iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: https://doi.org/10.22214/ijraset.2026.81545

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

Deep Learning Based Time-Series Anomaly Detection Using LSTM Autoencoders

Arpit Agrawal[1], Ramnaresh Sharma[2]
Centre of Artificial Intelligence, Madhav Institute of Technology & Science, India

***Abstract:* Continuous monitoring of industrial and networked systems generates large volumes of timestamped sensor data, within which anomalies represent critical events requiring prompt identification. This paper introduces an unsupervised anomaly detection framework built upon an LSTM Autoencoder, trained exclusively on normal operational sequences. The architecture encodes input windows into a compressed latent representation and reconstructs them; deviations between input and output serve as anomaly indicators. A threshold derived from the mean reconstruction error plus two standard deviations governs classification, removing any dependency on labeled anomaly samples. Experimental training confirmed stable convergence, with reconstruction error remaining low for normal sequences and markedly elevated for anomalous ones, demonstrating the discriminative capacity of the proposed approach.**

***Keywords:* Anomaly Detection, LSTM Autoencoder, Time-Series, Unsupervised Learning, Reconstruction Error, Threshold-Based Classification.**

I. INTRODUCTION

Networked infrastructure spanning industrial control systems, cloud platforms, and healthcare monitoring devices continuously emits high-frequency, time-stamped measurements. Irregular patterns embedded within these streams may reflect equipment degradation, security incidents, or patient health changes, making their early detection operationally significant [4].

Rule-based alerting, though widely deployed, demands expert configuration of per-signal bounds and grows fragile as system behavior evolves. Supervised machine learning methods provide stronger adaptability but require balanced, labeled corpora that are rarely available when anomalous events are infrequent, varied, or entirely novel [1].

Long Short-Term Memory networks [3] excel at capturing multi-step temporal dependencies in sequential data, overcoming the limitations of conventional feature pipelines. When an LSTM encoder is paired with a symmetric decoder in an autoencoder architecture, the resulting system can characterize normal behavior from unlabeled data alone. Sequences that deviate from learned norms incur elevated reconstruction error and can therefore be identified without any anomaly annotation.

This paper presents an LSTM Autoencoder-based anomaly detection system with four primary contributions: (i) a preprocessing pipeline using min-max normalization and sliding-window segmentation; (ii) a symmetric LSTM encoder-decoder architecture; (iii) a statistically-derived detection threshold at mean plus two standard deviations of training reconstruction error; and (iv) a qualitative evaluation confirming clear separation between normal and anomalous error profiles. The proposed approach can be applied to real-world scenarios such as industrial monitoring, fraud detection, and system health analysis.

II. RELATED WORK

Chandola et al. [4] offer a foundational taxonomy of anomaly detection techniques, grouping approaches into classification-based, clustering-based, statistical, and information-theoretic categories. A central finding of their survey is that supervised methods underperform significantly when labeled anomaly instances are scarce, a condition that characterizes most real-world deployments.

Statistical approaches such as ARIMA capture linear temporal dependencies and are computationally efficient for stationary, low-dimensional signals. Their core assumptions, however, break down in the presence of non-linearity, high dimensionality, or non-stationarity, all of which are common in modern sensor data. Malhotra et al. [7] pioneered the use of LSTM networks for unsupervised anomaly detection by showing that prediction error at anomalous time steps consistently exceeds that at normal ones, even without any labeled training examples. This observation laid the conceptual groundwork for subsequent encoder-decoder formulations. Park et al. [8] extended this line of work by introducing a variational LSTM autoencoder that quantifies uncertainty in the latent space, yielding calibrated anomaly scores rather than binary classifications. Transformer-based models [9] have since pushed detection capability further for long-range dependencies, though their data and compute requirements favor the LSTM autoencoder for resource-constrained or data-limited scenarios such as the present study.

III. METHODOLOGY

A. Data Preprocessing

The dataset consists of univariate time-series observations recorded at uniform intervals, divided into a training partition containing only normal-regime samples and a test partition that includes both normal and anomalous segments. Min-max normalization is applied to rescale all values to the unit interval, ensuring no single channel disproportionately influences model training.

Overlapping windows of fixed length are produced using a sliding-window technique, transforming the continuous signal into a structured set of sequential samples. This approach captures local temporal context effectively and expands the number of training instances available from a single recording.

B. Model Architecture

The encoder comprises two stacked LSTM layers that progressively distill the input window into a compact latent representation. This bottleneck vector encodes the essential temporal characteristics associated with normal system behavior.

The decoder reconstructs the original window from the latent representation using a symmetric LSTM stack, with a time-distributed dense output layer generating the final reconstructed sequence. Mean Squared Error between the input and its reconstruction serves as the training loss.

Training is carried out over 100 epochs using the Adam optimizer. Early stopping based on validation loss is employed to halt training before overfitting occurs, ensuring that the model generalizes well to unseen normal sequences.

C. Anomaly Detection Threshold

Following training, the reconstruction error is computed for every window in the training set. The detection threshold θ is set at $\mu + 2\sigma$, where μ is the mean and σ is the standard deviation of those training-set errors. Any test window exceeding θ is flagged as anomalous. This approach is entirely data-driven, requires no labeled anomaly examples, and statistically constrains the false-positive rate on normal inputs.

IV. IMPLEMENTATION

The system is implemented in Python using TensorFlow and the Keras API for model construction and training. Data preprocessing and sequence generation rely on NumPy and Pandas, while Matplotlib supports visualization of training curves and reconstruction error distributions. All development and experimentation were carried out on Google Colab, providing a cloud-based notebook environment accessible without specialized local hardware.

V. RESULTS AND ANALYSIS

Table I presents the key observations recorded during model training and evaluation. Training loss declined consistently across all 100 epochs with no instability, confirming that the model successfully learned to reconstruct normal sequences. The model was evaluated on a publicly available time-series dataset containing both normal and anomalous segments.

Table I
Performance Metrics Of Lstm Autoencoder Model

Metric	Value / Observation
Training Loss (Final Epoch)	Decreased significantly
Reconstruction Error (Normal)	Low
Reconstruction Error (Anomaly)	High
Threshold Method	Mean + 2×Std Dev
Learning Type	Unsupervised
Model Convergence	Stable

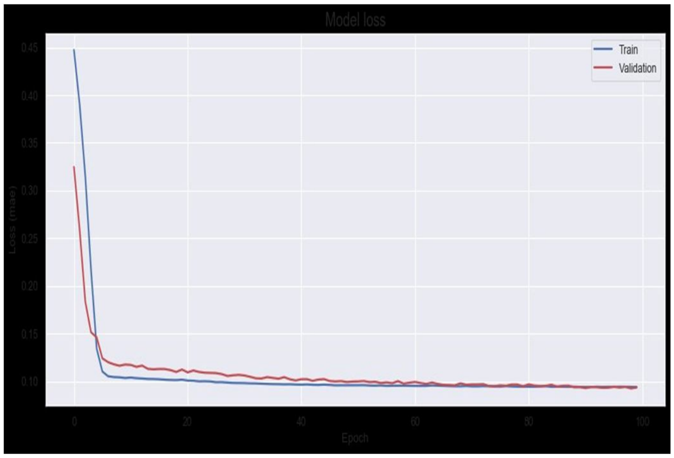

Fig. 1. Training and Validation Loss

Post-training evaluation revealed a clear and consistent gap between reconstruction errors of normal and anomalous windows. Normal sequences generated errors tightly concentrated near the training mean, while anomalous sequences produced errors substantially above the $\mu + 2\sigma$ threshold. This separation demonstrates that the encoder's latent space represents normal operational structure effectively and fails to compactly encode out-of-distribution patterns.

These findings are consistent with the unsupervised LSTM-based detection results reported by Malhotra et al. [7] and confirm the practical utility of reconstruction-error thresholding for anomaly identification in the absence of labeled data.

VI. CONCLUSION

This paper presented an LSTM Autoencoder trained exclusively on normal time-series data for unsupervised anomaly detection. Stable model convergence and a clear reconstruction-error gap between normal and anomalous sequences were observed, validating the core design premise. The $\mu + 2\sigma$ threshold provides an objective, label-free classification criterion suitable for diverse monitoring applications.

The current study is limited by the absence of labeled data, which restricts the use of standard evaluation metrics such as precision and recall. Future work will explore attention mechanisms to capture long-range anomaly signatures, multivariate extensions for correlated sensor systems, and adaptive threshold recalibration for non-stationary environments.

VII. ACKNOWLEDGMENT

The author thanks the project supervisor and the Centre of Artificial Intelligence, Madhav Institute of Technology & Science, for their guidance and computational support throughout this work. This paper is based on a student academic project undertaken as part of the undergraduate program at the institute.

REFERENCES

[1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
[2] F. Chollet, *Deep Learning with Python*. Manning Publications, 2017.
[3] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
[4] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly Detection: A Survey," *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, 2009.
[5] TensorFlow Development Team. TensorFlow. [Online]. Available: https://www.tensorflow.org
[6] F. Chollet et al. Keras Documentation. [Online]. Available: https://keras.io
[7] P. Malhotra, A. Ramakrishnan, G. Anand, L. Vig, P. Agarwal, and G. Shroff, "LSTM-based Encoder-Decoder for Multi-sensor Anomaly Detection," in *Proc. ICML Anomaly Detection Workshop*, 2016.
[8] D. Park, Y. Hoshi, and C. C. Kemp, "A Multimodal Anomaly Detector for Robot-Assisted Feeding Using an LSTM-based Variational Autoencoder," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1544–1551, 2018.
[9] J. Xu, H. Wu, J. Wang, and M. Long, "Anoformer: Unsupervised Anomaly Detection via Transformer," in *Proc. IJCAI*, 2022.