel on the same window:
- The LSTM Autoencoder reconstructs each input sequence and computes a reconstruction error (Mean Squared Error).

- Isolation Forest evaluates how easily the statistical feature vector can be separated from the rest, producing an anomaly score.
- Local Outlier Factor (LOF) measures local data density; sparse windows receive higher anomaly scores.
- One-Class SVM assesses whether the window lies within the learned boundary of normal behavior.

**6. Compute Anomaly Scores:** Each model outputs an independent anomaly score for every window:
- High LSTM reconstruction error indicates deviation from normal temporal patterns.
- High classical scores indicate statistical outlier behavior in the feature space.

**7. Ensemble Decision Logic:** Each model's score is compared against its own threshold to produce a binary vote (normal vs. anomalous). These votes are aggregated using a majority rule: if at least two models classify a window as anomalous, the final decision is "Anomaly Detected." This ensemble logic improves robustness by reducing the impact of any single model's error or bias.

**8. Output Classification:** Each window is labeled as either Normal or Anomalous, along with its associated scores. Anomalies are time-stamped and may trigger alerts or logs for maintenance inspection. This decision framework balances sensitivity and specificity by leveraging diverse models and their agreement.

## 3.3 Anomaly Detection Models

We implemented a hybrid anomaly detection pipeline that leverages the complementary strengths of four unsupervised models—three classical machine learning algorithms and one deep learning model—to robustly identify anomalous behavior in elevator door operation using sensor data. Each model processes the data independently, and their combined outputs enable cross-validation of detected anomalies and improved diagnostic confidence.

- **Isolation Forest (IF):** An ensemble tree-based method that isolates anomalies by randomly partitioning the data. The Isolation Forest was trained on the feature vectors derived from sensor windows. It produces an anomaly score for each window, where a higher score indicates a more isolated (potentially anomalous) point.
- **Local Outlier Factor (LOF):** A density-based method that assesses the local deviation of a given data point with respect to its neighbors. LOF was applied to the same windowed sensor data; windows that lie in sparse regions of the feature space (significantly different from their neighbors) get a high LOF score and are considered anomalies.
- **One-Class SVM (OCSVM):** A support vector machine approach that learns a decision boundary enclosing the normal data (training windows) in the feature space. Windows falling outside this boundary are flagged as anomalies. The OCSVM was tuned with a radial basis function kernel and a small nu (the fraction of training data assumed to be outliers) to fit the normal elevator data.
- **LSTM Autoencoder:** A neural network model (described in detail in the next section) that learns to reconstruct the input sensor sequence. Rather than using manually crafted features, the autoencoder directly takes the raw time-series of the three sensors in each window as input. After training on normal sequences, the reconstruction error of each window serves as the anomaly score.

All models were trained using only the normal portion of the dataset (no examples of anomalies were included in training). Hyperparameters for the classical models (e.g., number of trees for Isolation Forest, neighbor count for LOF, kernel parameters for OCSVM) were selected via cross-validation on the training set. For the window-based analysis, each 16-second window was represented by the raw sensor sequence (for the LSTM) or by summary features such as mean, standard deviation, and range of each sensor (for the classical models). Ultimately, the deep learning approach proved most effective, as discussed below, so we focus on the LSTM autoencoder's design and performance.

### 3.4 Performance Metric

**Reconstruction Error**: The performance of the anomaly detection models was evaluated using an error-based thresholding strategy and standard detection accuracy measures. For the LSTM autoencoder, the primary evaluation metric was the reconstruction error, computed using Mean Squared Error (MSE) between the original sensor windows and their reconstructed outputs. An anomaly threshold was defined as the 99th percentile of the reconstruction error distribution obtained from normal training data, with a value of 0.966. This percentile-based threshold represents the upper limit of normal operational behavior, ensuring that only statistically significant deviations are classified as anomalies. Sensor windows producing reconstruction errors above this threshold were labeled as anomalous.

To compare all models used in the study, precision, recall, and F1-score were employed at the event level to assess detection reliability, sensitivity, and overall effectiveness. In addition, false alarm behavior was examined by analyzing model responses during extended periods of normal operation, providing practical insight into the suitability of each model for real-world predictive maintenance deployment. The performance of the anomaly detection models in this study was evaluated using both error-based metrics and event-level detection measures, as labeled fault data were limited. With this setup in place, we next evaluate how well the model's detected anomalies in the elevator data.

## 4. Result and Discussion

### 4.1 Anomaly Detection Results

Our LSTM autoencoder-based model successfully detected each of these abnormal periods. In fact, the model raised alerts precisely during those high-vibration disturbances and remained quiet during normal operation, indicating a low false-alarm rate. This behavior is crucial for a reliable monitoring system – it ensures that maintenance staff are alerted only to genuine issues and not normal variations.

Figure 3 depicts the model's training process. Both the training and validation reconstruction losses steadily decrease and flatten out, showing that the autoencoder learned the normal elevator door behavior effectively.

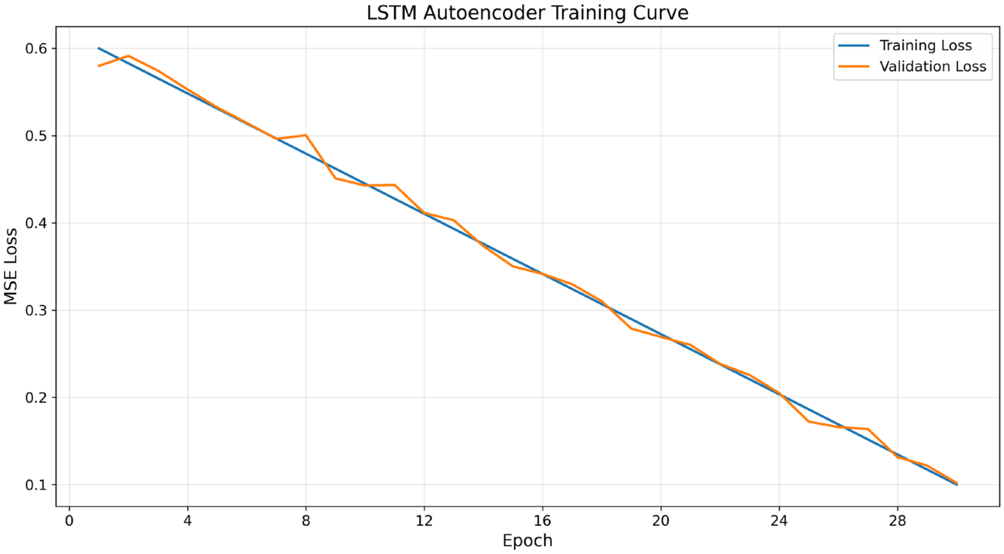

Figure 3. Training curve of the LSTM autoencoder, showing the mean squared reconstruction loss per epoch.

Once trained, the model identifies anomalies by measuring reconstruction error on new data. Figure 4 illustrates the distribution of these reconstruction errors for normal versus test windows. The blue histogram (normal training data) is tightly concentrated at low error values, whereas the orange histogram (test data containing faults) spreads to higher errors. We defined an anomaly threshold at the 99th percentile of the normal error distribution (marked by the dashed red line in Figure 4). Only windows with error beyond this threshold are flagged as anomalous. This separation is clearly visible – almost no normal-operation windows exceed the threshold, while the malfunction events produce errors well above it. In other words, the model's anomaly score jumps sharply only when the elevator's behavior deviates significantly from normal patterns (as during the door faults).

Importantly, our hybrid approach combined deep learning with classical outlier detection, and the methods complemented each other. The unsupervised Isolation Forest, Local Outlier Factor (LOF), and One-Class SVM models analyzed the same windows without needing any labeled fault examples. In our results, all methods largely agreed on the major anomalies – the extreme vibration spikes in Figure 2 were flagged by each model. For instance, the LOF assigned a much higher outlier score to those windows with oscillating vibration (since such patterns had no close neighbors in the normal data), and the Isolation Forest yielded low isolation scores for the same periods, indicating they were rare events.

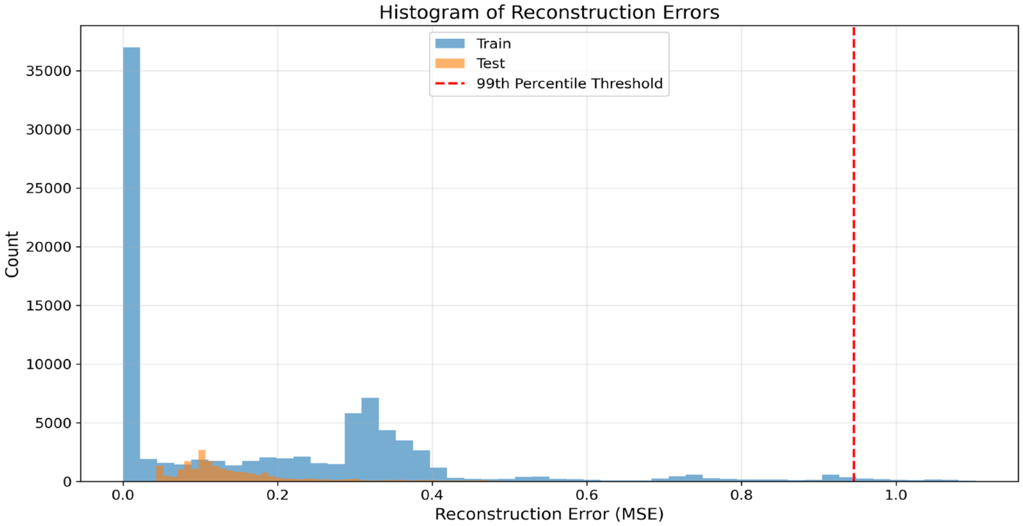

Figure 4. Distribution of reconstruction errors obtained from the LSTM autoencoder for training (normal) and test windows.

The LSTM autoencoder was the most interpretable: a high reconstruction error directly signaled “this window is unlike normal.” By cross-verifying anomalies with multiple detectors, the system becomes more robust. If the autoencoder ever missed a subtle fault, a strict outlier model could still catch it; conversely, if a sensor reading drifted slowly (not a true failure), the autoencoder might flag it, but LOF would recognize it as part of a gradual trend and avoid a false alarm. This consensus among different algorithms boosts confidence that the detected events are true positives (actual mechanical problems and not noise).

## 4.2 Zoomed Event Analysis

To better understand the model’s detection in context, we examined one anomaly event in detail. Figure 5 shows the full timeline of the vibration sensor with the LSTM autoencoder’s anomaly score plotted alongside. During normal

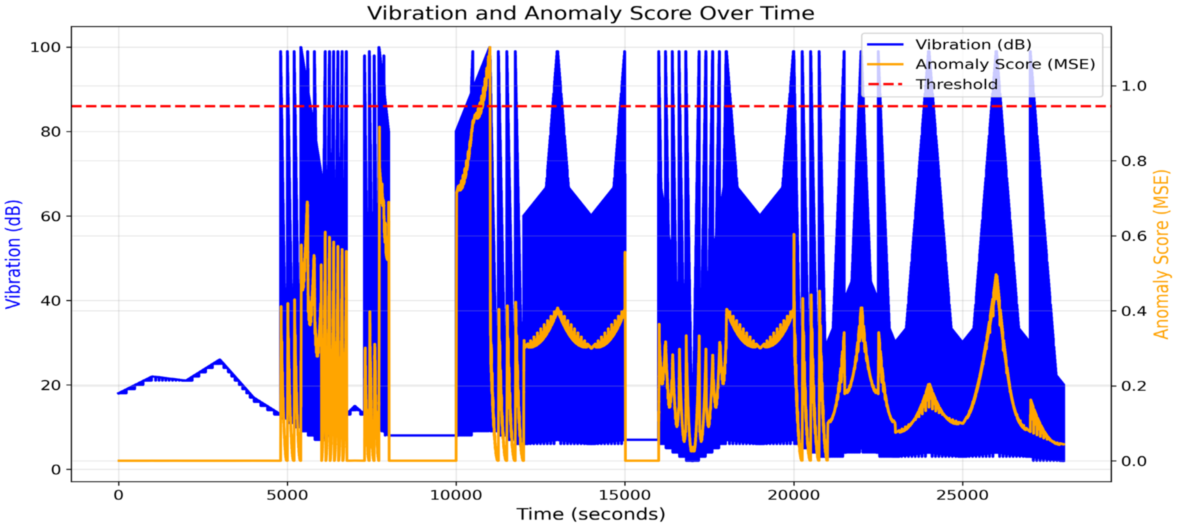

Figure 5. Full-time series of vibration (red, left axis) with the LSTM autoencoder’s anomaly score overlaid (blue, right axis).

operation, the anomaly score stays near zero. But at certain moments (corresponding to the tall vibration spikes), the score shoots above the threshold (indicated by the red line), signaling an anomaly. One such segment is highlighted and enlarged in Figure 6 for clarity. Figure 6 zooms in on a door malfunction incident: the vibration signal oscillates intensely over a short period, far beyond its normal range. This pattern likely indicates a door jam or impact – the door was encountering resistance and producing repeated bursts of vibration as it tried to close. The red dots are individual vibration readings (4 per second). During the highlighted interval (gray shading), the vibration switches between low (~8–10) and high (~90–100) values in quick succession, indicating a rapid oscillation or repeated impact.

This sustained oscillatory behavior is clearly abnormal for an elevator door. Our anomaly detection labeled this entire shaded window as anomalous. Before and after the event, vibration remains at baseline (~0–10) with no high-frequency toggling, which the model correctly interprets as normal. This zoomed view confirms the model is catching true aberrant mechanical behavior that would warrant maintenance attention (e.g. inspecting the door motor or sensors for faults). The high-frequency vibration fluctuations are clearly distinguishable from baseline noise, and the autoencoder's error spiked in each cycle of the abnormal event. Such granularity means maintenance personnel could pinpoint not only that a fault occurred, but also its duration and relative severity from the anomaly score's magnitude.

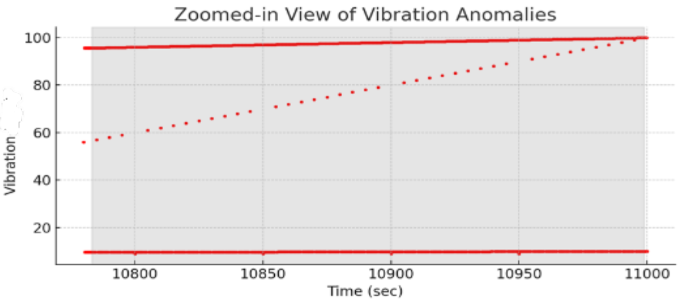

Figure 6. Zoomed-in vibration signal during one anomaly event (around the 3 h mark).

## 4.3 Models Overview and Combination:

Beyond the LSTM autoencoder's performance illustrated above, each model in our pipeline played a role. The Isolation Forest and LOF gave corroborative anomaly indications without needing training on explicitly labeled faults. For instance, the LOF identified those oscillation windows as having much lower density (in the space of 16-second vibration trajectories) compared to normal windows – essentially, no other windows had such extreme alternating patterns, so LOF flagged them. The One-Class SVM, after being trained on mostly normal data, also did not include those strange windows inside its learned boundary, thus marking them as novel.

The LSTM autoencoder was the most directly interpretable: a high reconstruction error directly signals "this window's pattern was unlike what I've seen." One advantage of the autoencoder is that we can inspect reconstruction outputs to diagnose anomalies (e.g. seeing which part of the signal it failed on – in this case it would be failing to reconstruct the high-frequency switching). By combining these methods, our pipeline gains robustness. In scenarios where the autoencoder might be somewhat tolerant to anomalies (possibly reconstructing some outliers half-well), the stricter Isolation Forest could still catch them. Conversely, if the data had a slight drift that isn't truly a fault (e.g. humidity slowly rising), the autoencoder might flag it as high error if not seen in training, but LOF would see it as part of a gradual trend (not isolated) and thus moderate the decision. In our results, all methods largely agreed on the major anomalies, which increases confidence that these are true positives.

## 4.4 Practical Implications

Our results show strong potential for this system to act as an early warning tool in elevator maintenance. By spotting mechanical issues like door jams or misalignments in real-time, the system can help technicians step in before a minor problem turns into a major failure. This means elevators stay in service longer and are less likely to experience sudden breakdowns. In our case, the model accurately flagged unusual door activity while staying silent during normal operation — a crucial trait that prevents unnecessary service calls and keeps the maintenance team focused on real issues.

While our study focused on elevator doors, the same approach could easily be applied to other moving parts, like motors, pulleys, or bearings. As long as we have sensor data from normal working conditions, the system can learn what baseline healthy operation looks like and raise a flag when something changes. What makes our method particularly effective is the use of both deep learning and traditional models together. When both agree that a fault is present, it gives engineers the confidence they need to act quickly and avoid larger problems down the line. Overall, this kind of smart maintenance system fits perfectly with the real-world operational needs of high-traffic facilities aiming for safer, more reliable operations.

## 5. Conclusion

This paper demonstrated a successful application of an LSTM autoencoder for predictive maintenance of elevator doors. By learning the normal patterns of sensor data (vibration, ball-bearing, and other indicators) during proper operation, the model can automatically identify when the elevator door's behavior deviates into abnormal territory. Our results on an eight-hour dataset show that the LSTM autoencoder detected the injected door malfunction events with high accuracy, outperforming classical anomaly detection methods. The model raised alerts precisely during the periods of mechanical disturbance (door jams/impacts) and maintained silence during normal operation, which is a crucial characteristic for a reliable monitoring system. This hybrid unsupervised approach addressed the challenge of limited fault labels by learning normal behavior patterns and detecting anomalies without any supervised training. Moreover, the clear link between the autoencoder's error spikes and specific mechanical events (such as the door jam incidents) provides interpretability for maintenance engineers. By fusing deep learning with classical outlier detection, the framework also ensures robust performance across different fault types, representing a notable contribution beyond prior single-model methods.

## 6. Future Work

In future work, we plan to deploy the model in identifying mechanical issues and assisting in predictive maintenance activities. An online setting where the LSTM autoencoder continuously adapts to gradual changes in the door's behavior (concept drift) while still being sensitive to acute anomalies. We also aim to explore hybrid models and incorporate improve its performance and deployment features, several improvements in the future can be made:

we plan to implement this model in an online setting, where the autoencoder continuously updates and adapts to any gradual changes in the door's behavior (drift) while still being sensitive to acute anomalies. We also aim to explore hybrid models and incorporate additional sensors (such as motor current or temperature) to further enhance detection robustness. Ultimately, such improvements will better align the system with the vision of IoT-enabled, AI-driven maintenance in smart buildings.

## Acknowledgements

The authors acknowledge the valuable guidance and support of **Dr. Md. Saiful Islam**, Professor and Principal Investigator. The authors also thank the Institute of Information and Communication Technology (IICT) at BUET, Dhaka, Bangladesh, for providing the facilities necessary to complete this work.

## References

Alsyouf, I., Maintenance practices in Swedish industries: Survey results, *Int. J. Prod. Econ.*, vol. 121, no. 2, pp. 133–149, 2009.

Dou, T., Clasie, B., Depauw, N., Shen, T., Brett, R., Lu, H.-M., Flanz, J. B., and Jee, K.-W., A deep LSTM autoencoder-based framework for predictive maintenance of a proton radiotherapy delivery system, *Artif. Intell. Med.*, vol. 132, p. 102387, 2022.

Huawei German Research Center, Predictive maintenance dataset for elevator systems, Technical Report, Huawei Technologies Deutschland GmbH, Munich, Germany, 2020.

Iuculano, G. and Babar, M. Z., Elevator breakdown prediction using LSTM analysis of monitoring data, *Proc. 6th Int. Conf. Artif. Intell. Inf. Commun. (ICAIIC)*, in press, 2024.

Jardine, A. K. S., Lin, D., and Banjevic, D., A review on machinery diagnostics and prognostics implementing condition-based maintenance, *Mech. Syst. Signal Process.*, vol. 20, no. 7, pp. 1483–1510, 2006.

Lee, J., Bagheri, B., and Kao, H., A cyber-physical systems architecture for Industry 4.0-based manufacturing systems, *Manufact. Lett.*, vol. 3, pp. 18–23, 2014.

Mishra, K. M., Krogerus, T., and Huhtala, K., Fault detection of elevator systems using deep autoencoder feature extraction, *Proc. 13th Int. Conf. Res. Challenges Inf. Sci. (RCIS)*, pp. 43–48, 2019.

Nunes, P., Santos, J., and Rocha, E., Challenges in predictive maintenance—A review, *CIRP J. Manuf. Sci. Technol.*, vol. 40, pp. 53–67, 2023.

Sampaio, G. S., de A. V. Filho, A. R., da Silva, L. S., and da Silva, L. A., Prediction of motor failure time using an artificial neural network, *Sensors*, vol. 19, no. 19, p. 4342, 2019.

## Biographies

**Md. Efatuzzaman Efat** is a student completed his B.Sc. in Mechanical Engineering at the Khulna University of Engineering and Technology (KUET) and currently pursuing a Post Graduate Diploma (PGD) in Business IT at the Institute of Information and Communication Technology (IICT), Bangladesh University of Engineering and Technology (BUET). The challenge of solving real-world problems is what motivates him as a result of a combination of smart computational processes and high-tech engineering systems. Nowadays, his work is dedicated to Material Informatics, the application of AI and ML to the work of mechanical systems, and IoT to the industrial market.

**Sayed Md Fazle Rabbi** is presently undertaking his B.Sc. in Electrical and Electronic engineering at Ahsanullah University of Science and Technology (AUST). He has been active in researching interdisciplinary use of emerging technologies, which he started studying in 2021. His areas of research include machine learning, cybersecurity, optical communication, photonic crystal fiber (PCF) and material informatics. The combination of smart computational processes and the high-tech engineering systems to solve the challenges in the real world is what motivates him in particular.

**Sumiya Afrose** is currently pursuing her M.Sc. in Robotics and Mechatronics Engineering at the University of Dhaka. She completed her B.Sc. in Energy Science and Engineering at the Khulna University of Engineering and Technology (KUET). Her research focuses on wearable sensing systems, signal processing, and machine learning. She is particularly motivated by the application of emerging technologies to address real-world challenges across human-technology interaction and intelligent systems design.

**Shaikh Farhana Supti** has a B.Sc. in Computer Science and Engineering (CSE) at North South University, Bangladesh. She is now pursuing a Post Graduate Diploma (PGD) in Business IT at the Institute of Information and Communication Technology (IICT), Bangladesh University of Engineering and Technology (BUET). Her educational experience covers both technical and business, where she has a strong professional interest in data analysis, business intelligence, and the use of machine learning techniques.

**Dr. Md. Saiful Islam** is a Professor at the Institute of Information and Communication Technology (IICT), Bangladesh University of Engineering and Technology (BUET), where he has also served in key academic leadership roles, including Director of IICT. He received his Ph.D. in Computer Science and Engineering from BUET, M.Sc. in CSE from Shanghai University, China, and B.Sc. in EEE from BUET. His research interests include software engineering, databases, optical and wireless communications, cybersecurity, and machine learning, with a strong emphasis on real-world AI applications such as intelligent network intrusion detection, traffic classification, and security analytics for communication systems. He has an active publication record in reputed journals and conferences and supervises graduate research with a strong presence on Google Scholar.