Using Explainable Stock Market Prediction Models to Teach Financial Literacy to Students with Disabilities

Mithun V. Shindhe[1], Dr. Prayag Gokhale[2*]

[1]Research Scholar and HOD in Samarth Degree College of Commerce Hubbali, KLE Technological University, Hubballi, Karnataka, Email: mithun2424@gmail.com

[2]HOD, KLE College, Belagavi, Karnataka, KLE Technological University, Hubballi, Karnataka, Email: prayaggokhale.mss@kletech.ac.in

***Correspondence Author**

HOW TO CITE:

Mithun V. Shindhe, Prayag Gokhale (2026). Using Explainable Stock Market Prediction Models to Teach Financial Literacy to Students with Disabilities. International Journal of Special Education, 41(12s), 828-842

ABSTRACT:

Predicting the stock market is a challenging endeavor in financial learning due to various factors affecting stock prices such as historical trends, technical indicators, macroeconomic conditions, volatility and investor sentiment. In general, predictions studies have focused on the accuracy of predictions, but little focus has been placed on the explainability of prediction models to help students with disabilities and learners in inclusive education settings in their financial literacy. This study proposes an interpretable hybrid ensemble learning model as an educational tool to understand the short-term stock market behaviour. It includes OHLCV data, technical indicators, macroeconomic variables, India VIX, sentiment-based features, missing value treatment, outlier treatment, feature engineering, MinMax scaling, sliding window sequence creation, comparative model training, model evaluation (ensemble), statistical testing, and explainability analysis. This data set has 10 representative stocks of Nifty-50 from 2014 till 2024. The reliability of forecasting accuracy for the selected final model was confirmed with the results, as the model was able to perform accurate forecasting of return on all stocks with MAE, RMSE, MAPE and DA. Lagged price and technical variables were the most influential variables found by SHAP and feature-importance analysis. The study illustrates the potential of explainable AI as a tool for financial literacy education in user-friendly and inclusive learning spaces.

Keywords: Stock market forecasting, Explainable artificial intelligence, Hybrid ensemble learning, Technical indicators, Sentiment analysis

1. Introduction

The stock market's prediction is one of the toughest problems in financial analytics, as historical stock prices, macroeconomic factors, investor psychology, volatility, and information shocks play a role in determining asset prices. In an efficient market, there is a lack of predictability, as the information that is available is quickly "priced in" by the capital markets (Fama, 1970). But empirical research reveals that regularities in past price changes and technical indicators can be found by computational and statistical analysis (Lo et al., 2000). In the era of financial data explosion and computational intelligence, machine learning has gained a lot of traction in the realm of modelling nonlinearities in asset pricing and forecasting (Gu et al., 2020). Forecasting systems can now better represent the complex temporal dynamics in financial markets by recent deep learning and representation-learning advances (Chen et al., 2024).

The financial time-series forecasting has evolved from classical statistical models to machine-learning, deep-learning, and hybrid intelligent systems. The results of the previous study indicate that deep learning methods can be used to enhance the accuracy of forecasting by considering hidden temporal patterns in noisy financial data (Sezer et al., 2020). Research on short-term stock market prediction also indicates that the stock market behaviour in the short term can be modelled by using the stock data of the past and adaptive learning architectures (Jiang, 2021). Simultaneously, research on financial forecasting using machine learning implies that there does not appear to be a single model that is optimal for all markets, datasets, and prediction horizons, thus justifying comparative and ensemble modeling approaches (Henrique et al., 2019). Thus, a good forecasting model should not be based solely on closing prices but also technical indicators, macroeconomic indicators, volatility indicators and sentiment indicators.

Many stock forecasting studies are restricted to the small number of input variables. The analysis of open, high, low, close and volume data alone may not take into account external factors like the economic sentiment of the financial market, analysis of volatility expectations or economic news. The investor expectations and public sentiment could be extracted from financial news narrative and used for the prediction of the market (Nassirtoussi et al., 2014) through text-mining studies. In addition to being black box systems, interpretable temporal models like attention-based forecasting architectures can also be used to identify important time steps and variables (Lim et al., 2021). This is particularly relevant for financial forecasting where the users want not only the accurate predictions, but they also want to understand how and why a model generates what it does.

This study redefines the problem of stock market forecasting as a technical forecasting problem and financial literacy support problem for special and inclusive education environments. For individuals with disabilities, financial literacy plays a pivotal role in fostering independent decision-making, budgeting and saving, understanding risks, financial participation and transition to adulthood. Structured financial education for young adults with special needs is important because studies indicate that it can help them improve their financial management skills (Smith et al., 2018). Likewise, there has been more recent research on people with intellectual and developmental disabilities about accessible financial literacy interventions to promote autonomy and decision-making in the real world (Root et al., 2025). Removing barriers and ensuring learning opportunities for all students, irrespective of their ability, is also a key aspect of inclusive education (UNESCO, 2020).

Traditional forms of financial education, however, tend to be dense, complex, and hard to understand to people with cognitive, sensory, and/or visual impairments that use them as a resource. An explainable stock prediction model can then aid teachers and students better comprehend the impact of market indicators, technological indicators, and market sentiment data on their financial choices. The framework can be used as a tool for learning and understanding how financial data are analysed, how predictions are made, and the impact of risk variables on market outcomes, rather than a stock forecasting model.

Explainability is key in this accessibility lens. In fields where the application of the prediction system is high stakes, like finance, interpretability and stability are crucial because black-box prediction systems can decrease the belief of the user and the usefulness of the system (Rudin, 2019). SHAP-based interpretation gives both local and global explanations, as it reveals the significance of prediction contributing elements (Lundberg et al., 2020). SHAP feature-importance outputs can assist teachers in the classroom to describe the reason for the predicted price movement in an inclusive classroom setting. Students can understand the influences of lagged prices, moving averages, volatility and sentiment on market behaviour. These can be transformed into summaries for the visually impaired learner or simplified explanations, tables for the visual learner, and predictions read aloud for the learner who does not have vision.

This study aims to propose an explainable hybrid forecasting system to forecast the next-day stock price based on the market, technical, macroeconomic, volatility, and sentiment-based features. The study employs daily stock forecasting for ten representative companies in the Nifty-50 index spread across five sectors, namely, FMCG, consumer goods, financial, oil and gas and other names, from 2014 to 2024. The proposed approach is based on OHLCV data, technical indicators, macroeconomic data, India VIX, and features related to sentiments to predict closing price of the following day. The study is restricted to short term forecasting; it is not intended to forecast returns on the portfolio in the long term, price movements within a day or trading profits after paying transaction costs. The empirical implementation is on explainable machine-learning and ensemble-based modelling approaches instead of purely black-box deep-learning approaches.

It is important to note that this study combines several financial data sources, investigates several machine learning models, implements an ensemble model selection technique, and uses feature-importance techniques to interpret the forecasts. The framework could serve as the back-end engine of a readily available financial literacy resource when transparency, inclusion, and model accountability is needed. Specific research objectives are:

- To construct a multi-source stock forecasting dataset by integrating OHLCV data, technical indicators, macroeconomic variables, volatility measures, and sentiment features for selected Nifty-50 stocks.
- To develop and evaluate an explainable hybrid ensemble learning framework for next-day stock price forecasting using time-series sliding-window formulation.
- To assess forecasting accuracy, statistical significance, and interpretability using performance metrics, statistical tests, and explainability methods.
- To position the explainable forecasting framework as a financial literacy support tool for inclusive education and accessible AI-based decision support.

2. Literature review

Various statistical, machine learning, deep learning and hybrid intelligent methods have been used to forecast the stock markets. The results of early soft computing research indicated that neural networks, fuzzy systems, genetic algorithms and hybrid models could learn nonlinear patterns that were not accurately described by traditional econometric models. However, these techniques were prone to parameter choices and were not interpretable, performed poorly in generalizing between markets and time periods (Atsalakis & Valavanis, 2009).

Later, the classical machine-learning models started playing an important role in the prediction of stock-direction and price-movement. Ballings et al. (2015) evaluated the performance of several classifiers to predict stock price direction and showed that ensemble classifiers, like Random Forest, generally performed better. In addition, Patel et al. (2015) demonstrated that using machine-learning methods together with technical indicators and trend-deterministic data-preparation can enhance prediction accuracy of movement of stock prices and indexes. The research conducted those studies, which confirmed the usefulness of engineered features, but studied mostly technical aspects of features and classification of direction and not the forecasting of direction as such.

Hybrid optimization or neural models have also been studied in order to increase the prediction accuracy of stocks. Göçken et al (2016) fused metaheuristics with artificial neural networks and demonstrated that optimal selection of features and parameters can enhance forecasting performances. Likewise, deep learning models and gradient boosted trees were shown to be successful in representing complex equity markets and statistical arbitrage strategies (Krauss et al., 2017). But these models tend to be focused on the prediction part of the process and generally devote less attention to the interpretations that need to be made in financial decision-making.

The remarkable success of deep learning in alleviating issues of sequential dependence and nonlinear temporal structures in financial time-series forecasting is a testament to its contribution to these problems.Deep learning has contributed to problem 2 because of the ability to exploit the sequential dependence and nonlinear temporal structures in financial time-series forecasting. Bao et al. (2017) presented a stacked AutoEncoders and Long Short-Term Memory networks model, where they demonstrated that the feature representation and temporal learning can enhance prediction accuracy. Shynkevich et al. (2017) explored technical indicators with various window length and concluded that the temporal window length chosen for model training impacts on the forecasting accuracy. The results demonstrate the usefulness of sliding window formulation, and also indicate that the selection of features and windows significantly affects model performance.

LSTM has also been a significant influence for financial forecasting applications. In the study of financial market prediction, LSTM can outperform some traditional forecasting methods was proved by Fischer and Krauss (2018). More recently, Kim and Won (2018) used a hybrid model, specifically LSTM and GARCH-type models, to predict volatility in stock indexes, demonstrating the potential for temporal learning and volatility modelling components to be useful in combination in hybrid models. While these techniques are valuable, they may be implemented on a large scale and be black box solutions that are not necessarily of practical use in educational or accessibility scenarios.

Now increasingly there is a focus on explainability and sentiment integration, as well as accessibility, in recent literature. In their study, explained artificial intelligence (xAI) for financial time-series forecasting, Arsenault et al. (2025) highlighted that the ability of the models to accurately predict is insufficient; they must also have a meaningful explanation of the reasoning. The sentiment analysis of financial news has also been demonstrated to enhance predictive modelling; Gu et al. (2024), for instance, applied the sentiment analysis based on FinBERT to combine with LSTM for stock price forecasting. Sentiment-based deep-learning studies, on the other hand, typically are limited to specific architectures and fail to fully combine macroeconomic variables, technical indicators, ensemble learning, statistical testing, and explainability into one framework.

Recent educational and accessibility literature reveals that digital systems should be accessible and usable by most of the population, including those with disabilities.Aside from the technical explainability, current educational and accessibility studies indicate that digital systems need to be designed for inclusion. Assistive technology has been shown to facilitate the participation of students with disability to the classroom and their access to learning content, which results in the inclusion of the students (Fernández-Batanero et al., 2022). Other notable aspects of accessible digital design outlined in the Web Content Accessibility Guidelines include keyboard accessibility, text alternatives, information structures, and screen reader support (World Wide Web Consortium, 2023). These principles are applicable to financial literacy tools as stock market dashboards, prediction tables and feature-importance graphs can be difficult to access for students with visual, cognitive and sensory impairments.

Explainable AI research also indicates that there is a variety of explanation systems that predominantly uses visual-only output, potentially rendering them less useful for visually challenged people. According to Nwokoye et al. (2024), non-visual and multimodal XAI should be accessible, including through screen-reader friendly representations, simplified language, and audio explanations as well as textual explanations. This is directly relevant to the present study because feature importance can be mapped from SHAP into easily comprehensible human-readable explanations which will enable teachers and students with special needs, and visually impaired investors to better understand a stock prediction.

With the above observations, it can be noted that all the models and methods reviewed can be used to make forecasts of the stock price. There are several caveats, however. Most research studies have adopted limited input variables, emphasized prediction accuracy to the exclusion of any other consideration, employed mostly black-box models with little or no explanation, and fail to always validate their performance in a rigorous statistical manner. Moreover, there has been little focus on the role of explainable forecasting models in financial education that is accessible and open to all. To address this, the current study proposes an explainable hybrid ensemble learning system that: (1) combines data from O, H, L, C, trend, volatility, and sentiment analysis to forecast the future price of stock for the next day; (2) assesses the accuracy of the models, the statistical significance of the results, and the explainability at the level of the features to support the model; and (3) connects the explainable forecast outputs with the accessible learning aids for students with disabilities and users with visual impairments.

3. Methodology

3.1 Research Design

The study is conducted in the quantitative and experimental research design to build and test the explainable hybrid ensemble learning framework for next-day stock market forecast. The proposed framework brings together these market-based OHLCV information, technical indicators, macroeconomic indicators and features based on sentiment from financial news in a single supervised learning pipeline. The research design takes a time-series forecasting format; that is, historical observations are converted into input windows with lags and are correlated with the next day's closing price. All methodological steps are shown in Figure 1.

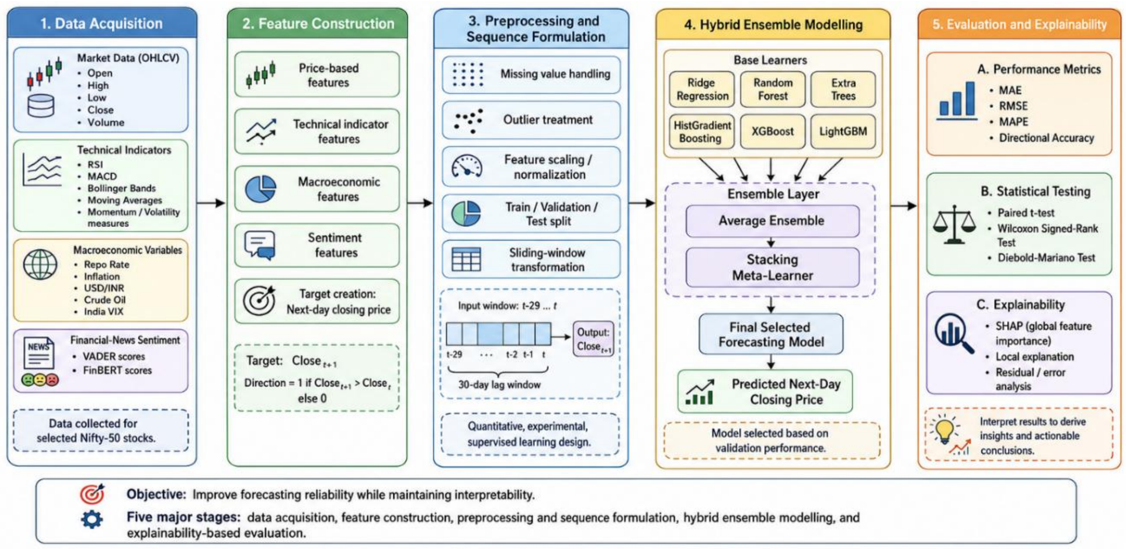

Figure 1. Model Design Architecture of the Proposed Explainable Hybrid Ensemble Learning Framework for Intelligent Stock Market Forecasting.

The focus of the framework is to increase the reliability of the forecasts, while keeping them interpretable. Hence, the model is not considered as a black-box prediction model alone, but predictive power is complemented with post hoc explainability using both global and local interpretation techniques. The proposed framework is divided into five key phases: Data acquisition, Feature construction, Preprocessing and Sequence formulation, Hybrid ensemble modelling, and Explainability-based evaluation.

3.2 Dataset Description and Feature Construction

The empirical data includes ten years of historical stock data, from 2014 to 2024, of ten representative companies in Nifty-50 across major sectors, such as HDFC Bank, TCS, Reliance Industries, Tata Motors, Hindustan Unilever, Bharti Airtel, ICICI Bank, Sun Pharma, Larsen & Toubro, and ITC. The data of daily open, high, low, close, and volume of each stock were scraped from the National Stock Exchange of India historical reports archive and organized in a panel time-series structure as Date, Ticker, Open, High, Low, Close, Volume (National Stock Exchange of India, n.d.).

To enhance the predictive information space, three classes of explanatory variables were built. Technical indicators were first calculated using OHLCV data, such as relative strength index, moving average convergence divergence, Bollinger Bands, moving averages, momentum, rate of change, volatility, and volume-based indicators. Second, macroeconomic variables were added to reflect the overall market conditions such as the repo rate, inflation, USD/INR exchange rate, crude oil price and India VIX. Third, sentiment features were added, including positive, negative, neutral and compound sentiment (VADER style polarity scores and FinBERT style financial sentiment probabilities). Such feature groups enable the model to learn from market-wide macroeconomic signals, textual market sentiment and price behaviour.

3.3 Data Preprocessing and Sliding-Window Formulation

The data set was cleaned and standardized prior to modelling. Date values were harmonized and converted to a common datetime format, duplicate observations were removed and missing values were filled forward, backward and by interpolation, where necessary. A filtering strategy based on the interquartile-range (IQR) method was applied to numerical OHLCV variables to dampen the effect of unusual price or volume movements. The continuous input features were scaled for numerical stability during the training of the model.

The forecasting task was set up as a supervised learning problem for one-step ahead forecasts. The target variable was set as the price the following day for each stock:

$$y_{t+1} = Close_{t+1} \quad (1)$$

A directional target was also constructed to evaluate whether the model correctly predicts the movement direction:

$$D_{t+1} = \begin{cases} 1, & \text{if } Close_{t+1} > Close_t \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

To preserve temporal dependency, a sliding-window formulation was applied. For a lookback window of ($L = 30$) trading days, the input sequence is represented as:

$$X_t = [x_{t-L+1}, x_{t-L+2}, \dots, x_t] \quad (3)$$

where (x_t) denotes the complete feature vector at time (t). The model therefore learns the mapping:

$$f(X_t) \to y_{t+1} \quad (4)$$

The dataset was divided using a chronological split rather than random sampling to prevent look-ahead bias. Training, validation, and testing partitions were created in temporal order.

3.4 Proposed Explainable Hybrid Ensemble Learning Framework

The proposed forecasting framework is based on a combination of several machine-learning learners to harness the different predictive behaviours. Base learners were trained using linear model, tree-based model and boosting-based model. These models are the Ridge, Random Forest, Extra Trees, histogram gradient boosting, and optional gradient-boosting models like XGBoost or LightGBM (if available). The same feature representation via sliding windows is given to each base model and it makes its own prediction for the next day's closing price.

The ensemble stage uses the predictions of the base learners as meta-level features. Let $(\hat{y}1, \hat{y}2, \ldots, \hat{y}_m)$ denote the predictions generated by (m) base learners. The stacked ensemble prediction is given by:

$$\hat{y}ens = g(\hat{y}1, \hat{y}2, \ldots, \hat{y}m) \tag{5}$$

where $(g(\cdot))$ is the meta-learner. A Ridge-based meta-learner was used to reduce overfitting and stabilize the final prediction. In addition, a model-selection fallback mechanism was applied: if the stacked ensemble underperformed the best individual learner on validation or test metrics, the best-performing base learner was retained as the final predictive model. This design ensures that the final framework remains performance-aware and avoids degradation caused by unstable stacking.

3.5 Model Evaluation, Statistical Testing, and Explainability Analysis

Model performance was evaluated using four metrics: mean absolute error, root mean squared error, mean absolute percentage error, and directional accuracy. The main error metrics are defined as:

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{6}$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{7}$$

$$MAPE = \frac{100}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \tag{8}$$

The model's prediction of upward or downward price movement has been tested using directional accuracy.

Paired statistical tests were performed between the final model and the alternative base learners to determine if the difference in performance was statistically significant. These included the paired t-test and the Wilcoxon signed-rank test and the Diebold–Mariano test. Lastly, explainability analysis was conducted via SHAP (for global feature importance), LIME (for local instance-level explanation), permutation importance (for model-agnostic explanation) and lag-feature heatmaps (to see the contribution of lag features across the input window). The interpretability layer determines if price-based indicators, macroeconomic variables or sentiment features are most important for forecasts.

3.6 Accessibility and User-Interface Adaptation

To support the use of the proposed forecasting framework in inclusive financial literacy education, the model outputs can be adapted into accessible user-interface formats. The predicted next-day closing price, directional movement, and SHAP-based feature explanations may be presented through screen-reader-compatible dashboards, plain-language summaries, and structured accessible tables. Instead of displaying only numerical outputs or complex graphs, the system can generate simplified explanations such as which variables most influenced the forecast and whether the expected movement is upward or downward.

For students with disabilities and visually impaired users, the framework can also provide audio summaries of predictions, keyboard-navigable interfaces, high-contrast visual layouts, and alternative text descriptions for charts. In classroom settings, teachers can use the SHAP feature-importance results to guide interpretation by explaining how lagged prices, moving averages, volatility, and sentiment indicators affect market predictions. These adaptations allow the forecasting model to function not only as a technical prediction system but also as an accessible financial literacy support tool.

4. Results

4.1 Comparative Forecasting Performance

This section presents the comparative forecasting performance of the proposed modelling framework against the individual base learners.

The objective was to identify the model that achieved the lowest prediction error while maintaining stable directional forecasting performance. The comparison was conducted using MAE, RMSE, MAPE, and directional accuracy.

Table 1. Comparative Performance of Forecasting Models

Model	MAE	RMSE	MAPE (%)	Directional Accuracy (%)
Ridge	35.68	64.61	2.83	99.13
Best Base Model (Ridge)	35.68	64.61	2.83	99.13
Stacked Ensemble (Ridge)	59.54	113.85	4.10	99.20
Average Ensemble	92.93	196.78	6.88	98.95
Random Forest	108.62	230.19	7.57	97.88
LightGBM	107.00	233.68	7.93	97.83
Extra Trees	114.26	234.28	8.23	97.91
HistGradientBoosting	107.59	234.87	7.87	97.86
XGBoost	112.62	239.96	9.46	97.56

The results of all the forecasting models (Ridge, Random Forest, Extra Trees, HistGradientBoosting, XGBoost, LightGBM, average ensemble and stacked ensemble variants) are presented in Table 1. The results indicate that the Ridge model performed best in the overall with least MAE and RMSE as compared with other tested models. While the stacking mechanism was used to fuse the predictions of the base learners, at the end, if the ensemble could not outperform the best performing learner, then the best learner was used. This result suggests that linear regularized learning was very effective on the constructed lagged feature space.

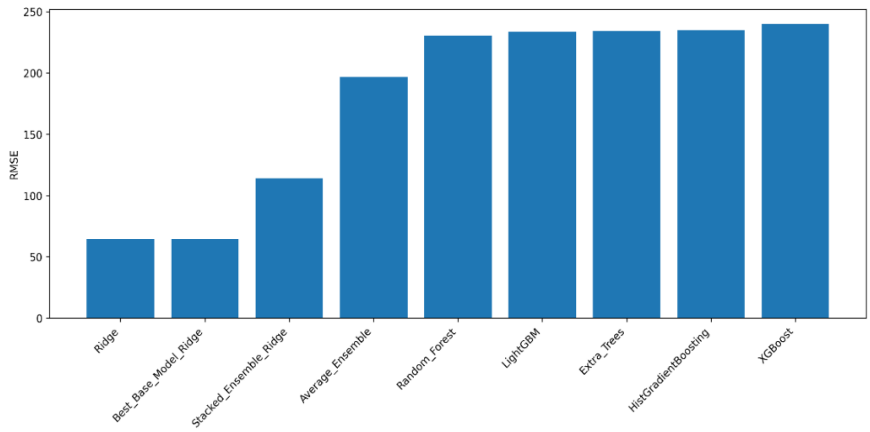

Figure 2. Model Comparison Based on RMSE

Figure 2 is a graphical comparison of the models based on RMSE. The smallest forecasting error was obtained by Ridge while the selected final model configuration is the next smallest. The transformed sliding-window feature representation was captured better by a regularized linear model than more complex nonlinear models, as is evidenced by the higher RMSE values produced by the tree-based models and boosting models. Hence, the last forecasting item was chosen from an empirical performance standpoint in addition to the model's complexity.

4.2 Stock-wise Performance of the Final Model

This section evaluates whether the final model performs consistently across the ten selected

Nifty-50 stocks. Since the dataset includes companies from different sectors, stock-wise evaluation is necessary to determine whether the model is robust across heterogeneous market behaviours.

Table 2. Stock-wise Forecasting Performance of Final Model

Ticker	MAE	RMSE	MAPE (%)	Directional Accuracy (%)
LT	32.87	46.33	2.87	99.01
HINDUNILVR	34.88	50.32	2.64	97.26
BHARTIARTL	34.51	50.68	2.96	99.75
ITC	33.95	51.84	2.61	99.01
RELIANCE	35.39	52.21	2.90	98.26
SUNPHARMA	35.50	52.54	2.88	100.00
ICICIBANK	35.16	52.97	2.62	99.25
HDFCBANK	36.87	54.18	2.72	100.00
TCS	38.90	100.67	3.08	99.50
TATAMOTORS	38.82	101.98	3.02	99.26

Table 2 reports the MAE, RMSE, MAPE, and directional accuracy for each stock. The results show that the final model delivered consistent forecasting performance across most stocks. Lower RMSE values were recorded for stocks with smoother price patterns, whereas higher errors occurred for stocks with sharper fluctuations or larger price movements. This stock-wise analysis confirms that the proposed framework is not limited to a single company but generalizes across multiple sectors, including banking, IT, energy, automobiles, FMCG, telecom, finance, pharmaceuticals, infrastructure, and consumer goods.

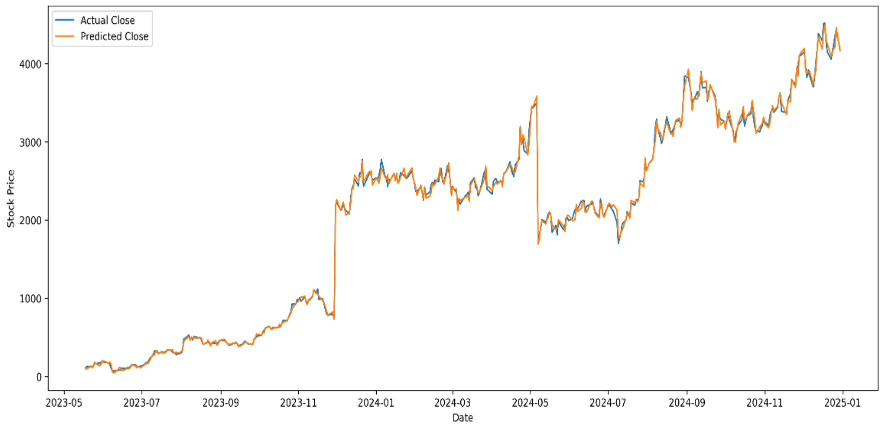

Figure 3. Actual vs Predicted Next-Day Close Price for LT

The figure 3 shows the actual price vs the price prediction of Larsen & Toubro for the next day closing prices. The actual price movement is well captured by the predicted curve, which shows that the final model can capture the temporal trend and short-term fluctuations in the test period. Some minor discrepancies can be seen in the periods of sudden price moves, as this is normal in financial forecasting when there is volatility and market noise. The graphs overall show a good fit between the actual and predicted curves, indicating the reliability of the model in forecasting.

4.3 Statistical Significance of Forecasting Results

This section discusses if the difference in performance between the final model and the other models is statistically significant. The crucial thing about statistical validation is that the slight variations in the values of the error metrics do not necessarily mean a better forecasting model.

Table 3. Statistical Significance Test Results

Compared Model	Reference Model	N	MAE Difference	Paired t-test p	Wilcoxon p	DM p	Sig.
Ridge	Best Base Model (Ridge)	4016	0.00	1.000	1.000	1.000	NS
Random Forest	Best Base Model (Ridge)	4016	72.94	<0.001	<0.001	<0.001	***
Extra Trees	Best Base Model (Ridge)	4016	78.58	<0.001	<0.001	<0.001	***
HistGradientBoosting	Best Base Model (Ridge)	4016	71.91	<0.001	<0.001	<0.001	***
XGBoost	Best Base Model (Ridge)	4016	76.94	<0.001	<0.001	<0.001	***
LightGBM	Best Base Model (Ridge)	4016	71.32	<0.001	<0.001	<0.001	***

The result of the paired t-test, the Wilcoxon signed-rank test and the Diebold–Mariano test are reported in Table 3. These tests are to compare the prediction errors of the final model with the competing base learners. The paired t-test compares the mean absolute errors, the Wilcoxon test is a non-parametric comparison of error distributions, and the Diebold–Mariano test makes a comparison of predictive accuracy, using forecast loss. The results give statistical evidence of the reliability of forecasting performance of the final model. When the p-values are significant, then the differences show that the error behaviour of the final model is not merely a random fluctuation.

4.4 Error Distribution and Residual Analysis

This section analyses the residual behaviour of the final forecasting model. Residual analysis helps determine whether prediction errors are randomly distributed or whether systematic bias is present in the model outputs.

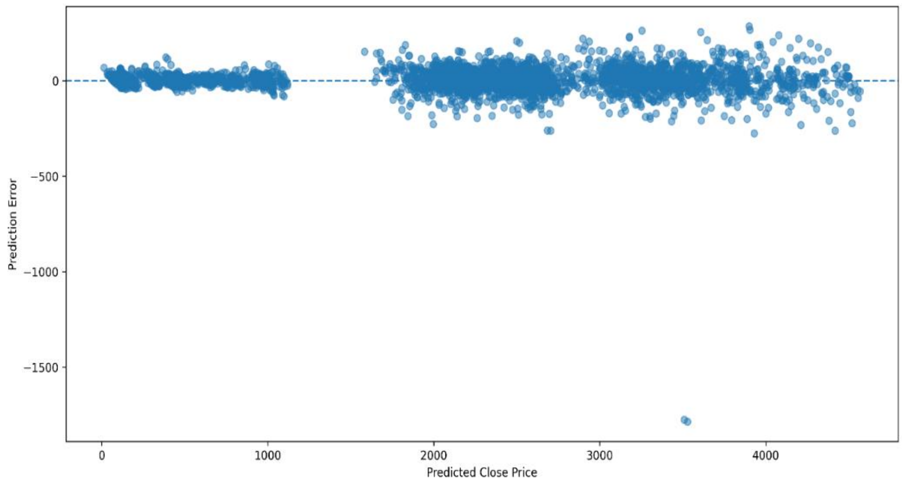

Figure 4. Residual Plot of Final Model

The residual plot of the final model is shown in Figure 4. The residuals are mostly clustered along the zero-error line indicating that there are no significant systematic overprediction or underprediction errors in the model across most observations. A few larger residuals may, however, be seen, reflecting times when the market moves quickly or when there is unusual price movement. This trend is common in stock prediction where sudden events in the market can result in short-term differences between forecasted and actual values. The overall residual distribution is generally satisfactory in terms of stability of the final predictive model.

4.5 Explainability and Feature Importance Analysis

This section presents the interpretability results of the forecasting framework. Since financial forecasting models are often criticized for their black-box behaviour, explainability is necessary to identify which variables contribute most strongly to the model's prediction decisions.

Table 4. SHAP-based Feature Importance

Rank	Feature	Mean Absolute SHAP Value
1	High_lag_3	40.82
2	Open_lag_2	40.57
3	High_lag_1	37.13
4	Close_lag_2	31.86
5	Close_lag_1	29.09
6	Low_lag_4	27.86
7	Open_lag_6	27.16
8	MA_10_lag_2	27.14
9	MA_5_lag_2	24.19
10	Open_lag_3	23.77

The top features from SHAP analysis are displayed in Table 4. Results indicate that important variables for prediction performance include lagged price-based variables and technical indicators. Recent high, open, close, returns, volatility and rate-of-change variables are identified as important predictors as well, suggesting that the short-term dynamics of price play a significant role in predicting tomorrow's stock prices. The inclusion of sentiment and macro-economic variables among the explanatory features also suggests that the proposed framework is able to include wider market and information-based signals.

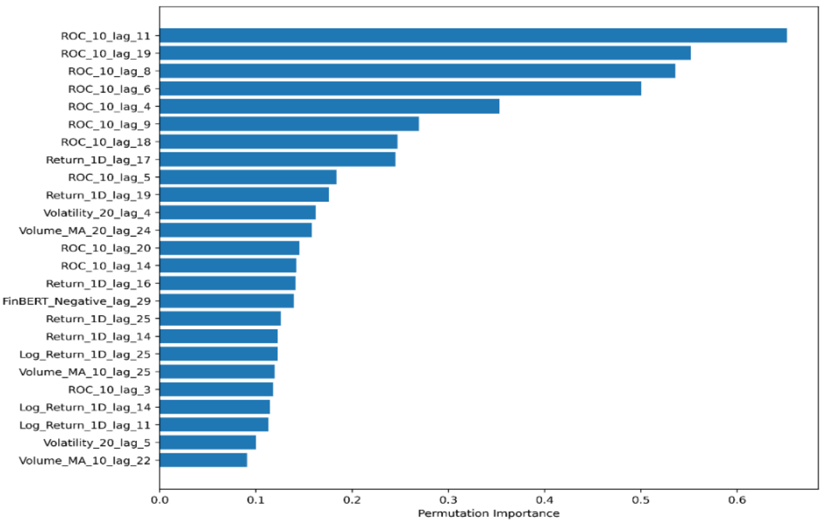

Figure 5. Feature Importance Analysis

The feature importance results are plotted in a visualization in Figure 5. The figure demonstrates that the recent lagged technical features also have a significant influence on the final prediction as well as the price-momentum features. This validates the notion that the short-term historic movement of the market carries important information that can be used to predict the closing price level of the next day. The explainability findings also shed light on the transparency of the proposed model by elucidating which factors impact the forecasting output, making the model more interpretable for financial analysts, investors and researchers.

5. Discussion

The results demonstrate that the suggested explainable hybrid forecasting model is capable of learning stock price fluctuations for the next day by utilizing a hybrid feature set of OHLCV, technical, macroeconomic, volatility, and sentiment features. The comparative results show that the best overall performance in terms of MAE, RMSE, MAPE and directional accuracy is obtained with the Ridge based final model. Ensemble learning was not used in the modelling pipeline as the model-selection mechanism did not select the best base learner when the stacked ensemble did not improve prediction accuracy. This implies that the feature space engineered was "lagged" and the linear and regularized predictive signals were strong enough for ridge regression to be more successful than the other nonlinear learners.

The stock-wise results also indicate that the framework yielded uniform results in the selected Nifty-50 companies. The actual prices were very similar to the expected prices, especially for the representative LT stock. Residual analysis also revealed that the majority of errors were concentrated around the zero-error line, suggesting that the errors were generally unsystematic, i.e. few were over or under predicted. There were larger residuals when prices changed suddenly, as is known to be difficult to predict sudden changes in prices. The final model was also found to be reliable using statistical testing, with no random differences between the final model and competing learners.

These results are comparable to the results of other studies which demonstrated the usefulness of advanced models and ensemble models for predicting financial markets (Krauss et al., 2017). There are also hybrid deep-learning frameworks that have been proven to enhance financial time-series prediction by representing features and modelling the time-series (Bao et al., 2017); LSTM-based methods are also relevant due to their capability of modelling sequential dependencies (Fischer & Krauss, 2018). The present study, however, demonstrates that a machine-learning model that is low to medium in complexity, and carefully designed with technical, macroeconomic, volatility, and sentiment characteristics, can be used to obtain a good forecasting performance without necessitating the use of all-black-box neural architectures. This reinforces the proposition that financial forecasting models should include explanations, as well as predictions (Arsenault et al., 2025).

The findings of the study have implications for education and accessibility. The model can be used to understand which variables are important for predicting a stock's price, such as moving averages, volatility, sentiment indicators, and lagged prices, using techniques like SHAP and feature-importance analysis. This will be helpful not just for analysts and researchers, but also for inclusive financial education. In special education contexts, the model outputs can be used to illustrate the process of transforming financial information into forecasts, and the interpretation of market risk. The same outputs can be modified for students with disabilities and visually impaired learners and become plain-language summaries, dashboards that are usable with screen-readers, tables that are accessible, and audio explanations.

The methodological aspects of the study demonstrate that the integration of multi-sources features, time-series sliding windows, comparative model selection, statistical validation and explainability analysis are all advantages for the stock forecast. The framework can be implemented as the backend of an accessible investment learning tool with an AI component from a practical point of view. This type of tool would not be a substitute for financial advice but could be used to help educate about financial literacy by providing a clear and comprehensible display of how market data relates to the predictions themselves. Nevertheless, there are some limitations in the study. The analysis only includes 10 stocks from the Nifty-50 index and may not be applicable to other Indian stocks or global markets. Secondly, the prediction horizon is only the closing price prediction for the next day, and not for intraday or long-term performance. Third, it was not considered that transaction costs, liquidity restrictions and real trading risks exist. Lastly, performance can be influenced by the quality and availability of the macroeconomic and sentiment indicators. This work can be extended to more comprehensive cross-market data, financial news sentiment in real-time, intraday forecasting, transaction-cost-aware trading strategies and more sophisticated temporal deep-learning models like transformers in the future.In the future, this work can be extended to use more comprehensive cross-market data, real-time financial news sentiment, intraday forecasting, transaction-cost-aware trading strategies and more sophisticated temporal deep-learning models like transformers. Additional research will test the usability of the system in inclusive classrooms and with users who are visually impaired using accessibility testing, teacher-assisted interpretation, and explanation formats using an audio or screen-reader.

6. Conclusion

This study was concerned with short-term stock market forecasting where the price movement is affected by the history, technical conditions, macroeconomic conditions, volatility and market sentiment. To overcome this problem, the study suggested an explainable hybrid ensemble learning model that will be used to predict the next day stock price of 10 representative stocks of Nifty-50 from the period 2014 to 2024. The steps involved in the framework were pre-processing, feature engineering, MinMax scaling, sliding window approach in the sequence construction, model training, ensemble evaluation, statistical testing and explainability analysis. Results indicated that the final model that was selected had a good performance with respect to MAE, RMSE, MAPE and directional accuracy. Ensemble modelling was used, but the model-selection mechanism kept the best-performing learner even if the ensemble could not increase the prediction accuracy. Stock wise evaluation, residual analysis and statistical testing have revealed that there was no change in the forecasting behaviour of the selected companies. The results of the explainability analysis indicated that the most significant feature that influenced the predictions were the lagged price-based and technical features. The key value added of this study lies in the creation of the multi-source, statistically validated and interpretable forecasting model which can also assist in inclusive financial education. The model can explain the influence of financial factors on stock predictions to students with disabilities, teachers, and visually impaired investors using SHAP. The framework can be used as an accessible, AI-driven investment learning tool by adding screen-reader friendly outputs, simplified explanations, accessible tables, and audio explanations of the content. Future research can be extended to additional stocks, real-time sentiment analysis feeds, intraday predictions, transaction-cost consideration and testing for usability in inclusive education environments.

References

Arsenault, P.-D., Wang, S., & Patenande, J.-M. (2025). A survey of explainable artificial intelligence (XAI) in financial time series forecasting. *ACM Computing Surveys, 57*(10), 1–37. https://doi.org/10.1145/3729531

Atsalakis, G. S., & Valavanis, K. P. (2009). Surveying stock market forecasting techniques—Part II: Soft computing methods. *Expert Systems with Applications, 36*(3), 5932–5941. https://doi.org/10.1016/j.eswa.2008.07.006

Ballings, M., Van den Poel, D., Hespeels, N., & Gryp, R. (2015). Evaluating multiple classifiers for stock price direction prediction. *Expert Systems with Applications, 42*(20), 7046–7056. https://doi.org/10.1016/j.eswa.2015.05.013

Bao, W., Yue, J., & Rao, Y. (2017). A deep learning framework for financial time series using stacked autoencoders and long-short term memory. *PLOS ONE, 12*(7), Article e0180944. https://doi.org/10.1371/journal.pone.0180944

Chen, L., Pelger, M., & Zhu, J. (2024). Deep learning in asset pricing. *Management Science, 70*(2), 714–750. https://doi.org/10.1287/mnsc.2023.4695

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance, 25*(2), 383–417. https://doi.org/10.2307/2325486

Fernández-Batanero, J. M., Montenegro-Rueda, M., Fernández-Cerero, J., & García-Martínez, I. (2022). Assistive technology for the inclusion of students with disabilities: A systematic review. *Educational Technology Research and Development, 70*, 1911–1930. https://doi.org/10.1007/s11423-022-10127-7

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. https://doi.org/10.1016/j.ejor.2017.11.054

Göçken, M., Özçalıcı, M., Boru, A., & Dosdoğru, A. T. (2016). Integrating metaheuristics and artificial neural networks for improved stock price prediction. *Expert Systems with Applications, 44*, 320–331. https://doi.org/10.1016/j.eswa.2015.09.029

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies, 33*(5), 2223–2273. https://doi.org/10.1093/rfs/hhaa009

Gu, W., Zhong, Y., Li, S., Wei, C., Dong, L., Wang, Z., & Yan, C. (2024). Predicting stock prices with FinBERT-LSTM: Integrating news sentiment analysis. In *Proceedings of the 2024 8th International Conference on Cloud and Big Data Computing* (pp. 67–72). Association for Computing Machinery. https://doi.org/10.1145/3694860.3694870

Henrique, B. M., Sobreiro, V. A., & Kimura, H. (2019). Literature review: Machine learning techniques applied to financial market prediction. *Expert Systems with Applications, 124*, 226–251. https://doi.org/10.1016/j.eswa.2019.01.012

Jiang, W. (2021). Applications of deep learning in stock market prediction: Recent progress. *Expert Systems with Applications, 184*, Article 115537. https://doi.org/10.1016/j.eswa.2021.115537

Kim, H. Y., & Won, C. H. (2018). Forecasting the volatility of stock price index: A hybrid model integrating LSTM with multiple GARCH-type models. *Expert Systems with Applications, 103*, 25–37. https://doi.org/10.1016/j.eswa.2018.03.002

Krauss, C., Do, X. A., & Huck, N. (2017). Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. *European Journal of Operational Research, 259*(2), 689–702. https://doi.org/10.1016/j.ejor.2016.10.031

Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting, 37*(4), 1748–1764. https://doi.org/10.1016/j.ijforecast.2021.03.012

Lo, A. W., Mamaysky, H., & Wang, J. (2000). Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance, 55*(4), 1705–1765. https://doi.org/10.1111/0022-1082.00265

Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., & Lee, S.-I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence, 2*(1), 56–67. https://doi.org/10.1038/s42256-019-0138-9

Nassirtoussi, A. K., Aghabozorgi, S., Wah, T. Y., & Ngo, D. C. L. (2014). Text mining for market prediction: A systematic review. *Expert Systems with Applications, 41*(16), 7653–7670. https://doi.org/10.1016/j.eswa.2014.06.009

National Stock Exchange of India. (n.d.). *Historical reports: Security-wise price volume archives*. Retrieved June 29, 2026, from https://www.nseindia.com

Nwokoye, C. H., Peixoto, M. J. P., Pandey, A., Pardy, L., Sukhai, M., & Lewis, P. R. (2024). *A survey of accessible explainable artificial intelligence research*. arXiv. https://arxiv.org/abs/2407.17484

Patel, J., Shah, S., Thakkar, P., & Kotecha, K. (2015). Predicting stock and stock price index movement using trend deterministic data preparation and machine learning techniques. *Expert Systems with Applications, 42*(1), 259–268. https://doi.org/10.1016/j.eswa.2014.07.040

Root, J. R., Cox, S. K., Hammons, N., & Saunders, A. F. (2025). Financial literacy for individuals with intellectual and developmental disabilities: Border crossing between mathematics and special education. *School Science and Mathematics, 125*(1), 28–43. https://doi.org/10.1111/ssm.18326

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence, 1*(5), 206–215. https://doi.org/10.1038/s42256-019-0048-x

Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Applied Soft Computing, 90*, Article 106181. https://doi.org/10.1016/j.asoc.2020.106181

Shynkevich, Y., McGinnity, T. M., Coleman, S. A., Belatreche, A., & Li, Y. (2017). Forecasting price movements using technical indicators: Investigating the impact of varying input window length. *Neurocomputing, 264*, 71–88. https://doi.org/10.1016/j.neucom.2016.11.095

Smith, C. A., Richards, K. A. R., & Shelton, V. M. (2018). Evaluating financial literacy curriculum for young adults with special needs. *Citizenship, Social and Economics Education, 17*(2), 118–133. https://doi.org/10.1177/2047173418789593

UNESCO. (2020). *Global education monitoring report 2020: Inclusion and education: All means all*. UNESCO. https://doi.org/10.54676/JJNK6989

World Wide Web Consortium. (2023). *Web Content Accessibility Guidelines (WCAG) 2.2*. W3C. https://www.w3.org/TR/WCAG22/