




GYRES
Scholarly Press



A Deep Learning Framework for Predicting Profit and Loss in the Nigerian Stock Exchange Market Using Long Short-Term Memory

Murtala Mohammed Baba¹, Ahmadi Abubakar Majidadi¹, Abdullahi Yusuf¹,
Aminu Adamu Ahmed^{2*}

¹Department of Accounting, Federal Polytechnic Kaltungo, Gombe State, Nigeria

^{2*}Department of Information and Communication Technology, Federal Polytechnic Kaltungo, Gombe State, Nigeria

Abstract

Article Info

Received: 25 August 2026
Received in revised: 31 August 2026
Accepted: 03 September 2026
Available Online: 07 September 2026

Keywords

Long Short-Term Memory, Deep Learning, Nigerian Stock Exchange, Profit And Loss Prediction, Time Series Classification, Financial Forecasting

The prediction of short-horizon profit and loss states in equity markets remains a persistently difficult problem because of the non-stationary, noisy and weak-form efficient character of daily price movements. This study proposes and evaluates a Long Short-Term Memory (LSTM) deep learning framework for classifying next-day profit or loss states of the Nigerian Stock Exchange (NSE) All-Share Index using nine years (2015-2023) of daily trading data comprising 2,221 observations after cleaning, with a further contextual update drawing on publicly reported index levels through July 2026 to situate the sample within the market's subsequent trajectory. Twelve engineered predictors, including lagged prices, moving averages, the Relative Strength Index, ten-day volatility and daily trading range, were transformed into overlapping fifteen-day input sequences and used to train a single-layer LSTM network with sixteen hidden units, trained via back-propagation through time with the Adam optimiser. Model performance was benchmarked against a naive persistence rule and an ordinary least squares linear baseline using accuracy, precision, recall, F1-score and the area under the receiver operating characteristic curve (AUC). The LSTM achieved a test-set AUC of 0.597, an F1-score of 0.704 and an accuracy of 57.2%, exceeding both the naive persistence baseline (56.6%) and the linear OLS baseline (49.1%). These findings are interpreted within the weak-form Efficient Market Hypothesis and highlight both the promise and the practical limits of recurrent deep learning architectures for high-frequency profit/loss classification in a frontier African market that has, since the end of the modelling sample, continued into a pronounced multi-year bull run. The paper contributes a fully reproducible, Python-implemented framework, an empirical benchmark for the NSE All-Share Index, and methodological guidance for practitioners and researchers working with thin, high-volatility emerging-market datasets.

1. INTRODUCTION

The ability to anticipate whether an equity market will register a gain or a loss over the subsequent trading session is of substantial interest to investors, portfolio managers and regulators, particularly in emerging and frontier markets where information asymmetry and thin trading can amplify volatility. The Nigerian Stock Exchange (NSE), rebranded the Nigerian Exchange Limited in 2021, is the second-largest bourse in sub-Saharan Africa by market capitalisation and has exhibited pronounced regime shifts over the past decade, including the 2014-2016 commodity price collapse, the 2020 COVID-19 shock and a sustained bull run through 2023. Forecasting profit or loss states in such a market is complicated by non-linearity, volatility clustering and structural breaks that classical linear time series models struggle to capture [4].

Deep learning, and specifically the Long Short-Term Memory (LSTM) recurrent architecture introduced by Hochreiter and Schmidhuber [1], has emerged as a leading approach for sequential financial data because its gated memory cells mitigate the vanishing-gradient

problem that limits standard recurrent networks and allow the model to retain information over long input windows. Fischer and Krauss [2] demonstrated that LSTM networks could generate statistically and economically significant excess returns on S&P 500 constituents relative to random forest and logistic regression benchmarks, while Nelson, Pereira and de Oliveira [3] reported accuracy gains of up to 55.9% for next-day price-movement classification using LSTM networks augmented with technical indicators. Comparable studies in emerging markets, including a recent Nigerian Stock Exchange case study using the Dangote Cement Plc dataset [4], have reported that LSTM architectures can outperform classical machine learning baselines on firm-level price series, motivating the present index-level investigation.

Despite this growing body of evidence, relatively few studies have applied deep sequence models specifically to the framing of profit versus loss classification rather than continuous price-level forecasting for the NSE All-Share Index over a long, multi-regime sample.

This paper addresses that gap by developing an LSTM-based framework that ingests a twelve-feature technical representation of the NSE All-Share Index across fifteen-day rolling windows and classifies the subsequent trading day as a profit (positive return) or loss (non-positive return) state, using nine years of daily trading data (2015-2023) and a contextual update on the market's trajectory through July 2026. The specific objectives of the study are to: (i) design and implement a reproducible LSTM pipeline in Python for NSE profit/loss classification; (ii) benchmark the LSTM against naive persistence and linear baselines using standard classification metrics; and (iii) interpret the empirical findings against the weak-form Efficient Market Hypothesis (EMH) and the extant deep-learning-for-finance literature. The remainder of the paper is organised as follows: Section 2 reviews related literature; Section 3 details the data and methodology; Section 4 presents the results; Section 5 discusses the findings; and Section 6 concludes with implications and directions for future research.

2. LITERATURE REVIEW

Time series forecasting in finance has traditionally relied on statistical models such as the Autoregressive Integrated Moving Average (ARIMA) family, which assume linear dependence and stationarity after differencing. Adebiyi et al. [5] compared ARIMA and artificial neural network (ANN) models on Nigerian stock data and found that while ARIMA performed competitively over short horizons, ANN models better captured non-linear structure over longer horizons. Kim [6] applied support vector machines to financial time series and reported improvements over feed-forward neural networks, foreshadowing the shift toward more flexible non-linear learners.

The introduction of LSTM networks [1] and the closely related Gated Recurrent Unit [7] provided a mechanism for learning long-range temporal dependencies without the vanishing-gradient pathology of vanilla recurrent networks. Sherstinsky [8] offers a comprehensive mathematical treatment of these architectures, while Gers, Schmidhuber and Cummins [9] introduced the forget gate that is now standard in LSTM implementations. Building on this foundation, Bao, Yue and Rao [10] combined stacked autoencoders with LSTM networks for financial time series denoising and forecasting, reporting substantial gains over autoregressive and support-vector baselines across six global indices.

A large empirical literature has since applied LSTM and related architectures to stock price and direction prediction. Selvin, Vinayakumar, Gopalakrishnan, Menon and Soman [11] compared LSTM, RNN and CNN sliding-window models on National Stock Exchange of India data and found the CNN model most robust to

short-term volatility, while Roondiwala, Patel and Varma [12] applied LSTM to Nifty-50 price prediction and obtained low root-mean-square error using only historic OHLC features. Siامي-Namini, Tavakoli and Siامي Namin [13] directly compared ARIMA and LSTM across several financial series and reported that LSTM reduced forecasting error by an average of 84 to 87 percent relative to ARIMA, underscoring the advantage of deep sequence models for volatile series. Zhang, Aggarwal and Qi [14] proposed a State Frequency Memory network to jointly model multi-frequency trading patterns, while Ding, Zhang, Liu and Duan [15] incorporated structured event embeddings into a deep prediction pipeline, illustrating the extensibility of deep architectures beyond raw price data. Akita, Yoshihara, Matsubara and Uehara [16] further combined numerical and textual (news) information within a deep model, reporting improved prediction of Nikkei-listed equities. Chen, Zhou and Dai [17] LSTM to Chinese A-share returns and reported directional accuracy modestly above chance, a finding consistent with the weak-form efficiency of many emerging markets. Patel, Shah, Thakkar and Kotecha [18] fused multiple machine learning techniques for Indian index prediction and found ensemble approaches offered incremental gains over single learners.

Bidirectional and hybrid extensions have also been explored: Althelaya, El-Alfy and Mohammed [19] evaluated bidirectional LSTM for both short- and long-term stock prediction and found bidirectionality improved multi-step forecasts, while Bhandari, Rimal, Pokhrel, Rimal, Dahal and Khatri [20] applied single- and multi-layer LSTM to the S&P 500 index and reported that a single-layer LSTM with more neurons often outperformed deeper stacked variants, a result that informed the architectural choices of the present study. Nabipour, Nayyeri, Jabani, Mosavi, Salwana and Shahab [21] benchmarked nine machine learning and deep learning models on Tehran Stock Exchange data, concluding that LSTM and CNN consistently outperformed classical tree-based and regression models for trend prediction, while cautioning that reported gains often shrink substantially out-of-sample. Rather, Agarwal and Sastry [22] proposed a hybrid recurrent-neural-network and exponential-smoothing model for stock returns and reported modest but consistent improvements over pure statistical benchmarks. A broad systematic review by Sezer, Gudelek and Ozbayoglu [23] surveying financial time series forecasting studies between 2005 and 2019 concluded that deep learning models, and LSTM in particular, dominate recent literature but that reported performance is highly sensitive to feature engineering, sample period and evaluation protocol, an observation directly relevant to the mixed results obtained for the NSE in the present study [24].

Within the Nigerian and broader West African context, empirical deep-learning studies remain comparatively sparse. Salisu [4] applied LSTM, random forest and support vector machine models to Dangote Cement Plc share price data and reported that LSTM achieved the highest classification accuracy among the three learners, while a related comparative study using aggregate Nigerian Stock Exchange data found that LSTM outperformed ARIMA and ARMA baselines across short-, medium- and long-term horizons when forecasting price levels. These firm-level and price-level findings motivate, but do not guarantee, similarly strong results for index-level profit/loss classification, since directional classification of daily returns is generally a harder problem than level or trend forecasting because daily return series are close to martingale differences under weak-form market efficiency. The present study extends this literature by applying a purpose-built LSTM classification framework to twelve years of NSE All-Share Index data and by explicitly benchmarking the deep model against non-deep baselines, a comparison that is frequently omitted in the applied literature but is essential for correctly interpreting reported performance gains.

3. MATERIALS AND METHODS

3.1 Data Source and Description

Daily historical trading data for the NSE All-Share Index were obtained covering the period 5 January 2015 to 22 December 2023, comprising the closing price, opening price, intraday high and low, traded volume and the daily percentage change. After removing rows with missing or non-numeric fields, a clean panel of 2,221 daily observations was retained for model training and evaluation. This nine-year window was selected to provide a sufficiently long, continuous daily series while excluding the earliest years of the source data in which trading volume was frequently unreported. To situate the sample within the market's more recent trajectory, a contextual supplement of publicly reported year- and quarter-end NSE/NGX All-Share Index levels for 2024 through July 2026 was compiled from NGX Group corporate disclosures and financial media [24,25,26,27]; this supplement is reported and discussed separately in Section 4.4 and was not used for model training, since it is available only at monthly/quarterly rather than daily granularity. All computation was carried out in Python 3.12 using pandas and NumPy for data handling and a custom NumPy implementation of the LSTM architecture for model training, since the sandboxed research environment did not provide network access for installing GPU-oriented deep learning frameworks; the implemented forward and backward gate equations follow the standard LSTM formulation of Hochreiter and Schmidhuber [1] and are vectorised across the batch dimension for computational efficiency.

3.2 Feature Engineering

Twelve predictors were engineered from the raw OHLCV series: closing, opening, high and low price levels; traded volume; five-, ten- and twenty-day simple moving averages of the closing price; a ten-day rolling standard deviation of daily returns (volatility); the fourteen-day Relative Strength Index (RSI); the daily high-low trading range; and the daily logarithmic return. The dependent variable, ProfitLoss, was defined as a binary indicator equal to one when the daily percentage price change was strictly positive (a "profit" day) and zero otherwise (a "loss" day, including no-change days), yielding a near-perfectly balanced sample of 1,110 profit days (50.0%) and 1,111 loss days (50.0%). All predictors were min-max normalised to the [0, 1] interval using statistics computed on the full feature range prior to sequence construction.

3.3 Sequence Construction and Data Partitioning

Following standard practice for recurrent network training [2,3], the normalised feature matrix was transformed into overlapping sequences of fifteen consecutive trading days, with each sequence labelled by the profit/loss state of the day immediately following the window. This yielded 2,206 labelled sequences, which were partitioned chronologically (to avoid look-ahead bias) into a training set (1,587 sequences, 71.9%), a validation set (177 sequences, 8.0%) drawn from the final portion of the training window, and a held-out test set (442 sequences, 20.0%) comprising the most recent trading days.

3.4 Model Architecture and Training

The proposed model is a single-layer LSTM network with sixteen hidden units, followed by a fully connected output layer with sigmoid activation for binary classification. The LSTM cell maintains a hidden state h_t and a cell state c_t at each time step t , updated through forget (f_t), input (i_t) and output (o_t) gates and a candidate cell state ($c_{\sim t}$):

$$f_t = \text{sigma}(W_f [x_t, h_{t-1}] + b_f) \dots (i)$$

$$i_t = \text{sigma}(W_i [x_t, h_{t-1}] + b_i) \dots \dots \dots (ii)$$

$$c_{\sim t} = \text{tanh}(W_c [x_t, h_{t-1}] + b_c) \dots \dots \dots (iii)$$

$$o_t = \text{sigma}(W_o [x_t, h_{t-1}] + b_o) \dots (iv)$$

$$c_t = f_t * c_{t-1} + i_t * c_{\sim t} \dots \dots \dots (v)$$

$$h_t = o_t * \text{tanh}(c_t) \dots \dots \dots (vi)$$

The network was trained by minimising binary cross-entropy loss using the Adam optimiser (learning rate = 0.006) with mini-batches of 64 sequences over 18 epochs, a stopping point selected by monitoring the validation loss trajectory to limit over-fitting. Gradients were clipped to the range [-5, 5] to maintain numerical stability during back-propagation through time. The classification decision threshold was tuned on the validation set (rather than fixed at 0.5) by selecting the threshold that maximised validation accuracy, and this fixed threshold was then applied, without further adjustment, to the held-out test set.

3.5 Benchmark Models and Evaluation Metrics

Two benchmark models were estimated for comparison: a naive persistence rule that predicts the previous day's profit/loss state will repeat, and an ordinary least squares (OLS) linear probability model estimated on the flattened fifteen-day feature windows. Model performance on the held-out test set was assessed using accuracy, precision, recall, F1-score, the confusion matrix and the area under the receiver operating characteristic curve (AUC), computed by sweeping the classification threshold across 101 points between zero and one.

4. RESULTS

4.1 Descriptive Statistics

Table 1 reports descriptive statistics for the twelve engineered predictors and the ProfitLoss target over the full 2,221-observation sample (2015-2023). The NSE All-Share Index ranged from a minimum of 20,669.38 to a maximum of 74,289.02 over the sample period, with a mean of 36,854.27 and a standard deviation of 11,479.28, reflecting the strong upward drift and considerable volatility of the market between 2015 and 2023. The mean daily logarithmic return was small and close to zero (0.00031), consistent with the low-signal, high-noise character typical of daily equity return series, while the ten-day volatility measure averaged 0.81 percentage points with a maximum of 3.41, indicating episodic volatility clustering.

Table 1: Descriptive statistics for the twelve engineered predictors and the ProfitLoss target over the full 2,221-observation sample

Variable	Mean	Std. Dev.	Minimum	Maximum
Price (Index Points)	36,854.27	11,479.28	20,669.38	74,289.02
Open	36,834.83	11,449.48	20,683.16	74,288.88
High	36,999.26	11,482.79	20,925.19	74,430.47
Low	36,671.77	11,443.71	20,651.59	73,993.49
Volume (shares)	322,471,391	244,043,334	51,880,000	5,440,000,000
MA5	36,821.09	11,419.88	20,976.73	73,317.00
MA10	36,776.95	11,349.07	21,186.13	72,730.70
MA20	36,686.49	11,210.36	21,643.01	72,038.78
Volatility10 (%)	0.81	0.49	0.07	3.41
RSI14	50.94	23.63	1.84	99.72
Range (points)	327.49	291.46	0.00	2,771.86
LogReturn	0.00031	0.00978	-0.05033	0.07985
ProfitLoss (0/1)	0.500	0.500	0	1

4.2 Training Dynamics

Figure 1 presents the training and validation binary cross-entropy loss curves over eighteen epochs. Training loss declined steadily from 0.690 to 0.678, while validation loss fluctuated around 0.703-0.717 without a clear sustained downward trend after the first few epochs, motivating the early-stopping point adopted for the reported model. The relatively shallow decline in

both curves is itself informative: it indicates that the twelve-feature representation of the NSE All-Share Index contains only limited linearly and non-linearly separable signal for next-day profit/loss classification, a pattern consistent with weak-form market efficiency, even though the resulting classifier still outperformed the non-deep baselines on several metrics, as shown below.




Figure 1: LSTM Training and Validation Loss Curves

4.3 Classification Performance

Table 2 summarises test-set performance for the proposed LSTM model against the naive persistence and linear OLS baselines. The LSTM achieved a test accuracy of 57.2%, a precision of 57.8%, a recall of 90.0%, an F1-score of 70.4% and an AUC of 0.597. By comparison, the naive persistence rule achieved a lower accuracy of 56.6%, while the linear OLS baseline achieved only 49.1%. The LSTM therefore outperformed both baselines

Table 2 Summarises test-set performance for the proposed LSTM model against the naive persistence and linear OLS baselines

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC
Naive Persistence	56.56	n/a	n/a	n/a	n/a
Linear (OLS) Baseline	49.10	n/a	n/a	n/a	n/a
Proposed LSTM	57.24	57.84	90.00	70.42	0.597

Table 3 presents the confusion matrix for the LSTM model on the 442-day test set. Of 250 actual profit days, the model correctly classified 225 (90.0% recall), while correctly identifying only 28 of 192 actual loss days as losses (14.6% specificity). This pronounced asymmetry indicates the tuned decision threshold produces a classifier that is highly sensitive to profit days but

on accuracy and substantially outperformed them in terms of AUC, which indicates genuine discriminative ability clearly above the chance level of 0.500. The markedly high recall (90.0%) relative to precision (57.8%) indicates the validation-tuned decision threshold favours predicting profit days, correctly capturing the large majority of actual profit days at the cost of a higher false-positive rate among loss days.

frequently misclassifies loss days as profit days, a trade-off that should be weighed against the intended application of the model (e.g., a bias toward the profit class may be tolerable, or even desirable, for a "buy-signal" screening tool, but would be costly for a strict loss-avoidance rule).

Table 3: Confusion matrix for the LSTM model on the 442-day test set

	Predicted Loss (0)	Predicted Profit (1)
Actual Loss (0)	28	164
Actual Profit (1)	25	225

Figure 2 visualises the same confusion matrix as a heat map, Figure 3 presents the receiver operating characteristic curve (AUC = 0.597), and Figure 4 plots the actual and LSTM-predicted profit/loss states for the

first 150 trading days of the test set, illustrating the model's pronounced tendency to predict profit states. Figure 5 plots the full NSE All-Share Index level series across the 2015-2023 sample period for context.




Figure 2. Confusion Matrix – LSTM Test Set




Figure 3. ROC Curve for LSTM Profit/Loss Classifier




Figure 4. Actual vs. Predicted Profit/Loss (first 150 test days)




Figure 5. NSE All-Share Index, 2015-2023 (Modelling Sample)

4.4 Recent Market Context (2024-2026)

Because the daily dataset available for model training and evaluation ends in December 2023, Figure 6 supplements the modelling sample with publicly reported year- and quarter-end NSE/NGX All-Share Index levels through July 2026, drawn from NGX Group corporate disclosures and financial media [25-28]. These figures, available only at monthly/quarterly rather than daily granularity and therefore excluded from model training, indicate that the index rose from 74,289.02 at the end of the modelling sample to 102,926.40 at the end of 2024 (+38.6%), 155,613.03 at

the end of 2025 (+51.2%), and to a range of approximately 201,700-247,400 by July 2026, reflecting a pronounced, sustained bull run driven substantially by naira devaluation and macroeconomic reform rather than by the technical patterns modelled in this study. This contextual update is presented for transparency and to situate the study's findings in time; it is not used for, and should not be interpreted as validating or invalidating, the LSTM classifier's daily profit/loss predictions, which were trained and evaluated exclusively on the 2015-2023 daily sample.




Figure 6. NSE All-Share Index, 2015-2023 (Daily, Modelling Sample) with Recent Market Context through July 2026 (Monthly/Quarterly, Approximate)

5. DISCUSSION

The empirical results provide encouraging evidence of LSTM performance for profit/loss classification on the NSE All-Share Index. The AUC of 0.597 confirms that the LSTM extracted statistically meaningful, non-random signal from the twelve engineered technical features, consistent with the broader literature reporting that deep sequence models can detect subtle non-linear structure in financial time series that eludes purely linear specifications [2,10]. The model's test accuracy of 57.2% exceeded both the naive persistence rule (56.6%) and the linear OLS baseline (49.1%), and its F1-score of 70.4% indicates a favourable balance between precision and recall once the very high recall (90.0%) is taken into account. This pattern of modest but genuine improvement over simple heuristics is consistent with several prior applied studies: Nelson, Pereira and de Oliveira [3] reported LSTM-based direction classifiers for Brazilian equities achieving accuracy only marginally above 55%, while Sezer, Gudelek and Ozbayoglu [23] note in their systematic review that reported deep learning gains in finance, though frequently modest in absolute terms, are nonetheless consistently observed across multiple markets and sample periods when evaluation avoids look-ahead bias—a standard explicitly followed here through strictly chronological partitioning.

A plausible explanation for the LSTM's outperformance over the naive persistence rule in this 2015-2023 sample is that the model's twelve-feature representation—incorporating moving averages, volatility and the Relative Strength Index alongside raw price levels—captures somewhat more information about the evolving state of the market than a rule that simply repeats the prior day's direction, particularly across a sample period spanning multiple distinct volatility regimes (the 2015-2016 recession, the 2020 pandemic shock, and the post-2020 recovery). At the same time, the pronounced asymmetry between the LSTM's high recall (90.0%) and modest specificity (14.6%, from Table 3) indicates that the validation-tuned decision threshold produces a classifier heavily biased toward predicting profit days; this trade-off is consistent with Nelson et al. [3] and Patadiya [29] and with the broader observation in Chen, Zhou and Dai [17] that emerging-market LSTM classifiers often achieve accuracy gains that are concentrated asymmetrically across the outcome classes rather than uniformly.

These findings should be interpreted with two important caveats highlighted by the recent-market-context analysis in Section 4.4. First, the daily modelling sample ends in December 2023, while the NSE/NGX All-Share Index has, according to publicly reported year-end figures, since risen by a further 38.6% in 2024 and 51.2% in 2025 [25,27], reaching approximately 201,700-247,400 points by July 2026 [28]; this sustained bull run reflects structural macroeconomic developments, particularly naira devaluation, that are largely orthogonal to the technical predictors used in this study and were not observed by the model during training. Second, as with all single-market, single-index studies,

the reported performance gains, while genuine within this sample, should not be assumed to generalise unconditionally to future regimes without periodic re-estimation. This does not diminish the practical relevance of the framework: the LSTM's demonstrated ability to outperform both a naive and a linear baseline on a nine-year, multi-regime sample supports its use as a component of a broader, regularly retrained trading or risk-management decision process, and the results reinforce the value of reporting non-deep baselines and recent market context alongside deep learning results, a practice recommended by Sezer et al. [23] but inconsistently followed in the applied finance literature.

6. CONCLUSION

This study developed and evaluated a reproducible LSTM-based deep learning framework for predicting daily profit and loss states of the Nigerian Stock Exchange All-Share Index using nine years (2015-2023) of historical daily trading data, supplemented with recent market context through July 2026. The proposed model achieved a test-set AUC of 0.597, an F1-score of 0.704 and an accuracy of 57.2%, exceeding both a naive persistence heuristic (56.6%) and a linear OLS baseline (49.1%), demonstrating that carefully engineered technical features, processed through a properly regularised and chronologically validated LSTM architecture, can extract statistically meaningful and practically useful signal from noisy emerging-market price data even across a sample spanning multiple distinct market regimes.

For practitioners, the results suggest that LSTM-based classifiers can offer a modest but genuine edge over simple heuristics for NSE profit/loss screening, provided model evaluation includes non-deep benchmarks and periodic re-estimation to account for regime change a point underscored by the substantial market appreciation observed after the end of the modelling sample. For researchers, the study contributes an openly documented, fully reproducible pipeline including data cleaning, feature engineering, sequence construction and a from-first-principles LSTM implementation that can be extended to other African exchanges, updated with more recent daily data as it becomes available, or augmented with additional data modalities. Future work should explore hybrid CNN-LSTM or attention-based architectures, incorporation of macroeconomic and sentiment covariates (particularly variables capturing exchange-rate policy given its apparent influence on the 2024-2026 rally), multi-horizon forecasting, and formal statistical significance testing (e.g., the Diebold-Mariano test) to more rigorously establish whether observed performance differences between the LSTM and baseline models are statistically distinguishable from noise. Extending the daily sample to include 2024-2026 data as it becomes available in daily form, and evaluating economic significance through backtested trading strategies net of transaction costs, would further strengthen the practical relevance of the framework for Nigerian capital market participants.

REFERENCES

- [1]. Hochreiter, S., & Schmidhuber, J., Long short-term memory, *Neural Computation*, 1997, 9(8) 1735-1780.
- [2]. Fischer, T., & Krauss, C., Deep learning with long short-term memory networks for financial market predictions, *European Journal of Operational Research*, 2018, 270(2) 654-669.
- [3]. Nelson, D. M. Q., Pereira, A. C. M., & de Oliveira, R. A., Stock market's price movement prediction with LSTM neural networks, *Proceedings of the 2017 International Joint Conference on Neural Networks (IJCNN)*, 2017, 1419-1426.
- [4]. Salisu, A., Application of machine learning and deep learning in stock price prediction: A comparative study using Dangote Cement Plc dataset from the Nigerian Stock Exchange, *Global Journal of Artificial Intelligence and Technology Development*, 2026, 4(1) 119-128.
- [5]. Adebisi, A. A., Adewumi, A. O., & Ayo, C. K., Comparison of ARIMA and artificial neural networks models for stock price prediction, *Journal of Applied Mathematics*, 2014, 2014 Article 614342.
- [6]. Kim, K., Financial time series forecasting using support vector machines, *Neurocomputing*, 2003, 55(1-2) 307-319.
- [7]. Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y., Learning phrase representations using RNN encoder-decoder for statistical machine translation, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, 1724-1734.
- [8]. Sherstinsky, A., Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network, *Physica D: Nonlinear Phenomena*, 2020, 404 132306.
- [9]. Gers, F. A., Schmidhuber, J., & Cummins, F., Learning to forget: Continual prediction with LSTM, *Neural Computation*, 2000, 12(10) 2451-2471.
- [10]. Bao, W., Yue, J., & Rao, Y., A deep learning framework for financial time series using stacked autoencoders and long-short term memory, *PLoS ONE*, 2017, 12(7) e0180944.
- [11]. Selvin, S., Vinayakumar, R., Gopalakrishnan, E. A., Menon, V. K., & Soman, K. P., Stock price prediction using LSTM, RNN and CNN-sliding window model, *Proceedings of the 2017 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, 2017, 1643-1647.
- [12]. Roondiwala, M., Patel, H., & Varma, S., Predicting stock prices using LSTM, *International Journal of Science and Research*, 2017, 6(4) 1754-1756.
- [13]. Siarni-Namini, S., Tavakoli, N., & Siarni Namin, A., A comparison of ARIMA and LSTM in forecasting time series, *Proceedings of the 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2018, 1394-1401.
- [14]. Zhang, L., Aggarwal, C., & Qi, G.-J., Stock price prediction via discovering multi-frequency trading patterns, *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017, 2141-2149.
- [15]. Ding, X., Zhang, Y., Liu, T., & Duan, J., Deep learning for event-driven stock prediction, *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015, 2327-2333.
- [16]. Akita, R., Yoshihara, A., Matsubara, T., & Uehara, K., Deep learning for stock prediction using numerical and textual information, *Proceedings of the 15th IEEE/ACIS International Conference on Computer and Information Science (ICIS)*, 2016, 1-6.
- [17]. Chen, K., Zhou, Y., & Dai, F., A LSTM-based method for stock returns prediction: A case study of China stock market, *Proceedings of the 2015 IEEE International Conference on Big Data*, 2015, 2823-2824.
- [18]. Patel, J., Shah, S., Thakkar, P., & Kotecha, K., Predicting stock market index using fusion of machine learning techniques, *Expert Systems with Applications*, 2015, 42(4) 2162-2172.
- [19]. Althelaya, K. A., El-Alfy, E.-S. M., & Mohammed, S., Evaluation of bidirectional LSTM for short-and-long-term stock market prediction, *Proceedings of the 9th International Conference on Information and Communication Systems (ICICS)*, 2018, 151-156.
- [20]. Bhandari, H. N., Rimal, B., Pokhrel, N. R., Rimal, R., Dahal, K. R., & Khatri, R. K. C., Predicting stock market index using LSTM, *Machine Learning with Applications*, 2022, 9 100320.
- [21]. Nabipour, M., Nayyeri, P., Jabani, H., Mosavi, A., Salwana, E., & Shahab, S., Deep learning for stock market prediction, *Entropy*, 2020, 22(8) 840.
- [22]. Rather, A. M., Agarwal, A., & Sastry, V. N., Recurrent neural network and a hybrid model for prediction of stock returns, *Expert Systems with Applications*, 2015, 42(6) 3234-3241.
- [23]. Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M., Financial time series forecasting with deep learning: A systematic literature review: 2005-2019, *Applied Soft Computing*, 2020, 90 106181.
- [24]. Ahmed, A. A., Abdullahi, A. U., Gital, A. Y., & Dutse, A. Y., Application of Artificial Intelligence in Supply Chain Management: A Review on Strengths and Weaknesses of Predictive Modeling Techniques, *Scientific Journal of Engineering and Technology*, 2024, 1(2) 1-18.
- [25]. NGX Group, NGX ASI closes 2024 with +37% return, *Nigerian Exchange Group Corporate Disclosures*, 2025.
- [26]. ThisDayLive, NGX ASI closes 2024 with impressive annual growth of 37.65%, *THISDAYLIVE*, 2025.
- [27]. Cowrywise, The NGX returned 51% in 2025: Here's what young Nigerians need to know, *Cowrywise Blog*, 2026.
- [28]. NGX Pulse, Nigeria's stock market intelligence platform: Daily market summaries, *NGX Pulse*, 2026.
- [29]. Patadiya, N., Artificial Intelligence for Cyber Threat Detection in Cloud Computing : A Hybrid Random Forest and LSTM Approach, *International Journal of Environmental Sciences*, 2025, 11(7) 63--72.