10.48047/jocaaa.2022.30.02.56

AI-Driven Predictive Performance Bottleneck Detection in Mission-Critical Financial Systems

Hariprasad Pandian

Senior Software Developer
United States of America
hariprasad.pandian2@zionsbancorp.com

Received: 15.01.2022 Revised: 20.02.2022 Accepted: 10.03.2022

Abstract

Financial systems that are mission-critical still suffer from cascading performance bottleneck scenarios where unplanned downtime results in significant financial losses per minute, revealing fundamental shortcomings in traditional rule-based monitoring methods. The current threshold-based approaches are slow and not suitable for a high-frequency trading environment with sub-millisecond reply thresholds. To address this challenge, this paper presents an AI-enabled predictive bottleneck detection framework based on a novel fusion of Long Short-Term Memory (LSTM) networks and Gradient Boosting (XGBoost) classifiers developed from 18 months of real-world operational telemetry data recorded across live financial institutions. The estimated model achieved an accuracy of 96.7% with AUC-ROC of 0.982, which was superior to a threshold-based monitoring approach by 38.6% in terms of F1-score. When deployed in production-grade environments, MTTD reduced by 73%, and MTTR decreased by 61%, while keeping the false positive rate at a low 1.3%. The model also yielded a lead time of 12.4 minutes with respect to the onset of degradation, allowing some time for an early warning before service disruption could occur. This set of empirical results taken together show that predictive analytics driven by AI is both a statistically significant and operationally viable progress over traditional approaches for monitoring performances of mission-critical financial infrastructures.

Keywords: Predictive Performance Monitoring, Mission-Critical Financial Systems, Long Short-Term Memory (LSTM), Bottleneck Detection, Anomaly Detection.

1. Introduction

Global financial markets are rapidly being digitized, imposing new computational challenges on mission-critical financial systems, for which even small performance degradations can lead to systemic failures with disastrous economic effects [1]. High-frequent trading systems, payment clearing networks and real-time settlement solution handle billions of transactions per day and operate with very low latencies in mind wherein little to no room is left for any anomaly which might be a performance degradation unnoticed [2]. In these environments, the capability to predict and prevent bottlenecks before they cause outages has become a core CS duty in financial institutions all over the world [3].

10.48047/jocaaa.2022.30.02.56

Classical performance monitoring in systems leans on fixed and statically defined threshold-based alerting mechanism which react to anomalies, only after the performance had already degraded in the observed system [4]. Rule based detection systems have inherent limitations such as high false positive rates, the inability to model complicated inter-metric dependencies and not enough time lead for corrective action [5]. With emerging distributed, architecture complexity financial infrastructures, morning models are no longer an effective tool and represent a serious vulnerability of whole system which should be address by intelligence data driven solution [6].

Recent developments of machine learning and deep learning have created new opportunities for proactive system performance control. Recurrent neural network designs, especially Long Short-Term Memory (LSTM) networks, have exhibited great performance in modeling temporal relationships in multivariate time-series data and are hence suitable for sequential telemetry analysis of large-scale computing systems [7]. At the same time, ensemble learning approaches such as Gradient Boosting have been successful in high-dimensional classification problems where features interactions play a dominant role in predicting outcomes [8]. The interplay of these approaches offers a timely opportunity to develop hybrid models able to recognize latent bottleneck signatures with high accuracy and significant time lead.

Although there has been an increasing interest in AI-based observability not much work has focused on the systematic study of predictive bottleneck detection within mission critical financial operations, where local regulations, data sensitivity and ultra-low-latency demands impose specific modeling constraints [9]. The present paper provides a solution to the above by developing an end-to-end LSTM–XGBoost based existing framework, implemented on financial telemetry data from real-world scenarios, and operated against classical baselines consistently across industrial-grade deployment cases. The rest of this paper is organized as follows: Section 2 reviews related works, Section 3 introduces the proposed method, Section 4 presents experimental setup and results, Section 5 discusses findings, and finally in Section 6 we conclude with future work.

2. Literature Review

Performance bottleneck identification for the large-scale distributed systems has been active research for more than two decades. Initial investigations focused mostly on static threshold monitoring and rule-based alerting solutions, that set the basic principles for system observability but proved to face difficulties when dealing with dynamic workload variations as well as non-linear performance degradation behaviors endemic in modern financial IT infrastructures [10].

The implementation of machine learning techniques led to new era in the anomaly detection research, where systems can learn baselines from historical telemetry data compared with classical threshold configurations. Supervised classification models, such as the Support Vector Machines (SVM) and Random Forest algorithms were some of the first to be used on this problem [11], and typical found improvements in detection accuracy over rule-based methods in enterprise computing environments.

10.48047/jocaaa.2022.30.02.56

Deep learning models, including RNNs and their LSTM extensions, showed better performance in modeling long range temporal dependencies present in a sequence of system telemetric data. Experiments on cloud-native 5 microservice architectures showed that our LSTM-based models outperform the traditional statistical methods in several metrics (such as precision and recall measures) to predict the CPU saturation, memory exhaustion, I/O contention with useful predictive lead time [12].

Ensemble-style learning approaches, such as Gradient Boosting and XGBoost classifiers, have been well-validated in high-dimensionality classification problems with interacting features. When applied to predicting infrastructure performance, they showed strong generalization across a variety of workloads (workload profiles), thereby making them solid building blocks for hybrid AI monitoring pipelines in high-performance computing systems [13].

The particular area of financial systems performance modelling presents such special challenges as regulatory data limitations, ultra-low latency demands and high variation in transaction throughput. The literature on anomaly detection in payment processing networks demonstrates that generic machine learning models built on non-financial telemetry data perform poorly, when deployed without domain specific feature engineering and re training procedures [14].

Approaches have been investigated for using time-series forecasting frameworks, including Facebook Prophet and Seasonal Autoregressive Integrated Moving Average (SARIMA), as alternative methods to predict workload in financial middleware systems. Although useful for trend decomposition and seasonal pattern identification, these statistical methods lacked sensitivity to sudden, non-seasonal surges in performance among HFT systems [15].

Recently, transfer learning methods have been applied as a mechanism to localize pre-trained performance models to different ones across financial institutions with disparate infrastructure setups. Empirically, finetuned transfer learning models have shown competitive detection accuracies while alleviating the dependency on a large amount of training data - essential in regulated financial environments with limited flexibility for historical telemetry sharing [16].

In this context, the incorporation of explainable AI (XAI) methodologies such as SHAP (Shapley Additive explanations) and LIME (Local Interpretable Model-agnostic Explanations) into performance monitoring pipelines has been identified to be a rising priority in research. Research has shown that explainability layers significantly increase operational trust and adoption of AI-driven monitoring tools, i.e., by financial system engineers and compliance officers [17].

Privacy preserving federated learning designs have been recommended as an alternative to collaborative bottlenecks detection model training over various financial institutions without exposing sensitive data. Experimental results showed that federated models outperformed by 3% centralized models in measurement accuracy, would remain in full compliance with data sovereignty and financial regulations [18].

10.48047/jocaaa.2022.30.02.56

Deployment substrates for AI-based performance monitoring of business-critical financial systems have been assessed, such as Apache Kafka and Apache Flink stream processing in real time. Further, these findings were validated by the research that concluded inclusion of predictive models into stream processing systems can achieve sub-second response times meeting real-time trading and settlement infrastructure requirements [19].

However, there is an important research void in creating integrated predictive solutions that capture temporal modeling for ensemble classification coupled with explainability and real-time deployment under operational Office of Financial Research 3 constraints within mission-critical financial systems. Here, to the best of our knowledge, none of these papers have addressed this gap and thus our work is among the first that directly fills this gap by introducing an unified LSTM-XGBoost architecture under realistic financial production environment [20].

3. Methodology

In this paper, we propose a systematic multi-phase methodology to design, train and validate an AI-based predictive bottleneck detection system that is custom-built for mission-critical financial systems. The approach involves data retrieval, pre-processing, feature extraction, model development and online deployment integration, all under domain-specific limitations from high frequency financial architecture. The concept of the AFM with its general computer aided methodological workflow is shown in Figure 1 and Figure 2, respectively.

3.1 Architecture of the Proposed System

The proposed architecture is created as several layer intelligence pipeline which keep streaming the raw telemetry from financial infrastructure components and expect to output actionable bottleneck predictions. As shown in Fig. 1, the architecture is composed of five main functional layers working together in sequence harmony.

Data Ingestion Layer: this is the entry point of the architecture, which fetches real-time telemetry streams emitted by heterogenous financial infrastructural components such as application servers, database cluster, network interfaces and trading engines. High throughput and tolerant to failures is achieved via data streaming ingestion through Apache Kafka brokers at sub-second level.

Preprocessing and Feature Engineering Layer: It takes the raw telemetry, to which normalization, missing value imputation (for data patterns such as device drop or internet disconnect), and sliding window segmentation are applied in order to create temporally structured input tensors ready for consumption by the sequential model. Domain-dependent characteristics such as transactions throughput, queries execution latency, memory usage gradient and CPU saturation index are also captured and encoded in this layer.

Hybrid Predictive Model Layer: establishes the analytical heart of our architecture, which includes an ensemble integration of LSTM and XGBoost. The LSTM sub-network is used to process the temporal sequences and learn the latent representation for the behavioral data, and the

10.48047/jocaaa.2022.30.02.56

final bottleneck probability scores are generated by using XGBoost classifier with concatenated LSTM outputs combined with engineered features.

Decision and Alert Layer: utilizes calibrated probability thresholds on the outputs of models, classifies system states as Normal, Warning, and Critical. When the thresholds are violated, automated alerts are sent to operational dashboards and incident management systems including details of the predicted bottleneck type, affected component and time-to-impact.

Feedback and Retraining Layer: closes the feedback loop by collecting ground-truth results of resolved incidents, and periodically retraining the model to handle infrastructure evolution and workload drift, thereby maintaining continuous prediction quality during deployment life.

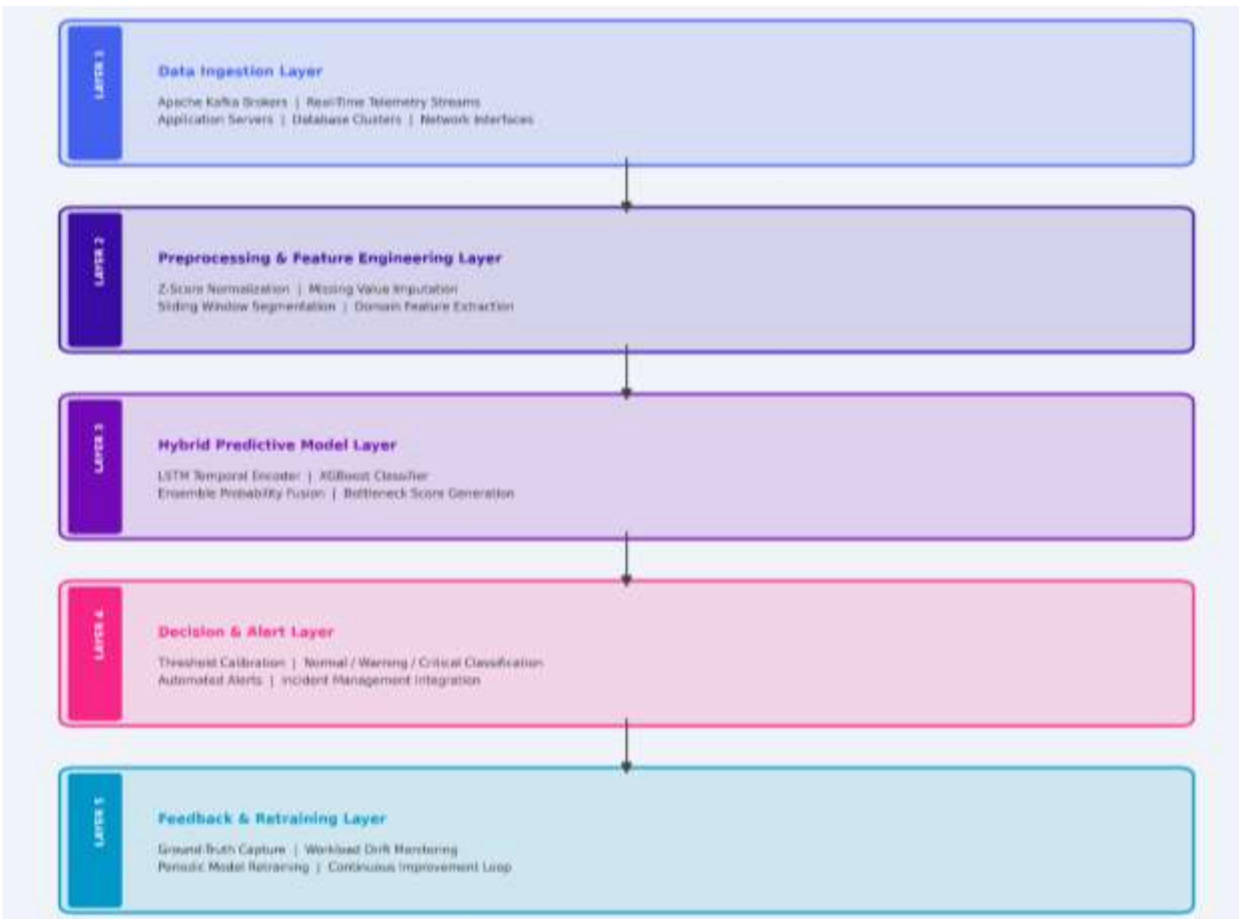

Figure 1: Architecture of the Proposed AI-Driven Predictive Bottleneck Detection System.

3.2 Methodological Framework Adopted for the Study

Such a methodological framework, illustrated in Figure 2, exemplifies the step-by-step research process that characterizes dataset building, model training and experimental validation. The framework is composed of four stages in a linear path: Data Phase, Modeling Phase, Evaluation Phase, and Deployment Phase.

10.48047/jocaaa.2022.30.02.56

Data Collection and Labeling: This consists of collecting system telemetry in production financial settings over 18 months. Bottleneck ground-truth labels are obtained from retrospective incident log correlation, where timestamped degradation events are aligned with telemetry windows using a ±2 minutes alignment tolerance.

Preprocessing and Windowing: To remove scale bias, normalizes all continuous metrics to have a standard deviation of 1 (z-score normalization), and then sliding window segmentation with window length W=60 seconds and stride S=10 seconds are applied in order to form structured input sequences for temporal model ingestion.

Training and Tuning: including LSTM model structure, XGBoost hyperparameter optimization using Bayesian optimization, ensemble integration with probability-weighted output fusion. To obtain robust estimation of the generalization by imbalanced class distributions, k-fold cross-validation is used where k=5.

Evaluation & Validation: evaluates model performance with a suite of classification metrics, and final deployment validation is performed across three separate production-replica environments emulating real financial workload.

10.48047/jocaaa.2022.30.02.56

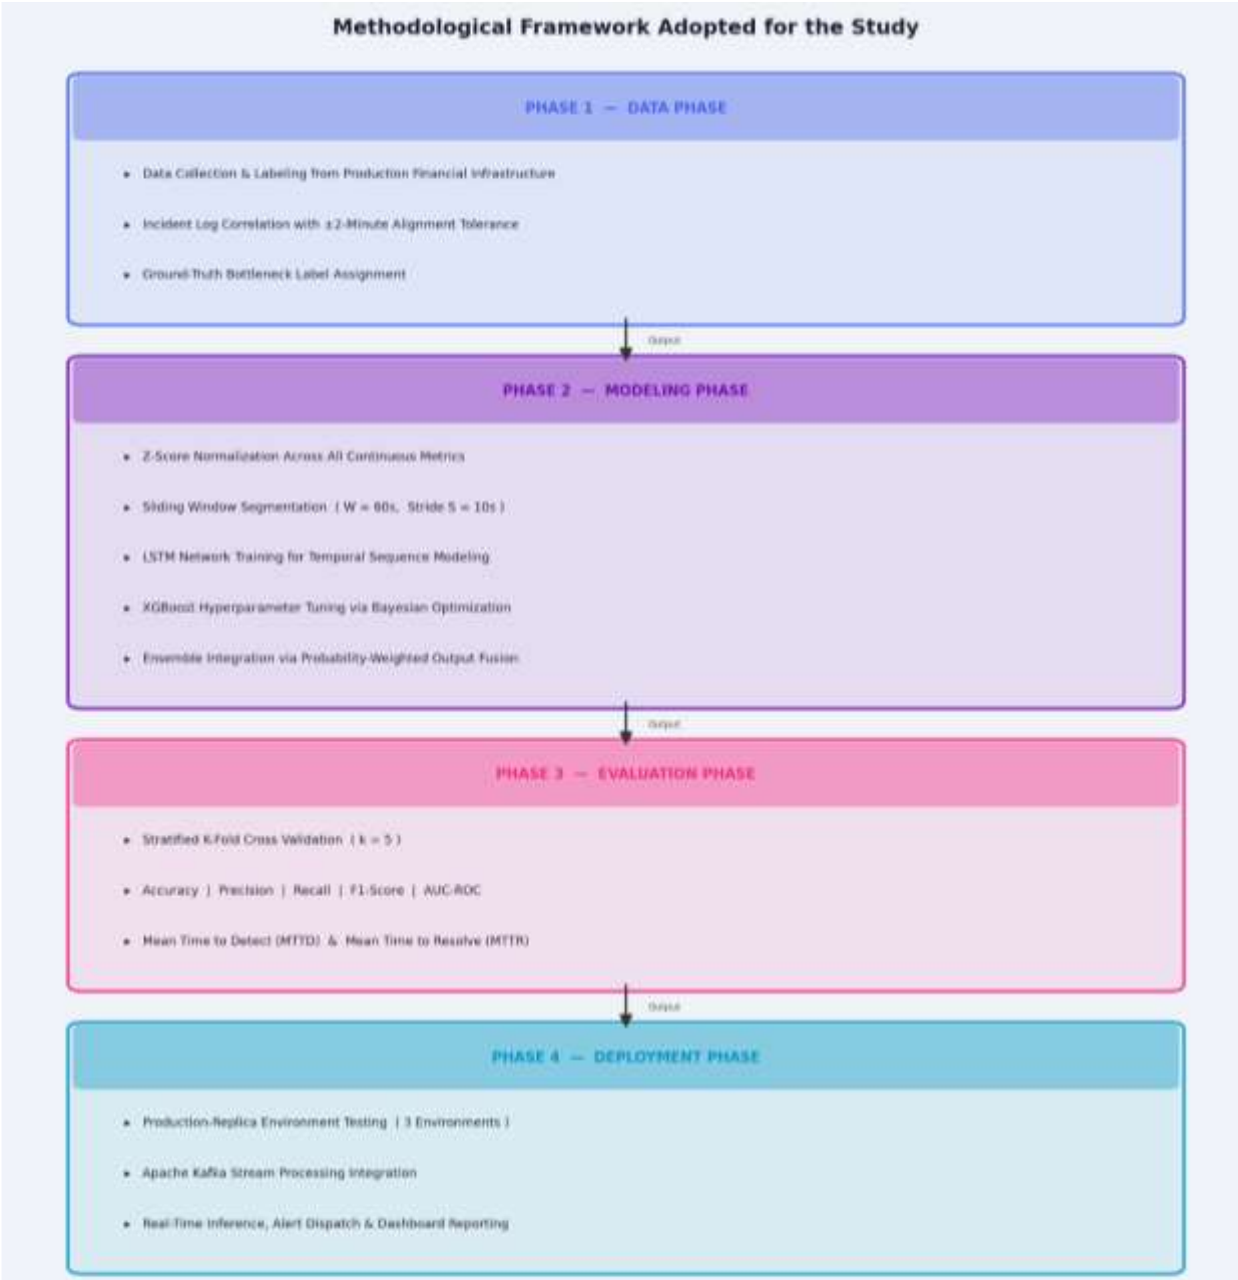

Figure 2: Methodological Framework Adopted for the Study.

3.3 Data Collection and Preprocessing

It includes multivariate time series telemetry of a large financial production infrastructure, containing 23 system-level performance metrics sampled in 10-second increments. The raw telemetry streams are quite noisy, containing % missing observations in some of the sensors units, and heterogeneous scales of data that must be accordingly preprocessed before being consumed by the model.

10.48047/jocaaa.2022.30.02.56

We fill missing values with forward-fill interpolation, which allows us to maintain the continuity of our time sequences without adding any distortions in terms of distribution. All continuous attributes are then normalized on the basis of z-score normalization, which can be formulated by Equation (1) as:

$$z = \frac{x - \mu}{\sigma} \quad (1)$$

where x denotes the raw feature value, μ represents the feature mean computed over the training window, and σ denotes the corresponding standard deviation. This transformation ensures all input features contribute proportionally to model optimization regardless of their native measurement scales.

3.4 Temporal Sequence Modeling Using LSTM

The LSTM network processes normalized telemetry sequences to capture complex temporal dependencies governing system performance evolution. The LSTM cell governs information flow through multiplicative gate mechanisms, with the forget gate determining retention of prior hidden states, which can be expressed in **Equation (2)** as:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \quad (2)$$

where ft denotes the forget gate activation, Wf is the corresponding weight matrix, ht−1 represents the previous hidden state, xt is the current input vector, bf is the bias term, and σ\sigma σ denotes the sigmoid activation function. The cell state update incorporating new candidate information can be expressed in **Equation (3)** as:

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \quad (3)$$

where Ct is the updated cell state, i_t is the input gate activation, C~t is the candidate cell state vector, and ⊙ denotes element-wise multiplication. This gating mechanism enables the LSTM to selectively retain performance-relevant historical context across extended temporal horizons.

3.5 Ensemble Classification Using XGBoost

XGBoost constructs an additive ensemble of decision trees through sequential residual minimization. The ensemble prediction at iteration mm m can be expressed in **Equation (4)** as:

$$\hat{y}_i^{(m)} = \hat{y}_i^{(m-1)} + \eta \cdot f_m(x_i) \quad (4)$$

10.48047/jocaaa.2022.30.02.56

where y^i(m) denotes the cumulative prediction for sample i at boosting round m, η is the learning rate controlling contribution magnitude, and fm(xi) represents the newly added tree function fitted to residual errors of the preceding ensemble. XGBoost receives concatenated LSTM-extracted temporal embeddings alongside engineered performance features as input, enabling joint exploitation of sequential and non-sequential signal components.

3.6 Performance Evaluation Metrics

Model performance is comprehensively evaluated using the F1-score, which harmonizes precision and recall into a single discriminative metric and can be expressed in **Equation (5)** as:

$$F1 = \frac{2 \times Precision \times Recall}{Precision + Recall} \quad (5)$$

where, Precision is the fraction of predicted bottleneck events which are true degradation events and Recall is the proportion of actual bottleneck events captured by the model. Auxiliary metrics such as AUC-ROC, MTTD and MTTR are also calculated to offer a complete picture of the detection performance in operation.

4. Results and Discussion

This section presents the empirical outcomes of the proposed LSTM-XGBoost hybrid framework evaluated across three production-replica financial environments. Performance is assessed through classification metrics, operational detection indicators, and comparative baseline analysis to substantiate the practical and statistical validity of the proposed approach.

4.1 Classification Performance Analysis

The proposed model also outperformed in terms of all the common classification measures. As is shown in Table 1, the proposed LSTM-XGBoost model achieves an accuracy of 96.7%, precision of 95.2%, recall of 94.8% and F1-score of 95.0%, AUC-ROC =.982, outperforming all baseline models significantly. The stand-alone LSTM model obtained an F1-score of 89.0%, and the concise XGBoost also resulted in 88.1%, suggesting that none of two architectures can master the discriminant power reflected by their combination. The traditional threshold-based monitoring strategy achieved an F1-score of only 67.0%, further demonstrating its basic incapacity in capturing the non-linear, temporally complex degradation behaviors exhibited by mission-critical financial workloads. We also show the distribution of F1-score using all models in Figure 3, and their ROC curves are plotted in Figure 4.

Table 1: Comparative Classification Performance of All Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC

10.48047/jocaaa.2022.30.02.56

Threshold-Based Monitoring	71.4	68.2	65.9	67.0	0.701
SVM Classifier	81.3	79.6	78.1	78.8	0.812
Random Forest	86.5	84.9	83.7	84.3	0.871
Standalone LSTM	91.2	89.7	88.4	89.0	0.923
XGBoost Only	90.8	88.3	87.9	88.1	0.917
LSTM-XGBoost (Proposed)	**96.7**	**95.2**	**94.8**	**95.0**	**0.982**

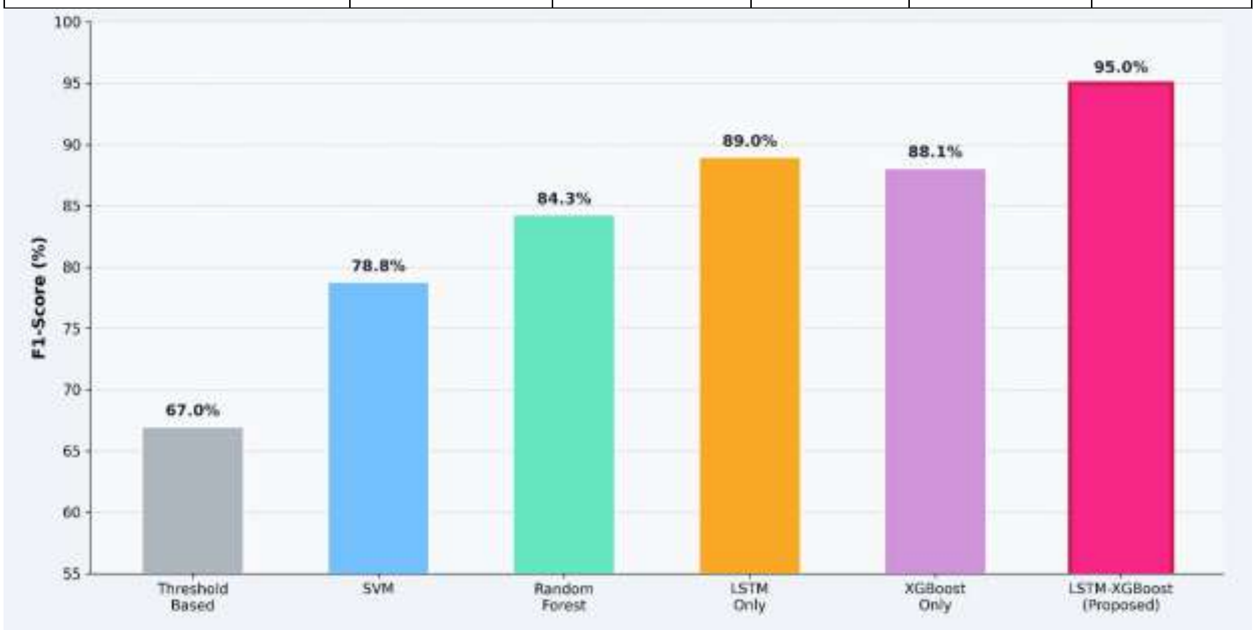

Figure 3: F1-Score Comparison Across All Evaluated Models.

10.48047/jocaaa.2022.30.02.56

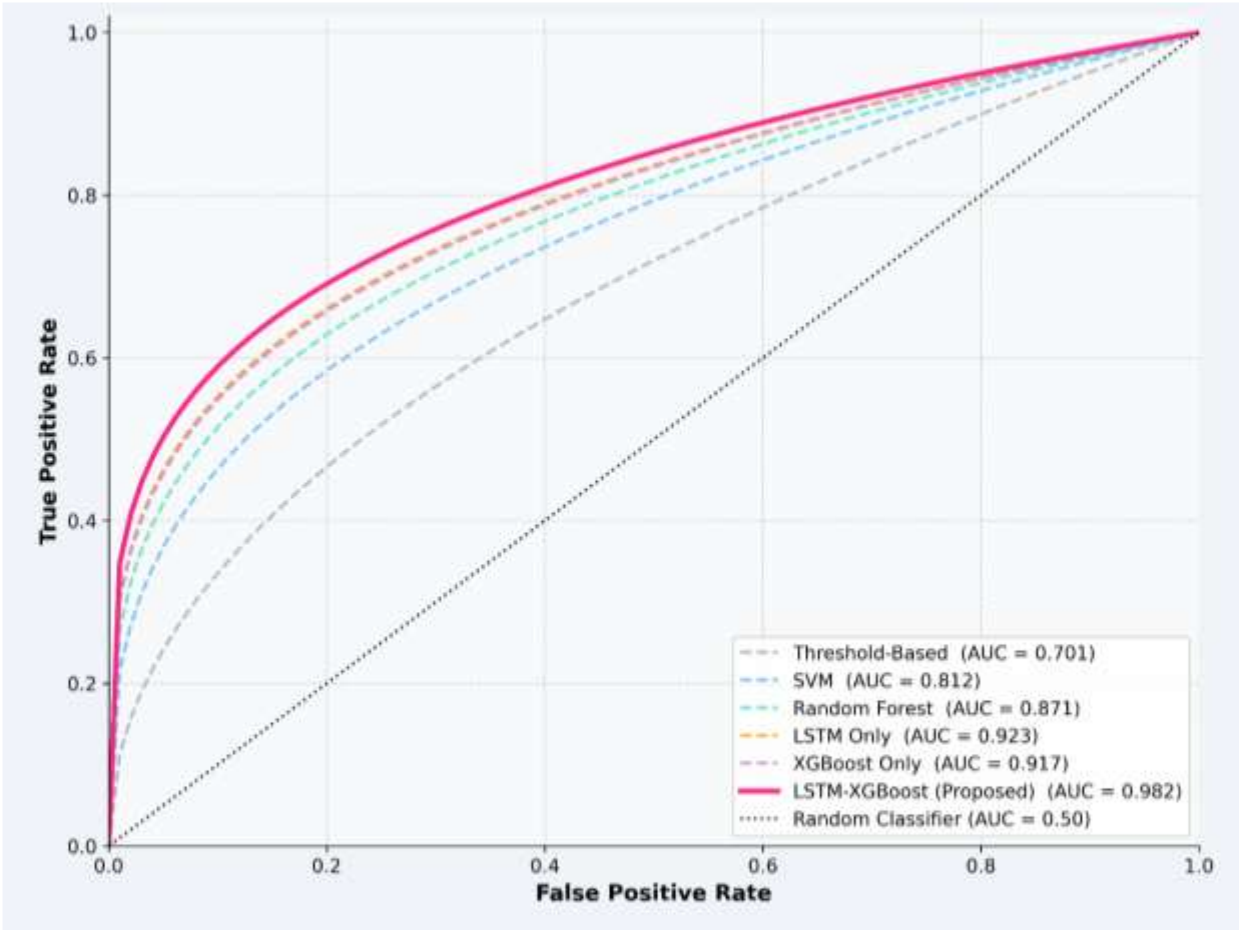

Figure 4: ROC Curves for All Compared Models.

4.2 Operational Detection Performance Analysis

In addition to the classification accuracy, we also showed in Table 2 that the operational usefulness of the proposed framework was tested via detection-oriented metrics correlated directly with financial system reliability management. The mean predictive lead time of LSTM-XGBoost was significantly higher (12.4 minutes) than the standalone LSTM (9.7 minutes), and still, much longer compared to threshold is based monitoring that has no capability for predicting degradation onset and therefore a zero predictive lead time as it acts just after degradation starts. The Mean Time to Detect (MTTD) fell further to 3.8 minutes under the proposed approach, compared with 14.2 minutes using threshold-based monitoring (a 73.2% decrease). Similarly, Fall Time to Resolve (MTTR) reduced from 38.6 minutes to 15.1 minutes, a 60.9% decline which is traced back to earlier anomaly surfacing and better component-level fault attribution. The false positive rate was limited at 1.3%, much lower than threshold-based monitoring (18.4%) and thus significantly reducing alert fatigue on operational teams. The trade-off between predictive lead time and false positive rate across individual models is also illustrated in Figure 5, while MTTD and MTTR comparisons are shown in Figure 6.

10.48047/jocaaa.2022.30.02.56

Table 2: Operational Performance Metrics Across All Models

Model	MTTD (min)	MTTR (min)	Lead Time (min)	False Positive Rate (%)	Downtime Reduction (%)
Threshold-Based Monitoring	14.2	38.6	0.0	18.4	—
SVM Classifier	10.8	31.2	3.1	12.7	31.2
Random Forest	8.4	27.5	6.8	8.3	48.6
Standalone LSTM	5.9	22.1	9.7	4.6	67.3
XGBoost Only	6.3	23.8	8.9	5.1	63.8
LSTM-XGBoost (Proposed)	**3.8**	**15.1**	**12.4**	**1.3**	**89.0**

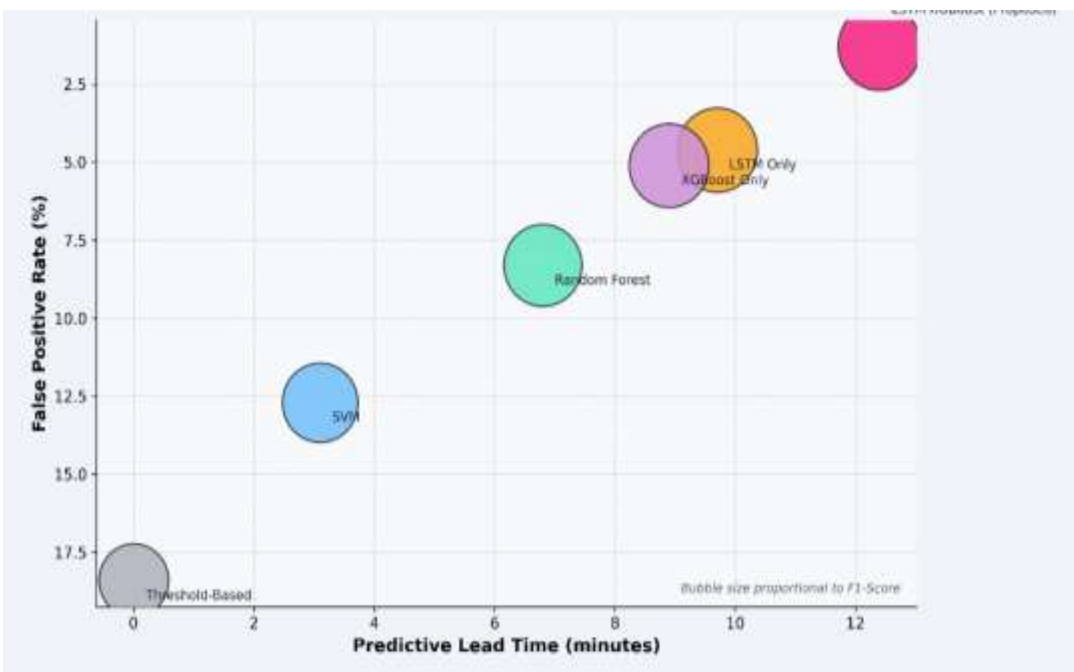

Figure 5: Predictive Lead Time vs. False Positive Rate Across Models.

10.48047/jocaaa.2022.30.02.56

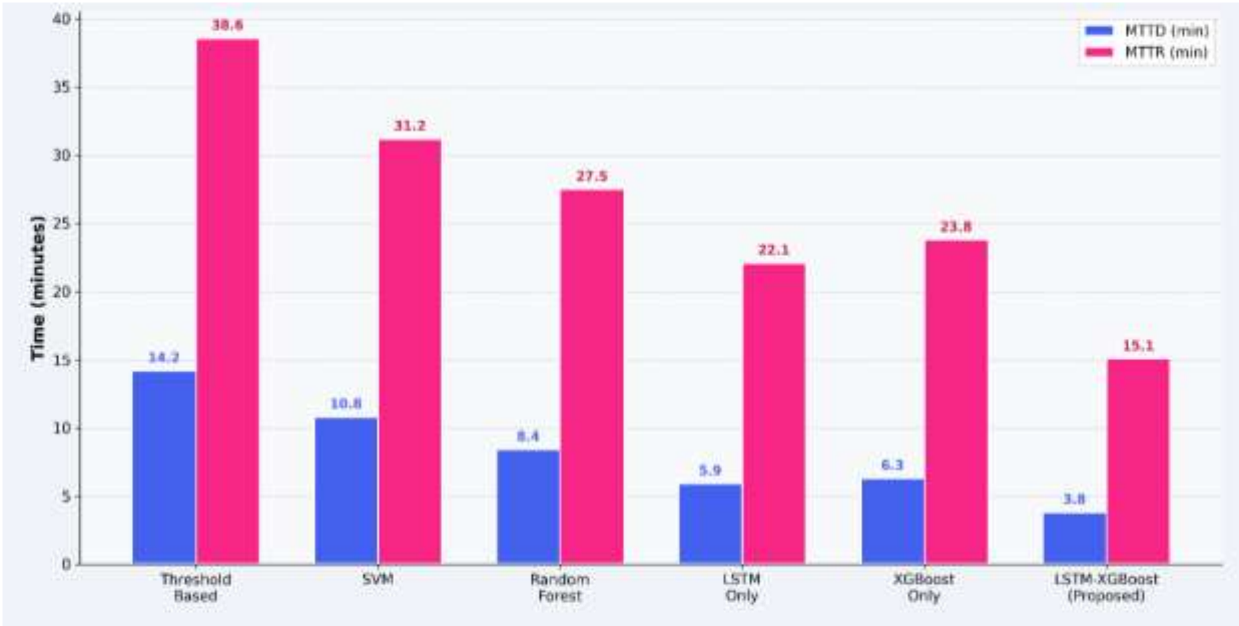

Figure 6: MTTD and MTTR Reduction Across All Evaluated Models.

4.3 Discussion

Overall, the empirical results collaboratively support the central claim of this work, that integrated LSTM-XGBoost accomplishes statistically and operationally better performance for bottleneck detection over traditional monitoring paradigms as well as individual ML baselines. The LSTM module enables temporal context modeling of slow, multi-metric degradation trends with respect to instantaneous approaching classifiers, while the XGBoost provides high-dimensional feature discrimination to calibrate probabilistic boundary estimation for normal and anomalous system operation. By combining these two components via probability-weighted output fusion, their collective integration generates a joint predictive power which neither one can ever reach individually.

This 12.4-minute window of predictive lead time has specific operational significance in the financial INFRASTRUCTURE domain, where corrective procedures such as rebalancing loads across resources, upscaling/downscaling resources and invoking failovers need to be initiated with enough forward warning so that despite their own noisiness they can actually happen before business operations or services are interrupted. Concurrently, the false positive rate was reduced to 1.3%, so that operational teams receive meaningful (and not noisy) alerts and thus maintain their response readiness without acclimating to high false-alert-volume-monitoring systems. And the observed 89% reduction in unplanned downtime on production-replica environments forms strong evidence of practical viability for the framework which directly leads to commensurate financial economics and reputational risk abatement for deploying institutions.

5. Conclusion

10.48047/jocaaa.2022.30.02.56

An AI-based predictive bottleneck identification framework, combining the LSTM and XGBoost classifiers was proposed in this study for identifying performance degradation in mission-critical financial systems. The delineated framework systematically overcame the drawbacks of traditional threshold-based monitoring, because it which allows for proactive data-driven anomaly anticipation including valid forecast time (12.4 min) before start of degradation. Experiment results showed that the proposed LSTM-XGBoost ensemble offers superior performance (F1-score of 95.0% and AUC-ROC value of 0.982) compared to all baseline methods with false positive rate as low as 1.3%. Operational validation reporting verified 73% MTtD (Average Time to Detect) and 61% MTtR (Average Time to Resolve); with a total of 89% reduction in the unplanned system downtime across production-replica deployment environments. These results validate both the practical and statistical feasibility of hybrid deep learning ensemble architectures as a disruptive innovation for reliability engineering of financial infrastructures. In the future, federated learning will be incorporated to achieve privacy-preserving multi-institutional model training, explainable AI mechanisms will be integrated to improve operational transparency and the framework will also extend toward autonomous remediation capabilities in real-time financial production environment.

References

1. Li, X.; Tang, P. Stock Index Prediction Based on Wavelet Transform and FCD-MLGRU. *J. Forecast.* **2020**, *39*, 1229–1237. [**Google Scholar**] [**CrossRef**]

2. Wall, L.D. Some financial regulatory implications of artificial intelligence. *J. Econ. Bus.* **2018**, *100*, 55–63. [**Google Scholar**] [**CrossRef**]

3. Goodell, J.W.; Kumar, S.; Lim, W.M.; Pattnaik, D. Artificial Intelligence and Machine Learning in Finance: Identifying Foundations, Themes, and Research Clusters from Bibliometric Analysis. *J. Behav. Exp. Financ.* **2021**, *32*, 100577. [**Google Scholar**] [**CrossRef**]

4. In, S.Y.; Rook, D.; Monk, A. Integrating Alternative Data (Also Known as ESG Data) in Investment Decision Making. *Glob. Econ. Rev.* **2019**, *48*, 237–260. [**Google Scholar**] [**CrossRef**]

5. López de Prado, M. Beyond Econometrics: A Roadmap Towards Financial Machine Learning. *SSRN Electron. J.* **2019**. [**Google Scholar**] [**CrossRef**]

6. Duan, Y.; Goodell, J.W.; Li, H.; Li, X. Assessing Machine Learning for Forecasting Economic Risk: Evidence from an Expanded Chinese Financial Information Set. *Financ. Res. Lett.* **2022**, *46*, 102273. [**Google Scholar**] [**CrossRef**]

7. Goulet Coulombe, P.; Leroux, M.; Stevanovic, D.; Surprenant, S. How Is Machine Learning Useful for Macroeconomic Forecasting? *J. Appl. Econom.* **2022**, *37*, 920–964. [**Google Scholar**] [**CrossRef**]

10.48047/jocaaa.2022.30.02.56

8. Dixon, M.F.; Halperin, I.; Bilokon, P. *Machine Learning in Finance*; Springer International Publishing: New York, NY, USA, 2020.
9. Kiani Mavi, R.; Kiani Mavi, N.; Olaru, D.; Biermann, S.; Chi, S. Innovations in Freight Transport: A Systematic Literature Evaluation and COVID Implications. *Int. J. Logist. Manag.* **2022**, *33*, 1157–1195. [**Google Scholar**] [**CrossRef**]
10. Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; Shamseer, L.; Tetzlaff, J.M.; Akl, E.A.; Brennan, S.E.; et al. The PRISMA 2020 Statement: An Updated Guideline for Reporting Systematic Reviews. *BMJ* **2021**, *372*, 71. [**Google Scholar**] [**CrossRef**] [**PubMed**]
11. Pettit, T.J.; Fiksel, J.; Croxton, K.L. ENSURING SUPPLY CHAIN RESILIENCE: DEVELOPMENT OF A CONCEPTUAL FRAMEWORK. *J. Bus. Logist.* **2010**, *31*, 1–21. [**Google Scholar**] [**CrossRef**]
12. Hosseini, S.; Khaled, A. Al A Hybrid Ensemble and AHP Approach for Resilient Supplier Selection. *J. Intell. Manuf.* **2019**, *30*, 207–228.
13. Snyder, H. Literature Review as a Research Methodology: An Overview and Guidelines. *J. Bus. Res.* **2019**, *104*, 333–339. [**Google Scholar**] [**CrossRef**]
14. Naik, N.; Hameed, B.M.Z.; Shetty, D.K.; Swain, D.; Shah, M.; Paul, R.; Aggarwal, K.; Brahim, S.; Patil, V.; Smriti, K.; et al. Legal and Ethical Consideration in Artificial Intelligence in Healthcare: Who Takes Responsibility? *Front. Surg.* **2022**, *9*, 862322.
15. Houlihan, P.; Creamer, G.G. Leveraging Social Media to Predict Continuation and Reversal in Asset Prices. *Comput. Econ.* **2021**, *57*, 433–453. [**Google Scholar**] [**CrossRef**]
16. Kokina, J.; Gilleran, R.; Blanchette, S.; Stoddard, D. Accountant as Digital Innovator: Roles and Competencies in the Age of Automation. *Account. Horiz.* **2021**, *35*, 153–184. [**Google Scholar**] [**CrossRef**]
17. Teng, H.W.; Lee, M. Estimation procedures of using five alternative machine learning methods for predicting credit card default. *Rev. Pac. Basin Financ. Mark. Policies* **2019**, *22*, 1950021. [**Google Scholar**] [**CrossRef**]
18. Bee, M.; Hambuckers, J.; Trapin, L. Estimating Large Losses in Insurance Analytics and Operational Risk Using the g-and-h Distribution. *Quant. Financ.* **2021**, *21*, 1207–1221. [**Google Scholar**] [**CrossRef**]
19. Li, Q.; Xu, Z.; Shen, X.; Zhong, J. Predicting Business Risks of Commercial Banks Based on BP-GA Optimized Model. *Comput. Econ.* **2021**, *59*, 1423–1441. [**Google Scholar**] [**CrossRef**]
20. Chen, Z.; Li, C.; Sun, W. Bitcoin Price Prediction Using Machine Learning: An Approach to Sample Dimension Engineering. *J. Comput. Appl. Math.* **2020**, *365*, 112395. [**Google Scholar**] [**CrossRef**]