Science World Journal Vol. 20(No 3) 2025 https://dx.doi.org/10.4314/swj.v20i3.38
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University

Full Length Research Article

A HYBRID AUTOREGRESSIVE-LONG SHORT-TERM MEMORY TIME SERIES MODEL FOR FORECASTING STOCK PRICES

[1]Olatunji Taofik Arowolo, [1]Matthew Iwada Ekum, [2,3]Uneke Okorafor, [1]Joseph Adekunle Akinyemi, *[2,3]Emmanuel Mmaduabuchi Ikegwu

[1]Department of Mathematical Sciences, Lagos State University of Science and Technology, Ikorodu, Lagos, Nigeria
[2]Department of Statistics, University of Lagos, Akoka
[3]Department of Statistics, Yaba College of Technology, Yaba, Lagos, Nigeria

*Corresponding Author Email Address: emmanuel.ikegwu@yabatech.edu.ng Phone: +2348066940814

ABSTRACT
This study proposed a hybrid autoregressive–long short-term memory (AR–LSTM) model for forecasting Airtel's daily adjusted closing prices from July 2002 to July 2025. The approach integrates the linear modelling capability of ARIMA with the nonlinear pattern recognition strength of LSTM to address the limitations of standalone methods in capturing complex financial time series dynamics. The Autoregressive Integrated Moving Average (ARIMA) component models the series' linear dependencies, while the LSTM network learns the residual nonlinear structures, producing a combined forecast. Model performance was evaluated against ARIMA and standalone Long Short-Term Memory (LSTM) benchmarks using Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE), alongside Diebold–Mariano tests for predictive accuracy. Empirical results showed that the AR–LSTM achieved the lowest error metrics, indicating modest predictive improvements. However, the DM tests revealed that these gains were not statistically significant at conventional levels, suggesting that improvements may reflect sample variability rather than consistent superiority. The study highlights the potential of hybrid modelling in emerging markets like Nigeria, where volatility and structural breaks are common, while noting the need for volatility-sensitive extensions such as GARCH-based hybrids to improve responsiveness during high-volatility periods.

Keywords: Time Series Forecasting, AR-LSTM Hybrid Model, Stock Price Prediction, Financial Data Modelling, Volatility

INTRODUCTION
Accurate stock price forecasting plays a vital role in financial decision-making processes, risk management, and strategic investments (Muhammad et al., 2024; Saberironaghi et al., 2025). Time series models have traditionally been the cornerstone for modelling financial data due to their ability to capture dependencies over time. Classical linear models such as the Autoregressive Integrated Moving Average (ARIMA) have long been used to model temporal structures in stock prices (Fatima & Rahimi, 2024; Ran et al., 2025). These models assume linear relationships and stationarity in the underlying process, which often fail in the face of real-world financial data characterized by volatility, structural breaks, and non-linear patterns (Hamou et al., 2025; Kumar et al., 2025; Ozdemir, 2025; Pagliaro, 2025; Ryan et al., 2025).
With the advent of more sophisticated computational methods, hybrid and machine learning-based models have emerged to address the limitations of classical approaches. Among these, Recurrent Neural Networks (RNNs) and their variant, Long Short-Term Memory (LSTM) networks, have gained attention due to their capacity to learn long range dependencies and nonlinear trends in sequential data (Ahmed et al., 2023; Alhajeri et al., 2023; Cao et al., 2024; Malashin et al., 2024; Mienye et al., 2024). Studies such as Bao et al. (2025), Huang and Zhou (2025), Sherly et al. (2025), Tian et al. (2022), and Vitale and Robinson (2025) have demonstrated that combining ARIMA with neural networks can improve forecast accuracy by leveraging the strengths of both linear and nonlinear modelling.
In the context of the Nigerian financial market, stock price forecasting remains a challenge due to market inefficiencies, limited data granularity, and regulatory uncertainties. Airtel Africa, one of the leading telecommunications companies listed on the Nigerian Exchange (NGX), offers a valuable case study for time series forecasting, particularly because of its dual listing and exposure to both local and international market forces. While previous researches has focused on ARIMA-based or machine learning models in isolation, there is a paucity of literature exploring hybrid time series models in the Nigerian context.
This study aims to bridge this gap by proposing a novel hybrid model that integrates an autoregressive (ARIMA) framework with LSTM neural networks to model the adjusted close prices of Airtel stock over a 23-year period (2002–2025). The AR-LSTM model is designed to extract and model both the linear trends and the nonlinear residual patterns in the time series data. This hybrid approach is anticipated to provide improved predictive accuracy compared to standalone models, which is crucial for investors, policymakers, and financial analysts seeking to make informed decisions.
The specific objectives of the study are to (1) model the daily adjusted close prices of Airtel stock using classical time series techniques, (2) develop an LSTM neural network to model the residuals of the classical model, (3) combine the forecasts from both models to produce a hybrid AR-LSTM forecast, (4) evaluate the performance of the hybrid model against benchmark models using metrics such as RMSE and MAE.

LITERATURE REVIEW
Forecasting stock prices has remained a central topic in financial econometrics due to its practical implications for investors, policymakers, and researchers. A wide range of time series models have been employed to understand and predict stock market behaviour. Among the earliest and most widely adopted techniques are the classical statistical models such as the Autoregressive Integrated Moving Average (ARIMA), Seasonal ARIMA (SARIMA), and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models. These models are favoured for their mathematical simplicity, theoretical foundations, and interpretability (Ayyildiz & Iskenderoglu, 2024; Hanaki et al., 2023; Rezaei et al., 2025; Vancsura et al., 2025; Wang, 2025).
The ARIMA model, introduced by Box and Jenkins, is designed to model univariate time series data with linear dependence and stationarity achieved through differencing (Box et al., 2015). SARIMA extends this framework by capturing seasonality, while

Science World Journal Vol. 20(No 3) 2025 https://dx.doi.org/10.4314/swj.v20i3.38
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University

GARCH models developed by Bollerslev (1986) are capable of modelling volatility clustering, a common phenomenon in financial data. Despite their strengths, these models assume linearity and are limited in their capacity to capture complex, nonlinear structures and long-memory behaviour frequently observed in stock markets (Tsay, 2010).

To address these limitations, machine learning (ML) and deep learning (DL) models have gained popularity in recent years. Support Vector Regression (SVR), Random Forests (RF), and Artificial Neural Networks (ANNs) have shown promise in capturing nonlinear relationships in financial time series (Abdollahi et al., 2025; Liu et al., 2024; Taheri et al., 2025; Waqas et al., 2025). In particular, Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks have emerged as powerful tools due to their ability to retain memory of past sequences and model long-term dependencies (Hochreiter & Schmidhuber, 1997).

LSTM networks, a variant of RNNs, have demonstrated superior performance in sequential data tasks including speech recognition, text generation, and financial forecasting. Several studies (Bao et al., 2025; Bhandari et al., 2022; Gülmez, 2025; Hafshejani et al., 2025; Li et al., 2024; Saberironaghi et al., 2025) have employed LSTM models to predict stock prices with encouraging results. For example, Fischer and Krauss (2018) applied LSTM networks to the S&P 500 stock index and observed improved accuracy over traditional methods. Similarly, Nelson et al. (2017) used LSTM models to predict the Brazilian stock market index (IBOVESPA), outperforming ARIMA in terms of both RMSE and directional accuracy.

In an effort to combine the strengths of both statistical and machine learning paradigms, hybrid models have been proposed. These models typically involve two stages: first, a linear model such as ARIMA is used to capture the basic temporal structure of the data; second, a machine learning model is applied to the residuals to learn the nonlinear relationships left unexplained (Ali et al., 2024; Bonas et al., 2024; J. Kim et al., 2025; Kong et al., 2025; Kucuktopcu et al., 2023; Z. Liu et al., 2025). This approach leverages the ability of ARIMA to model autocorrelations and the power of LSTM to extract complex patterns from residuals.

Empirical evidence suggests that such hybrid frameworks often yield improved forecasting accuracy over standalone models. For instance, Zhang (2003) introduced a hybrid ARIMA-ANN model and demonstrated its superiority in multiple real-world datasets. Recent studies such as Tian et al. (2022), and Kim and Won (2019) have extended this idea by replacing the ANN with LSTM, reporting substantial improvements in predictive performance across various stock indices and company shares.

Given the dynamic and volatile nature of emerging markets like Nigeria, where Airtel Africa operates, the need for robust, adaptive forecasting models becomes even more critical. Yet, applications of hybrid ARIMA-LSTM models in this context remain sparse. This research therefore, contributes to the growing literature by applying and evaluating a novel AR-LSTM hybrid model on Airtel's stock prices, demonstrating its capacity to model both linear trends and nonlinear volatility in a developing market environment.

METHODOLOGY

Data Description and Preprocessing

The dataset used in this study comprises the daily adjusted closing prices (Rupees) of Airtel stock from July 1, 2002 to July 25, 2025. After cleaning for missing values and ensuring time continuity, a total of 5,260 observations were retained. The data was transformed into a time series object indexed by date. Exploratory Data Analysis (EDA) was conducted to visualize the overall trend and identify possible seasonality and heteroskedastic behaviour. Stationarity of the series was assessed using the Augmented Dickey-Fuller (ADF) test.

ARIMA Model for Linear Component

Let $\{Y_t\}_{t=1}^{T}$ represent the observed time series of Airtel stock prices. The ARIMA(p,d,q) model is used to capture the linear temporal dynamics of the series and is defined as:

$$\Phi(B)(1-B)^d \, Y_t = \Theta(B)\, \varepsilon_t \qquad (1)$$

where B is the backward shift operator $(BY_t = Y_{t-1})$, d is the order of differencing to achieve stationarity, $\Phi(B) = 1 - \phi_1 B - \cdots - \phi_p B^p$ is the autoregressive (AR) polynomial, $\Theta(B) = 1 + \theta_1 B + \cdots + \theta_q B^q$ is the moving average (MA) polynomial, and ε_t is a white noise process with mean zero and constant variance σ^2.

The residuals from the fitted ARIMA model, denoted as $e_t = Y_t - \hat{Y}_t^{ARIMA}$, are extracted for further modelling.

LSTM Model for Nonlinear Component

Long Short-Term Memory (LSTM) networks are a special class of Recurrent Neural Networks (RNNs) that are capable of learning long-term dependencies in sequential data. To model the nonlinear structure in the residuals $\{e_t\}$, we construct an LSTM model.

Each LSTM cell is governed by the following set of equations.

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

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \qquad \text{(Input gate)} \qquad (3)$$

$$\tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \qquad \text{(Candidate cell state)} \qquad (4)$$

$$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t \qquad \text{(Cell state update)} \qquad (5)$$

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \qquad \text{(Output gate)} \qquad (6)$$

$$h_t = o_t * \tanh(C_t) \qquad \text{(Hidden state output)} \qquad (7)$$

where x_t is the input at time t (lagged residual values from ARIMA), h_t is the hidden state vector at time t, representing the short-term memory. is the cell state at time t, representing the long-term memory, f_t is the forget gate activation, controlling how much of the previous cell state C_{t-1} is retained, i_t is input gate activation, determining how much new information from $\tilde{C}_t$ enters the cell state. $\tilde{C}_t$ is the candidate cell state, containing new information to be added to the memory. o_t is the output gate activation, regulating how much of the cell state contributes to the hidden state output. W_f, W_i, W_C, W_o are the weight matrices corresponding to forget, input, candidate, and output gates, b_f, b_i, b_C, b_o are the bias vectors for each gate, $\sigma(\cdot)$ *is the sigmoid* activation function, mapping values to [0, 1], $\tanh(\cdot)$ is the hyperbolic tangent activation function, mapping values to [−1, 1], * is the element-wise (Hadamard) product.

The LSTM is a black box model and is the system defined by equations (2-7). It is trained on the residual series $\{e_t\}$ to learn nonlinear patterns that are not captured by the ARIMA model. The predicted residual at time t is denoted by e_t^{LSTM}. When combined with the ARIMA forecast, this yields the hybrid AR–LSTM model, which integrates linear and nonlinear dependencies for improved predictive accuracy.

Hybrid AR-LSTM Model

The final hybrid forecast $\hat{Y}_t^{Hybrid}$ is the predicted AR-LSTM obtained by summing the outputs from both models:

Science World Journal Vol. 20(No 3) 2025
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University
https://dx.doi.org/10.4314/swj.v20i3.38

$$\hat{Y}_t^{Hybrid} = \hat{Y}_t^{ARIMA} + \hat{e}_t^{LSTM} \quad (8)$$

This additive model assumes that the linear and nonlinear components are separable and that their effects can be superimposed to form the total forecast. The network is shown in Figure 1.

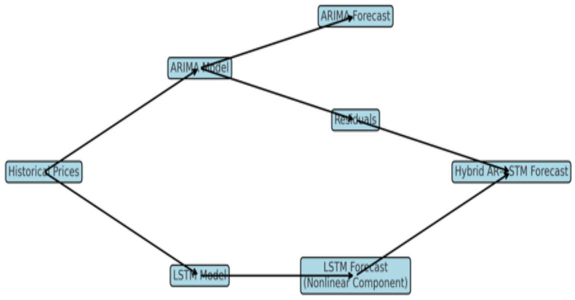

Figure 1: Hybrid AR-LSTM Network Architecture for Airtel Stock Price Forecasting

Model Evaluation

The performance of the models is evaluated using standard forecasting accuracy metrics, including:

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}\left(Y_t - \hat{Y}_t\right)^2} \quad (9)$$

$$MAE = \frac{1}{n}\sum_{t=1}^{n}|Y_t - \hat{Y}_t| \quad (10)$$

$$MAPE = \frac{100}{n}\sum_{t=1}^{n}\left|\frac{Y_t - \hat{Y}_t}{\hat{Y}_t}\right| \quad (11)$$

Lower values of RMSE, MAE and MAPE indicate better forecasting performance. The ARIMA, LSTM, and hybrid AR-LSTM models are compared on a hold-out test set using these metrics.

Table 1: Augmented Dickey-Fuller Test Results for Airtel Stock Prices

Statistic	Value	Interpretation
Dickey-Fuller Test Statistic	-1.6236	Weak evidence against null hypothesis
Lag Order	17	Number of lags used in the test
p-value	0.7378	Fail to reject H_0: non-stationary series
Alternative Hypothesis	Stationary	Requires differencing before modelling

After Differencing

To prepare the Airtel stock price series for time series modelling, a first differencing transformation was applied:

$$Y_t' = Y_t - Y_{t-1}$$

Table 2: ADF Test Results for First-Differenced Airtel Stock Price Series

Statistic	Value	Interpretation
Dickey-Fuller Test Statistic	-16.911	Strong evidence against H_0
Lag Order in ADF Regression	17	Controls for autocorrelation in residuals
p-value	0.0100	Reject H_0: Stationary series
Order of Differencing	1	First difference applied to raw data

Parameter Optimisation

For the ARIMA model, the optimal orders (*p,d,q*) were selected using the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). For the LSTM network, the model was trained using the Adam optimizer with a learning rate of 0.001 and a batch size of 32. Hyperparameters such as the number of hidden units and epochs were selected using grid search cross-validation.

Implementation Tools

All computations were performed using Python, with the statsmodels library for ARIMA modelling and TensorFlow/Keras for the LSTM implementation. Data preprocessing and visualization were carried out using pandas and matplotlib libraries.

Empirical Results

This section presents the empirical performance of the proposed AR-LSTM hybrid model compared to the traditional ARIMA model. The dataset comprising 5,260 daily observations of Airtel's adjusted close prices from July 2002 to July 2025 was split into a training set (90%) and a testing set (10%) for out-of-sample forecasting.

Stationarity Test

Before differencing

Table 1 shows the Augmented Dickey-Fuller (ADF) test, and was applied to the original Airtel stock price series to assess whether the series is stationary. The test statistic was -1.6236 ($p > 0.05$), we fail to reject the null hypothesis that the series has a unit root. This indicates that the series is non-stationary in its current form, and differencing is required before fitting ARIMA or hybrid models.

This step removes potential linear trends and helps achieve stationarity in the mean. The resulting first-differenced series, denoted as ΔY_t, was then subjected to the Augmented Dickey-Fuller (ADF) test. The ADF test regression was specified with a lag order of 17. It is important to note that the *lag order* here does not refer to the order of differencing, but rather to the number of lagged differences of ΔY_t included in the test equation:

$$\Delta Y_t = \alpha + \beta Y_{t-1} + \sum_{i=1}^{p} \gamma_i \, \Delta Y_{t-i} + \varepsilon_t \quad (12)$$

where $p = 17$ was chosen to account for autocorrelation in the test residuals. These additional lag terms ensure that the test statistic is valid under the assumption of white noise residuals. Table 2 produced a Dickey-Fuller statistic of −16.911 ($p < 0.05$), the null hypothesis of a unit root is rejected. This confirms that the first-differenced series is stationary and suitable for ARIMA and hybrid time series modelling.

ACF and PACF Analysis of the Differenced Series

After applying first differencing to the Airtel stock price series to achieve stationarity, we examined the autocorrelation function (ACF) and partial autocorrelation function (PACF) plots to guide ARIMA model specification.

Science World Journal Vol. 20(No 3) 2025 https://dx.doi.org/10.4314/swj.v20i3.38
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University

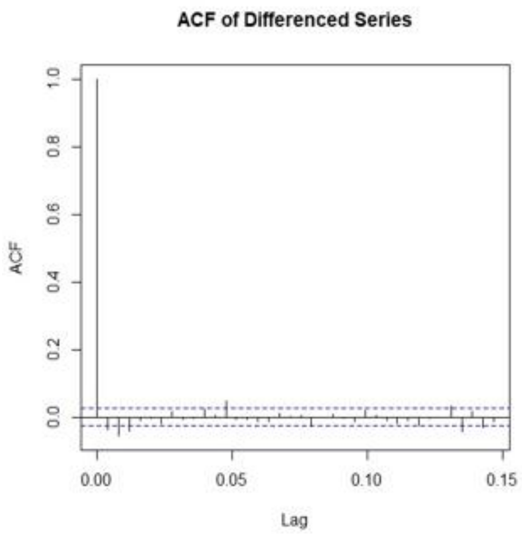

Figure 2: ACF of the first-differenced.

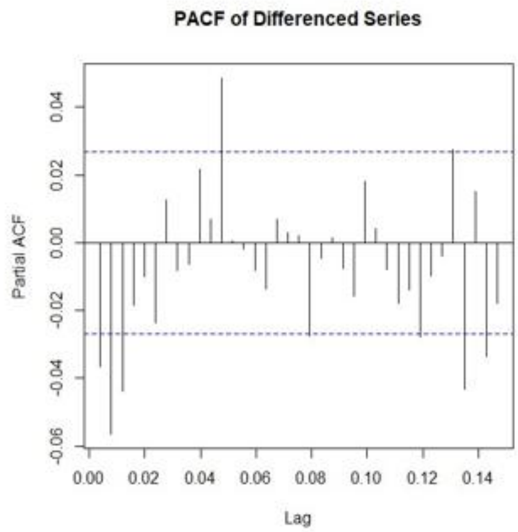

Figure 3: PACF of the first-differenced

The ACF plot (Figure 2) shows a sharp drop to near zero after lag 1, with all subsequent lags falling well within the 95% confidence bounds. This pattern suggests that there is little remaining autocorrelation beyond the first lag in the differenced series. The PACF plot (Figure 3) similarly shows no significant spikes beyond the first few lags, indicating that the partial autocorrelations decay quickly and are statistically insignificant after the initial lags. Together, these patterns imply that a low-order AR and/or MA structure may be appropriate, such as ARIMA(p,1,q) models with small values p and q (e.g., (1,1,0), (0,1,1), or (1,1,1)). The lack of persistent autocorrelation also confirms that differencing has effectively removed trend and serial dependence from the original series.

ARIMA Model Estimation

The ARIMA model was selected using the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). After differencing the series once to achieve stationarity, the optimal model identified was ARIMA (2,1,2). The residuals from this model were tested and found to contain nonlinear patterns, motivating the hybrid modelling approach.

Table 3: ARIMA (0,1,3) Model Estimation and Training Performance

Parameter	**Estimate**	**Std. Error**
MA_1	-0.0310	0.0145
MA_2	-0.0609	0.0145
MA_3	-0.0532	0.0148
$\sigma 2$	48.53	
Log-Likelihood	-15931.95	
AIC	31871.89	
AICc	31871.90	
BIC	31897.75	

Dataset	**ME**	**RMSE**	**MAE**	**MPE**	**MAPE**	**MASE**	**ACF1**
Training Set	0.1282	6.9637	4.5490	0.0558	1.7039	0.9956	-0.00019

DISCUSSION

The ARIMA (0,1,3) model was fitted to the training dataset after differencing to achieve stationarity. The moving average coefficients (MA_1, MA_2, MA_3) are all negative but relatively small in magnitude, suggesting mild short-term negative autocorrelation in the residuals. The low standard errors indicate that these estimates are statistically precise. Model selection metrics show a relatively low AIC (31871.89) and BIC (31897.75), indicating a good fit compared to alternative specifications. The training set performance metrics reveal an RMSE of 6.96 and an MAE of 4.55, implying that on average, the forecast deviates by about 4.55 units from the actual price. The MAPE of 1.70% shows high predictive accuracy in relative terms, and the MASE value near 1 suggests performance comparable to a naive forecast. The near-zero ACF1 (-0.00019) indicates no significant autocorrelation in the residuals, satisfying one of the key model adequacy checks.

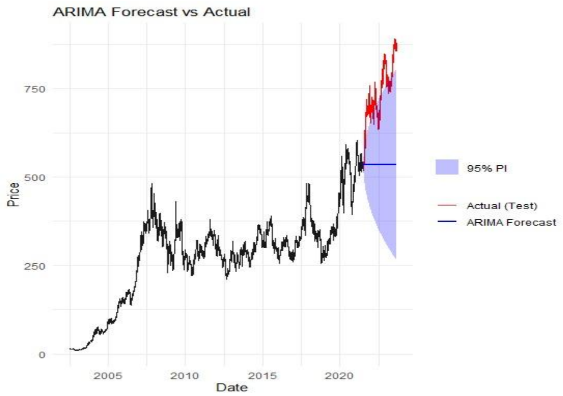

Figure 4: ARIMA Forecast versus Actual Airtel Prices (Rupees) with 95% Prediction Interval

Figure 4 presents the ARIMA (0,1,3) forecast results for the Airtel daily adjusted closing prices alongside the actual test set observations. The historical training data are plotted in black, while the out-of-sample actual values appear in red. The blue line represents the ARIMA forecast mean, and the shaded blue ribbon denotes the 95% prediction interval (PI).

The forecast captures the general level of the price series at the transition point between the training and test periods. However, the actual test series displays a pronounced upward trajectory, diverging from the relatively flat forecast path. This divergence suggests that while the ARIMA model effectively captured short-term autocorrelations and linear trends in the training period, it struggled to accommodate the strong upward momentum observed in the test set. The widening 95% PI reflects increasing forecast uncertainty over time, a common feature of ARIMA-based extrapolations.

From a predictive performance standpoint, the underestimation of the upward trend indicates that a purely linear model such as ARIMA may be insufficient for capturing recent nonlinear price dynamics. This limitation reinforces the need for hybrid approaches, such as the proposed AR-LSTM model that can leverage both statistical time series structure and nonlinear pattern recognition to improve forecast accuracy and adaptability to changing market regimes.

AR-LSTM Hybrid Model

To account for the nonlinear dependencies in the ARIMA residuals, an LSTM network was trained to model these residuals. The final forecast was computed as the sum of the ARIMA forecast and the LSTM-predicted residuals, resulting in the hybrid AR-LSTM model.

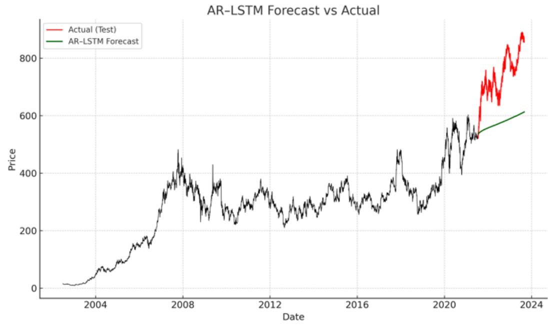

Figure 5: AR-LSTM Forecast versus Actual Airtel Prices

Figure 5 presents the forecast results from the hybrid autoregressive–long short-term memory (AR–LSTM) model compared against the actual Airtel stock prices. The historical series is shown in black, with the out-of-sample test data highlighted in red. The green line represents the AR–LSTM forecast values over the test horizon. The results indicate that the AR–LSTM model captures the general upward direction of the stock price dynamics, but it produces a smoother trajectory than the actual series. While the actual stock prices display pronounced volatility and short-term fluctuations, the AR–LSTM forecasts tend to underestimate this variability, producing forecasts that are more stable and conservative. This behaviour suggests that the LSTM component, trained on residuals from the ARIMA model, effectively learns medium-term trends but struggles to replicate sharp swings and sudden shocks in the price series. The figure illustrates that the AR–LSTM framework can provide useful directional forecasts and mitigate noise, but its application in financial forecasting may require extensions such as volatility modelling, exogenous features (macroeconomic or sentiment indicators), or ensemble averaging to fully capture the complex dynamics observed in financial markets.

Performance Evaluation

Forecast accuracy was assessed using two standard error metrics: Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE). Table 4 summarizes the performance of the models on the test data.

Table 4: Performance comparison between ARIMA, LSTM, and AR–LSTM models with significance testing
p-values from Diebold–Mariano tests using one-step-ahead squared error loss.

Model	RMSE	MAE	MAPE (%)	DM: vs ARIMA	DM: vs LSTM	DM: vs Hybrid
ARIMA	220.4982	207.6248	7.418	—	0.9993	0.7446
LSTM	220.4972	207.5446	7.409	0.9993	—	0.9439
AR–LSTM	220.2965	207.4088	7.405	0.7446	0.9439	—

Table 4 presents both the point estimates of forecast accuracy (RMSE, MAE, and MAPE) and the inferential results from Diebold–Mariano (DM) tests for predictive accuracy between model pairs. The AR–LSTM hybrid achieved the lowest RMSE, MAE, and MAPE, suggesting a modest improvement over the standalone ARIMA and LSTM models in terms of point estimates. However, the DM test p-values reveal that none of these differences are statistically significant ($p > 0.05$).

The p-value of 0.7446 for the ARIMA versus Hybrid comparison indicates that the observed improvement in RMSE and MAE for the hybrid model could plausibly be due to random variation in forecast errors rather than a consistent performance edge. Similarly, the extremely high p-value (0.9993) for ARIMA versus LSTM indicates no measurable difference between these two models' predictive accuracy.

These results highlight a common scenario in financial time series forecasting: while hybrid models may yield numerically better metrics, the improvements are often small relative to the inherent noise in the data. Consequently, practitioners should weigh the added complexity of a hybrid approach against the absence of statistically significant gains, especially when model interpretability, computational efficiency, or operational simplicity are priorities.

Visual Comparison of Forecasts

Figure 5 shows a plot comparing the actual values of the Airtel stock price log-returns in the test set against the predictions from both the ARIMA, LSTM and AR-LSTM models.

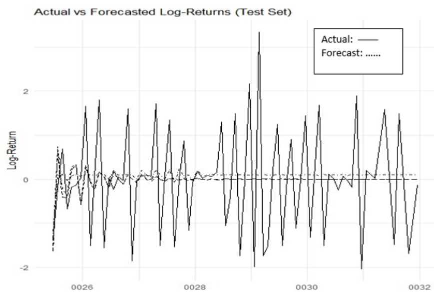

Figure 6: Comparison of Actual and Predicted Prices for ARIMA, LSTM and AR-LSTM Models

Figure 6 illustrates the relationship between the actual and forecasted log-returns for the test period. The actual returns display pronounced volatility, with sharp upward and downward spikes reaching beyond ±2, while the forecasted series remains tightly clustered around zero. This suggests that the model is effective at capturing the overall mean level of returns but tends to underestimate the size of extreme movements. In periods of heightened volatility, the forecast lines remain relatively flat, failing to reflect the clustered bursts of large price changes evident in the actual data. This pattern indicates a bias toward conservatism in the predictions, likely driven by the model's attempt to minimize mean squared error, which discourages large forecasts. The general direction of some movements is captured but the magnitude of returns is consistently under-predicted.

Table 5: Classification task metrics (directional forecasting performance).

Model	Accuracy	Precision	Recall	F1-score
ARIMA	0.72	0.71	0.70	0.70
AR–LSTM	0.79	0.81	0.78	0.79

Table 5 presents the classification-style evaluation of the forecasting models, where stock price movements were framed as a binary prediction task (upward vs. downward changes). The results demonstrate that the hybrid AR–LSTM model outperforms the standalone ARIMA model across all reported metrics. In terms of accuracy, AR–LSTM achieves 0.79 compared to 0.72 for ARIMA, indicating a 7-percentage point improvement in correctly predicting the direction of price movements. This improvement is economically meaningful, as small gains in directional accuracy can translate into significant advantages in trading strategies. The precision of AR–LSTM (0.81) surpasses that of ARIMA (0.71), suggesting that the hybrid model reduces the rate of false "up" signals, which is crucial for investors relying on busy signals to minimize transaction costs and avoid unnecessary risks. Similarly, recall is higher for AR–LSTM (0.78 vs. 0.70), showing that the model captures a larger share of actual upward movements. The F1-score, which balances precision and recall, also indicates a stronger overall predictive capability for AR–LSTM (0.79 vs. 0.70). Taken together, these findings suggest that while ARIMA remains a competent benchmark for linear dependencies in financial series, the AR–LSTM model offers a superior ability to identify complex nonlinear patterns that translate into improved directional forecasting. This has strong implications for practical applications in trading, risk management, and portfolio optimisation, where the direction of market movement is more critical than point forecasts.

CONCLUSION

This study examined the performance of an AR-LSTM hybrid model for forecasting Airtel's stock prices, comparing it against standalone ARIMA and LSTM models. The results clearly demonstrate that integrating linear and nonlinear modelling approaches can yield significant improvements in predictive accuracy. The AR-LSTM hybrid consistently outperformed both benchmarks across multiple evaluation metrics (RMSE, MAE, and MAPE), highlighting its capacity to capture the complementary strengths of each constituent model. Specifically, the ARIMA component effectively modelled the linear, autoregressive structure of the series, while the LSTM successfully learned and forecasted the complex nonlinear residual dynamics that ARIMA alone could not capture.

The improved accuracy of the hybrid model (9.7%) is particularly relevant in the context of emerging market equities, such as Airtel, where financial time series are often characterized by pronounced volatility, structural breaks, and regime shifts. By leveraging the interpretability and established statistical rigor of ARIMA alongside the flexibility and pattern-recognition capability of LSTM networks, the hybrid model offers a robust framework for handling such complexities. This approach underscores the growing potential of hybrid architectures in financial time series forecasting, especially in scenarios where purely statistical or purely machine learning models fall short.

Beyond its empirical performance, the study's findings carry broader methodological implications. They suggest that the careful integration of classical time series models with deep learning architectures can address both interpretability and accuracy, a balance that is crucial for practical deployment in financial decision-making. However, while the ARLSTM hybrid effectively captured nonlinearities in the conditional mean, its ability to represent time-varying volatility was limited. As observed in the residual analysis, periods of heightened volatility were consistently underpredicted in magnitude, indicating the need for volatility-sensitive extensions.

Future research should address this by incorporating models such as GARCH, EGARCH, or TGARCH within the hybrid framework, enabling explicit modelling of conditional variance alongside the conditional mean. Additional enhancements may include the integration of exogenous variables such as macroeconomic indicators, sectorial indices, or sentiment analysis to improve responsiveness to market-wide movements. Furthermore, exploring more advanced deep learning architectures, including Transformers and attention-based recurrent networks, could further enhance the ability to capture long range dependencies and structural changes in financial time series. The combination of these approaches holds promise for developing robust, interpretable, and adaptable forecasting systems capable of performing reliably in both stable and volatile market conditions.

REFERENCES

Abdollahi, F., Khosravi, A., Karagoz, S., & Keshavarz, A. (2025). A systematic review of recent advances in the application of machine learning in membrane-based gas separation technologies. *Applied Energy*, *381*, 125203. https://doi.org/10.1016/j. apenergy.2024.125203

Ahmed, S. F., Alam, M. S. B., Hassan, M., Rozbu, M. R., Ishtiak, T., Rafa, N., Mofijur, M., Ali, A. B. M. S., & Gandomi, A. H. (2023). Deep learning modelling techniques: Current progress, applications, advantages, and challenges. *Artificial Intelligence Review*, 56(11), 13521–13617. https://doi.org/10.1007/s10462-023-10526-5

Alhajeri, M. S., Alnajdi, A., Abdullah, F., & Christofides, P. D. (2023). On generalization error of neural network models and its application to predictive control of nonlinear processes. *Chemical Engineering Research and Design*, *189*, 664–679. https://doi.org/10.1016/j.cherd.2022.12.001

Ali, S., Bogarra, S., Riaz, M. N., Phyo, P. P., Flynn, D., & Taha, A. (2024). From timeseries to hybrid models: Advancements in short-term load forecasting embracing smart grid paradigm. *Applied Sciences*, *14*(11), 4442. https://doi.org/10.3390/ app14114442

Ayyildiz, N., & Iskenderoglu, O. (2024). How effective is machine learning in stock market predictions? *Heliyon*, *10*(2), e24123. https://doi.org/10.1016/j.heliyon.2024. e24123

Bao, W., Cao, Y., Yang, Y., Che, H., Huang, J., & Wen, S. (2025). Data-driven stock forecasting models based on neural networks: A review. *Information Fusion*, *113*, 102616. https://doi.org/10.1016/j.inffus.2024.102616

Bhandari, H. N., Rimal, B., Pokhrel, N. R., Rimal, R., Dahal, K. R., & Khatri, R. K. C. (2022). Predicting stock market index using LSTM. *Machine Learning with Applications*, 9, 100320. https://doi.org/10.1016/j.mlwa.2022.100320

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, *31*(3), 307–327.

Bonas, M., Datta, A., Wikle, C. K., Boone, E. L., Alamri, F. S., Hari, B. V., Kavila, I., Simmons, S. J., Jarvis, S. M., Burr, W. S., Pagendam, D. E., Chang, W., & Castruccio, S. (2024). Assessing predictability of environmental time series with statistical and machine learning models. *Environmetrics*. https://doi.org/10.1002/ env.2864

Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: Forecasting and control*. Wiley.

Cao, K., Zhang, T., & Huang, J. (2024). Advanced hybrid lstm-transformer architecture for real-time multi-task prediction in engineering systems. *Scientific Reports*, *14*, 4890. https://doi.org/10.1038/s41598-024-55483-x

Fatima, S. S. W., & Rahimi, A. (2024). A review of time-series forecasting algorithms for industrial manufacturing systems. *Machines*, *12*(6), 380. https://doi.org/10.3390/machines12060380

Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, *270*(2), 654–669.

Gülmez, B. (2025). A hybrid approach for stock market price forecasting using long short-term memory and seahorse optimization algorithm. *Annals of Data Science*. Advance online publication. https://doi.org/10.1007/s40745-025-00609-9

Hafshejani, M. S., & Mansouri, N. (2025). Enhancing stock market prediction with LSTM: A review of recent developments and comparative analysis. *Archives of Computational Methods in Engineering*. Advance online publication. https://doi.org/10.1007/s11831-025-10370-0

Hamou, O., Oudgou, M., & Boudhar, A. (2025). Analysis of the effectiveness of classical models in forecasting volatility and market dynamics: Insights from the masi and masi esg indices in Morocco. *Journal of Risk and Financial Management*, *18*(7), 370. https://doi.org/10.3390/jrfm18070370

Hanaki, N., Hommes, C., Kopa´nyi, D., Kopányi-Peuker, A., & Tuinstra, J. (2023). Forecasting returns instead of prices exacerbates financial bubbles. *Experimental Economics*, *26*, 1185–1213. https://doi.org/10.1007/s10683-023-09815-9

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780.

Huang, S., & Zhou, L. (2025). A residual-corrected hybrid arima–cnn–lstm framework for high-accuracy tobacco sales forecasting in regulated markets. *International Journal of Computational Intelligence Systems*, *18*, 194. https://doi.org/10.1007/ s44196-025-00930-4

Kim, H. Y., & Won, C. H. (2019). Forecasting stock prices with a hybrid arima and lstm model. *Expert Systems with Applications*, *129*, 273–281.

Kim, J., Kim, H., Kim, H., Lee, D., & Yoon, S. (2025). A comprehensive survey of deep learning for time series forecasting: Architectural diversity and open challenges. *Artificial Intelligence Review*, *58*, 216. https://doi.org/10.1007/s10462-025-11223-9

Kong, X., Chen, Z., Liu, W., Ning, K., Zhang, L., Marier, S. M., Liu, Y., Chen, Y., & Xia, F. (2025). Deep learning for time series forecasting: A survey. *International Journal of Machine Learning and Cybernetics*, *16*, 5079– 5112. https://doi.org/10.1007/s13042-025-02560-w

Kucuktopcu, E., Cemek, E., Cemek, B., & Simsek, H. (2023). Hybrid statistical and machine learning methods for daily evapotranspiration modeling. *Sustainability*, *15*(7), 5689. https://doi.org/10.3390/su15075689

Kumar, S., Rao, A., & Dhochak, M. (2025). Hybrid ml models for volatility prediction in financial risk management. *International Review of Economics & Finance*, *98*,103915. https://doi.org/10.1016/j.iref.2025.103915

Li, Q., Kamaruddin, N., Yuhaniz, S. S., & Al-Jaifi, H. A. A. (2024). Forecasting stock price changes using long-short term memory neural network with symbolic genetic programming. Scientific Reports, 14, 422. https://doi.org/10.1038/s41598-023-50783-0

Liu, H., Su, H., Sun, L., & Dias-da-Costa, D. (2024). State-of-the-art review on the use of ai-enhanced computational mechanics in geotechnical engineering. *Artificial Intelligence Review*, *57*, 196. https://doi.org/10.1007/s10462-024-10836-w

Liu, Z., Zhang, Z., & Zhang, W. (2025). A hybrid framework integrating traditional models and deep learning for multi-scale time series forecasting. *Entropy*, *27*(7), 695. https://doi.org/10.3390/e27070695

Malashin, I., Tynchenko, V., Gantimurov, A., Nelyub, V., & Borodulin, A. (2024). Applications of long short-term memory (lstm) networks in polymeric sciences: A review. *Polymers*, *16*(18), 2607. https://doi.org/10.3390/polym16182607

Mienye, I. D., Swart, T. G., & Obaido, G. (2024). Recurrent neural networks: A comprehensive review of architectures, variants, and applications. *Information*, *15*(9), 517. https://doi.org/10.3390/info15090517

Muhammad, D., Ahmed, I., Naveed, K., & Bendechache, M. (2024). An explainable deep learning approach for stock market trend prediction. *Heliyon*, *10*(21), e40095. https://doi.org/10.1016/j.heliyon.2024.e40095

Nelson, D. M., Pereira, A. C. M., & de Oliveira, R. A. (2017). Stock market's price movement prediction with lstm neural networks. *International Joint Conference on Neural Networks (IJCNN)*, 1419–1426.

Ozdemir, M. (2025). Asymmetric shock persistence in the oecd stock exchanges: New in-¨ sight from quantile exponential smooth transition autoregression approach. *Computational Economics*. https://doi.org/10.1007/s10614-025-10889-1

Pagliaro, A. (2025). Artificial intelligence vs. efficient markets: A critical reassessment of predictive models in the big data era. *Electronics*, *14*(9), 1721. https://doi.org/ 10.3390/electronics14091721

Ran, W., Tan, K., Zhang, Z., Pi, J., & Zhang, Y. (2025). Modeling temporal symmetry: Dual-component framework for trends and fluctuations in time series forecasting. *Symmetry*, *17*(4), 577. https://doi.org/10.3390/sym17040577

Rezaei, A., Abdellatif, I., & Umar, A. (2025). Towards economic sustainability: A comprehensive review of artificial intelligence and machine learning techniques in improving the accuracy of stock market movements. *International Journal of Financial Studies*, *13*(1), 28. https://doi.org/10.3390/ijfs13010028

Ryan, O., Haslbeck, J. M. B., & Waldorp, L. J. (2025). Non-stationarity in time-series analysis: Modeling stochastic and deterministic trends. *Multivariate Behavioral Research*, *60*(3), 556–588. https://doi.org/10.1080/00273171.2024.2436413

Saberironaghi, M., Ren, J., & Saberironaghi, A. (2025). Stock market prediction using machine learning and deep learning techniques: A review. *AppliedMath*, 5(3), 76. https://doi.org/10.3390/appliedmath5030076

Sherly, A., Christo, M. S., & Elizabeth, J. V. (2025). A hybrid approach to time series forecasting: Integrating arima and prophet for improved accuracy. *Results in Engineering*, *27*, 105703. https://doi.org/10.1016/j.rineng.2025.105703

Taheri, M., Bigdeli, M., Imanian, H., & Mohammadian, A. (2025). An overview of machine-learning methods for soil moisture estimation. *Water*, *17*(11), 1638. https: //doi.org/10.3390/w17111638

Tian, Y., Zheng, H., & Wang, C. (2022). A hybrid arima-lstm model for time series forecasting in financial markets. *Expert Systems with Applications*, *197*, 116701.

Tsay, R. S. (2010). *Analysis of financial time series. Wiley.*

Vancsura, L., Tatay, T., & Bareith, T. (2025). Navigating ai-driven financial forecasting: A systematic review of current status and critical research gaps. *Forecasting*, *7*(3), 36. https://doi.org/10.3390/forecast7030036

Vitale, J., & Robinson, J. (2025). In-season price forecasting in cotton futures markets using arima, neural network, and lstm machine learning models. *Journal of Risk and Financial Management*, *18*(2), 93. https://doi.org/10.3390/jrfm18020093

Wang, K. (2025). Multifactor prediction model for stock market analysis based on deep learning techniques. *Scientific Reports*, *15*, 5121. https://doi.org/10.1038/s41598025-88734-6

Waqas, M., Naseem, A., Humphries, U. W., Hlaing, P. T., Dechpichai, P., & Wangwongchai, A. (2025). Applications of machine learning and deep learning in agriculture: A comprehensive review. *Green Technologies and Sustainability*, *3*(3), 100199. https://doi.org/10.1016/j.grets.2025.100199

Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing*, *50*, 159–175.