International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Stock Market Prediction Using Machine Learning: A Review of Techniques, Applications, Challenges, and Future Directions

Manisha Devi[1], Ashish Kumar[2]
[1]Department of Mathematics, [2]Department of Computer Science
[1,2]P.I.G. Govt. College for Women, Jind, Haryana, India

Abstract

Stock market prediction is an interesting subject of financial science. It is worth noting that correct forecasts of the movements of stocks can help in the analysis of investments, portfolio management, and risk assessment in finance. Predicting the stock market is a difficult task because the factors influencing financial data are too numerous. Among such factors, one can mention company performance, economic state, investors' sentiments, government regulation, unforeseen events, etc. Statistical methods were very helpful for financial forecasting in the past, but their efficiency in capturing non-linear relationships between data can be restricted in some cases. Machine learning offers an alternative approach that allows identifying the patterns in historical data and analyzing relationships between several financial variables at once. The current review covers the issue of using machine learning methods for stock market prediction. It will focus on the use of Support Vector Machines, Decision Trees, Random Forest, Artificial Neural Networks, Gradient Boosting, and Long Short-Term Memory networks. Moreover, this review touches upon the importance of data preprocessing, technical indicators, fundamental variables, financial news, and sentiment analysis. Emphasis is especially laid on the difference between stock price prediction, return forecasting, direction forecasting, and development of profitable trading strategies. According to the review, there is no universally best algorithm as prediction efficiency depends on the data set, forecasting period, feature selection, and validation methodology. In addition, the review emphasizes the necessity to use chronological validation, make comparisons with relevant benchmarks, and realistically take into account transaction costs. Lastly, the review outlines possible areas of future studies that include explainable artificial intelligence, hybrid modeling, multimodal financial data, and forecasting systems.

Keywords: Stock Market Prediction, Machine Learning, Financial Forecasting, Artificial Neural Networks, Support Vector Machines, Random Forest, Deep Learning, LSTM, Sentiment Analysis, Predictive Analytics.

1. Introduction

The stock market is an integral component of the current financial world. The stock market helps companies to raise money, while investors are able to grow financially as well. Stock market data is used by individuals, financial organizations, portfolio managers, researchers and even policymakers for understanding the behavior of the stock market and making financial decisions. Since the price of the stock fluctuates all the time, forecasting becomes a challenging research topic.

Predicting stock market implies making projections about future stock prices, returns or trends based on existing information. These predictions may involve the prediction of the closing price of a company for the next trading day, the return for a certain period or whether the market is going up or down. These predictions are useful for investment purposes; however, their practical value is subject to the accuracy, reliability and generalizability of the predictions.

Forecasting financial markets is difficult since there are many interdependent variables that affect stock prices. This list includes company earnings, industry news, interest rates, inflation, exchange rates, politics, expectations of investors, and the state of the economy on a global level. Additionally, financial markets can be very sensitive to any new information, making past observations less valuable for forecasting purposes. The same interdependence of market variables behaves differently under different economic circumstances.

Traditional forecasting methods, including regression analysis and time-series models, have been widely used to study financial data. These approaches remain valuable because they provide useful statistical foundations and can be relatively easy to interpret. Nevertheless, some financial relationships may be nonlinear or may involve interactions among several variables. Machine learning offers methods that can learn such relationships from data without requiring every relationship to be specified in advance.

As a result, machine learning has now become an area of significant research importance for making predictions in finance. The algorithms like Support Vector Machine, Random Forest, Gradient Boosting, and Artificial Neural Network can be used to predict numeric values or to classify future movement in the market. Recurrent Neural Network and LSTM network are some deep learning techniques that have gained much attention recently.

Kumbure et al. (2022) reviewed 138 journal articles on machine learning techniques and data used for stock market forecasting. Their review examined the markets studied, the input variables used, and the methods employed by researchers. It highlighted the widespread use of technical indicators and the importance of neural networks and support vector methods. Sonkavde et al. (2023) subsequently reviewed machine learning and deep learning approaches, considering their applications, performance, and implications for financial forecasting.

Research has also examined the usefulness of machine learning in asset pricing and return prediction. Gu et al. (2020) investigated machine learning methods using a broad collection of financial characteristics and macroeconomic variables. More recent work by Kang, Ryu, and Webb (2025) compared machine learning approaches for predicting stock returns in the Korean market and examined the importance of explanatory variables. Such research demonstrates the value of evaluating algorithms comparatively rather than assuming that one technique will always produce the best predictions.

Nevertheless, the advent of such methods does not mean that the unpredictability of financial markets is capable of being solved through the use of machine learning techniques. The model can find patterns in historical data that will not be valid in the future, or the model can deliver seemingly good results due to overfitting, incorrect data partitioning, or inclusion of data that was not known at the time of the prediction.

In this paper, the main machine learning methods utilized in predicting the stock market will be reviewed. Also, the types of data and features used in stock market prediction will be examined. The pros and cons of the various methods will also be compared and some major research challenges highlighted. The main objectives of this review are:

- To examine the development of machine learning methods for stock market prediction.
- To discuss the major algorithms used in financial forecasting.
- To review the role of historical prices, technical indicators, fundamental data, and investor sentiment.
- To explain the evaluation measures used to assess predictive performance.
- To identify limitations in existing approaches and common methodological problems.
- To discuss emerging research directions for developing more reliable and interpretable forecasting systems.

2. Understanding Stock Market Prediction

2.1 Meaning and scope

The concept of stock market prediction consists of forecasting future performance of the stock market based on past data and other pertinent information. The domain of prediction varies based on the aim of the study. While some predict the stock value numerically, others are interested in determining if the stock will increase or decrease.

Price prediction is concerned with the forecasting of future prices. Return prediction forecasts the change in value for a specific period. Directional prediction divides the future behavior of the stock into groups like an increase or decrease. Volatility prediction forecasts future price changes.

However, while related, they do differ from each other. An accurate predictive model for price levels may not be a profitable strategy. Likewise, a correct directional prediction does not mean good trading performance if the size of the price movement cannot be predicted or if high trading costs occur. Thus, the research problem needs to be defined prior to the selection of the algorithm.

2.2 Factors influencing stock market movements

Stock prices are affected by several categories of information.

i. Company-specific factors: Earnings, revenue growth, profitability, debt levels, management decisions, product developments, and business performance can influence investors' expectations about a company's future prospects.

ii. Macroeconomic factors: Interest rates, inflation, economic growth, currency movements, and monetary policy can influence corporate earnings, borrowing costs, and investment decisions.
iii. Market-related factors: Trading volume, historical returns, volatility, momentum, and broader market movements may provide information about recent trading behaviour.
iv. Investor sentiment: Expectations, uncertainty, confidence, fear, and reactions to news can influence demand and supply for securities.
v. Unexpected events: Geopolitical developments, natural disasters, regulatory announcements, and other unforeseen events can produce substantial price movements that are difficult to anticipate from historical observations alone.

Such factors could be included in machine learning models, given the availability of appropriate data. However, the addition of other factors does not necessarily improve prediction accuracy. The significance of the factors needs to be demonstrated by rigorous evaluation.

2.3 Financial market uncertainty

There is a high degree of uncertainty and interrelationships in financial markets. There may be some patterns in historical data, but such patterns become less strong when there are changes in the economy or behavior of investors. It becomes quite clear that forecasting is a completely different concept from prediction under stable conditions.

One of the theoretical approaches to market predictability is the Efficient Market Hypothesis introduced by Fama (1970). The hypothesis implies that the degree of informational reflection in the price level influences earning of abnormal profits from using this information.

Machine learning techniques can test whether there is anything useful to be learned from available data. Yet, finding a useful pattern in past observations is not necessarily an indicator that the pattern will be useful in the future as well.

3. Traditional Approaches to Stock Market Prediction

Before the widespread use of machine learning, stock market forecasting relied primarily on statistical models and technical analysis. These approaches remain relevant because they provide foundations for understanding market behaviour and benchmarks for evaluating more complex algorithms.

3.1 Statistical forecasting

The statistical forecasting techniques are based on using mathematics to model the observed data in the past and predict future values. Linear regression, autoregressive models, moving averages, and other time series forecasting techniques can be applied in finance.

Linear regression attempts to find a relationship between explanatory variables and a dependent variable. Autoregressive models attempt to predict future value from historical values, whereas moving averages represent price action. There are numerous benefits in using these techniques. First of all, they are usually much simpler to apply than complicated deep learning algorithms, and it is easy to verify assumptions underlying these methods. However, they may not work well for nonlinear models.

Statistical methods should not be regarded as obsolete. They remain useful benchmarks for determining whether a machine learning model provides meaningful improvement.

3.2 Technical analysis

The idea of technical analysis is based on studying historical price fluctuations, trading volume and technical indicators that may signal possible trends or trade signals. The use of machine learning allows to build an algorithm, which would combine several indicators and learn correlations between them. In contrast to the use of exclusively trading rules defined manually, a model would allow to assess which combination of features correlates with the desired prediction object.

Nevertheless, technical indicators are usually based on the historical observations of the market. The majority of indicators are overlapped, and the predictive power of indicators is specific for each asset, period and forecasting task.

4. Machine Learning Techniques for Stock Market Prediction

4.1 Support Vector Machines

Support Vector Machines (SVM) are supervised learning techniques applicable for classification and regression analysis. For stock market prediction, SVM classification can determine whether a future price change would belong to either the up or down category. SVM Regression can be applied to predict numerical target values such as future prices.

The algorithm attempts to find a decision boundary or prediction function that optimizes the balance between complexity and prediction error. Kernel tricks enable SVM to model nonlinear relationships between inputs and output.

One such kernel function is the radial basis function kernel, which is capable of capturing the nonlinear relationship present in the inputs. Features used by SVM might include historical returns, moving averages, RSI, MACD, volatility and volume-related features.

Some of the strengths of SVM are that it performs well on problems with a moderate sample size and well-chosen features. The kernel trick is another strength of the technique, especially where the linear relation does not hold.

However, SVM performance is dependent on the choice of parameters and feature scaling. The technique may turn out to be computationally intensive for large data sets and may be sensitive to the kernel function chosen.

4.2 Decision Trees

A Decision Tree separates the observations according to a series of decisions on certain variables. In each step, the decision depends on one of the features, and the ultimate decision relies on the path through the tree.

In the case of stock prediction, the decision tree finds correlations between momentum, volatility, volume, and price movements. For instance, it may find out that there is a different relationship between the recent returns and the future price movement depending on volatility. Compared to some other models, a decision tree is quite understandable. It can account for nonlinear correlations and interaction effects without transforming the data.

The major drawback of this algorithm is instability. A slight change in the sample may result in a completely different tree, especially in the case of a complex one. Random Forest and Gradient Boosting were proposed to enhance the predictive power of a single decision tree.

4.3 Random Forest

In the Random Forest approach, several decision trees are combined together to generate a final prediction. The trees are generated based on randomized samples or feature sets, and their predictions are combined to get the output. In case of a numerical prediction, usually the averages of predictions of each tree are considered. In the case of classification, the class predictions or probabilities from each tree are combined.

Random Forest is advantageous for predicting stocks because it helps in dealing with non-linear relations between various financial factors. Moreover, it can be applied to cases where there are several technical indicators without the requirement of specifying all the relations manually.

One more advantage of Random Forest is that it helps in estimating the feature importance. This will help in studying whether the variables like momentum, volatility, or trading volume are helpful in generating the predictions. But, feature importance cannot establish causality. Correlated variables can have similar predictive power, which makes it difficult to give the contribution of each indicator separately.

Kang et al. (2025) reported comparatively strong performance from tree-based approaches in their analysis of the Korean stock market. Their results provide a useful example of the importance of comparative evaluation, but they should not be interpreted as proof that Random Forest or other tree-based models will always outperform alternative methods.

4.4 Artificial Neural Networks

Artificial Neural Network (ANN) is a computational model made up of connected units that are organized in layers. This type of model learns the relationship between inputs and outputs by making iterations of its internal parameters. In

forecasting of stock markets, an ANN can utilize previous prices, returns, technical indicators, and other financial information to predict future values or directions in the market.

The neural network can approximate highly non-linear functions and also take into account several variables. It can be used when the relations between financial information are complicated to describe with mathematical equations. However, an ANN has certain requirements for choosing architecture, learning rate, number of hidden units, amount of training, and regularization. The model becomes ineffective if the complexity of the model exceeds the amount of information. Also, the neural networks can be hard to interpret directly.

For this reason, neural networks should be compared with other models and tested on independent observations.

4.5 Gradient Boosting

Gradient Boosting learns an ensemble of weak prediction models step-by-step. Subsequent models seek to correct the errors committed by the ensemble already learned. Tree-based versions such as XGBoost and LightGBM find their application in structured-data tasks. In finance forecasting, they may use historical returns, technical and fundamental data, as well as other quantitative features.

Gradient Boosting can capture complicated interactions and non-linear dependencies. This flexibility allows using Gradient Boosting in comparison of different feature sets and detection of patterns in financial data. Nevertheless, this technique requires hyperparameters tuning, including learning rate, depth of trees, and number of boosting rounds. Too complicated structure may lead to overfitting, especially for small samples.

It means that the appropriate model specification must be chosen based on training and validation datasets. The test sample should be kept apart.

4.6 Recurrent Neural Networks

Recurrent Neural Networks (RNNs) are intended for working on sequences of observations. In contrast to standard feedforward networks, RNNs keep information from previous time steps when processing the following ones. This feature makes RNN applicable to problems with financial time series, since the order of observations can include some valuable information. A series of daily returns, prices or other technical features can be used to predict a future market situation.

However, simple recurrent neural networks might fail to detect dependencies that occur far back in time. The learning process can also become unstable if the sequence is too long or noisy. That is why there appeared such structures as Long Short-Term Memory networks.

4.7 Long Short-Term Memory Networks

Long Short-Term Memory networks are recurrent neural networks aimed at learning dependencies in a sequence of observations. The gated architecture of the network controls the flow of information stored and updated during processing of the sequence. In case of stock price prediction, the LSTM model may accept a sequence of price history, return observations, volume observations, or technical indicators and predict the next price or return.

Fischer and Krauss (2018) examined LSTM networks for their applicability in making financial market predictions and compared their performance with other techniques. This research highlighted the usefulness of sequence-based neural networks in this field.

However, LSTM is not guaranteed to perform better than other solutions. It will depend on the availability of data, length of sequences, architecture of the model, and its training process. LSTM networks may also need more computing power and tuning. The applicability of LSTM is justified by the case when a sequential observation provides predictive power, which cannot be obtained through simpler models and proper feature engineering.

4.8 Hybrid and ensemble models

Hybrid models use multiple algorithms or multiple information sources in order to leverage each other's strengths. For instance, a system can use statistical prediction and a neural network, or it can blend pricing features with financial

sentiments. The concept of ensemble models is to blend the results of several predictive models. The idea here is to avoid dependency on a single model and possibly increase the stability of the predictions.

Hybrid models can be effective when different relationships are learned by the underlying models. However, the use of multiple algorithms adds complexity and might lead to overfitting.

Thus, the hybrid model needs to be benchmarked against its separate parts using the same data and the same validation process. Additional complexity should only come with a consistent improvement.

5. Data and Features Used in Stock Market Prediction

The performance of machine learning models depends not only on the selected algorithm but also on the quality, relevance, and timing of the input data.

5.1 Historical market data

Typically, historical data from the market includes the opening price, maximum price, minimum price, closing price, adjusted closing price, and volume of trades. The observations can be converted into returns, price ranges, momentum ratios, etc., that serve as the input variables for predictions. Since historical prices are easily available, researchers usually begin their studies on predicting stock markets using historical prices.

However, care should be taken while collecting such data since splits, dividend payments, missing values, symbol changes, etc., might affect the historical comparisons. Sometimes, the adjusted prices may have to be considered when computing returns over a long period of time.

5.2 Technical indicators

Technical indicators summarize aspects of historical price and volume behaviour. Common indicators include:

Simple Moving Average (SMA).

Exponential Moving Average (EMA).

Relative Strength Index (RSI).

Moving Average Convergence Divergence (MACD).

Bollinger Bands.

Average True Range (ATR).

Rolling return and volatility measures.

Moving averages summarize recent price levels. RSI describes the strength of recent upward and downward movements, while MACD represents differences between exponential moving averages. Bollinger Bands and ATR provide information about price variability.

These indicators can be supplied to machine learning models as numerical features. However, many indicators are related to one another because they are derived from the same underlying data. Including more indicators does not necessarily provide more useful information. Feature selection should therefore be based on a clear research objective and evaluated without using future test observations.

5.3 Fundamental information

Basic variables reflect the financial state and performance of a firm. Such basic variables include such values as earnings, sales, profitability, debt, valuation ratios, and cash flow metrics. Basic data can serve as a supplementary source for technical indicators because it reflects the characteristics of the firm's performance that can affect investors' expectations.

But the timeliness of the basic data is crucial. Financial data of a firm becomes available for investors only when it is published. The prediction model has to take into account the date of publication and not the moment when all the data on the reporting period became available.

5.4 Macroeconomic variables

Macroeconomic variables include interest rates, inflation, exchange rates, industrial production, and economic growth. These factors can affect corporate earnings, borrowing costs, investor expectations, and the overall financial environment.

Gu et al. (2020) examined machine learning methods using firm characteristics and macroeconomic variables in empirical asset pricing. Their work illustrates the importance of investigating financial information beyond historical prices alone.

Nevertheless, macroeconomic variables may be released with delays or revised after their initial publication. Their use requires careful alignment with the information actually available at each prediction date.

5.5 News and sentiment analysis

Financial news serves to provide information concerning corporate events, earnings, economic events, and others that could affect market behaviour. Sentiment analysis is an application of natural language processing to assess the positive, negative or neutral sentiment expressed by a given piece of textual data. More sophisticated approaches could include categorization of financial events or representation of news articles in terms of numerical embeddings.

Sentiment features might supplement historical prices especially where the information alters market expectations. Their effectiveness depends on the quality of the text, source credibility, timeliness of the news, and the method of deriving the sentiment score. Another issue to keep in mind is the duplication of news across sources and the fact that sentiment does not necessarily translate into market action.

6. Evaluation of Machine Learning Models

Evaluating a stock market prediction model requires more than calculating a single accuracy value. Different measures capture different aspects of predictive performance.

6.1 Mean Absolute Error

Mean Absolute Error (MAE) measures the average absolute difference between actual and predicted values. A lower MAE indicates smaller average absolute errors. The metric is straightforward to interpret, but it does not distinguish between overestimation and underestimation.

6.2 Root Mean Squared Error

Because errors are squared before averaging, RMSE gives greater weight to large errors. It is useful when unusually large prediction errors are a particular concern. MAE and RMSE should be calculated using the same target and test observations when comparing models.

6.3 Directional accuracy

Directional accuracy reflects the number of times a model can successfully predict the direction of the next price change.

The prediction is said to be accurate if the direction predicted coincides with the actual one. Directional accuracy is appropriate for systems that try to classify whether the next move will be positive or negative. Nevertheless, it does not reflect the size of profits or losses from the correct or incorrect predictions. Thus, even an accurate prediction model can yield very bad results in terms of profitability.

6.4 Precision, recall, and F1-score

Precision is the fraction of true positive movements out of all movements classified by the algorithm as positives in case of classification problem. Recall is the fraction of true positive movements among those classified as positives by the classifier. F1-score incorporates both precision and recall into one metric.

Such metrics may provide additional insight when an algorithm provides trading signals or there is some imbalance between positive and negative outcomes. Area under ROC curve (AUC) is another metric, which can be used to assess the classification algorithm's ability to discriminate positive from negative outcomes at various thresholds.

Choice of metrics is determined by prediction goal and its practical significance.

6.5 Trading performance

When using forecasts to develop an investment strategy, other metrics have to be used besides predictions themselves. They can consist of net return, max drawdown, turnover, portfolio volatility, and Sharpe ratio.

The assessment must define such factors as trading costs, slippage, position size, and execution time. The model which generates many trading signals might look appealing until such costs are taken into account. Predictions generated by the model must be evaluated independently of the trading strategy based on those predictions.

6.6 Chronological validation

Financial forecasting requires careful separation of historical training data and future evaluation data. Randomly dividing time-series observations can allow information from later periods to influence a model evaluated on earlier observations.

A chronological split uses earlier observations for training, a subsequent period for validation, and a later period for final testing. Expanding-window validation can further examine whether performance remains stable across several test periods.

Preprocessing, feature selection, and hyperparameter optimization must be carried out without using information from the final test period. These safeguards reduce the risk of overly optimistic results.

7. Comparative Analysis of Machine Learning Techniques

Different machine learning algorithms offer different advantages, and their suitability depends on the structure of the forecasting problem.

Technique	Main strength	Main limitation
Linear Regression	Simple and interpretable baseline	Limited representation of complex nonlinear relationships
Support Vector Machine	Flexible classification and regression	Sensitive to scaling and hyperparameter selection
Decision Tree	Easy to understand and capable of representing nonlinear rules	Can be unstable and prone to overfitting
Random Forest	Captures feature interactions and reduces dependence on one tree	Feature importance can be difficult to interpret with correlated inputs
Gradient Boosting	Strong modelling flexibility for structured data	Requires careful tuning to control complexity
Artificial Neural Network	Learns complex nonlinear relationships	Requires careful architecture selection and training
LSTM	Processes sequences and temporal dependencies	Computationally demanding and sensitive to data and sequence design
Hybrid Models	Can combine complementary sources of predictive information	Increased complexity and risk of overfitting

This comparison should not be interpreted as a universal ranking. A method's performance depends on the dataset, prediction horizon, feature set, validation strategy, and market conditions.

Kumbure et al. (2022) and Sonkavde et al. (2023) document the diversity of methods used in the literature. Kang et al. (2025) further demonstrate the importance of direct comparison and interpretability in a specific financial-market setting.

The central implication is that algorithm selection should be evidence-based. Researchers should compare alternative methods against appropriate baselines rather than assume that the most complex model will provide the best result.

8. Major Challenges in Stock Market Prediction

8.1 Noisy and uncertain data

There is always some degree of volatility in financial information that does not always correspond to any stable predictive relation. Prices react to both known and unexpected news, hence it is hard to differentiate between randomness and persistency.

A very good model can be obtained for past information, but it may be unable to forecast future observations. Noise implies that regularization and benchmarking are necessary.

8.2 Non-stationarity

The behavior of the market evolves with time. Relationships which are relevant in one point in time may lose their relevance at another due to different economic situations, regulations, liquidity or investor behavior.

A model which does not undergo periodic evaluation may lose its relevance due to changing market conditions. Therefore, researchers need to analyze the model’s performance in several periods of time.

8.3 Overfitting

Overfitting is when a machine learning model learns the pattern from the training dataset instead of learning generalizable relations. The complexity of the model and exhaustive hyperparameter optimization increases the chance of overfitting. In some cases, overfitting happens if researchers conduct multiple experiments and publish the best one.

Chronological validation, limited model complexity, and clear reporting alleviate these issues.

8.4 Data leakage

Data leakage takes place in situations where information not available at the time of prediction becomes part of the modeling process. Such issues can arise due to inappropriate alignment of features, wrong preprocessing techniques, wrong time series partitioning, or future knowledge of financial data. For instance, a technical indicator supposed to predict tomorrow’s return should not rely on tomorrow’s closing price. In addition, a company’s financial statement cannot be treated as available prior to its actual release.

Leakage avoidance is crucial for getting reliable forecasts.

8.5 Limited interpretability

Some machine learning models produce predictions that are difficult to explain directly. This can limit their usefulness for financial analysis and make it difficult to identify unstable or implausible relationships. Explainable artificial intelligence can help identify influential variables and examine how model predictions change with the input features. However, explanatory measures should not be interpreted as proof that a variable causes future stock movements.

8.6 Differences between markets

Markets differ in trading activity, investor composition, regulation, sector structure, and access to information. A model developed for one market may not transfer effectively to another.

For this reason, results from one country or stock index should not be generalized automatically to all financial markets. Market-specific evaluation remains necessary.

8.7 The gap between prediction and profitability

A statistically accurate forecast does not automatically generate a profitable investment strategy. Trading costs, execution delays, liquidity, and risk exposure can reduce or eliminate apparent benefits. Research should therefore distinguish predictive performance from the economic performance of a strategy based on the predictions.

9. Explainable Artificial Intelligence in Stock Market Prediction

Explainable Artificial Intelligence (XAI) seeks to make the behaviour of machine learning systems more understandable. In stock market prediction, interpretability can help researchers investigate which variables influence forecasts and whether their contributions change across different observations.

Permutation feature importance evaluates how model performance changes when a feature's values are shuffled. A substantial decline in performance may indicate that the feature contains useful information for the fitted model.

SHAP, introduced in a unified framework by Lundberg and Lee (2017), estimates contributions of input features to individual predictions. LIME, developed by Ribeiro et al. (2016), approximates a model locally to provide explanations for particular predictions.

Partial dependence plots can summarize the average relationship between a feature and the model's predictions. Individual Conditional Expectation plots show how predictions vary across observations as a feature changes. Accumulated Local Effects plots provide another approach for examining feature relationships, particularly when predictors are correlated (Apley & Zhu, 2020).

These tools can help researchers understand complex models, but their interpretation requires caution. A feature may be important to a fitted model because it is correlated with another variable. Its importance does not establish a causal relationship with stock returns.

Kang et al. (2025) demonstrate how several interpretability techniques can be used alongside model comparison in financial forecasting. Their approach highlights the value of examining both the quality of predictions and the relationships learned by the model.

Future stock market prediction studies can benefit from combining performance measures with interpretability and stability analysis. This would provide a more complete understanding of how models behave and whether their predictions depend on robust relationships.

10. Future Research Directions

10.1 Integration of multiple data sources

Future research can examine whether combining historical prices with fundamental data, macroeconomic indicators, and financial news improves prediction. A comparison between price-only models and models incorporating additional variables would help determine whether the extra information provides meaningful predictive value.

Such studies must ensure that every variable is aligned with its actual availability date.

10.2 Improved hybrid models

Hybrid approaches may combine the strengths of statistical methods, tree-based algorithms, and neural networks. They may also integrate numerical financial data with sentiment or textual information.

However, hybrid models should not be considered superior merely because they combine multiple techniques. Their additional complexity should be justified by improvements on unseen observations.

10.3 Market-regime analysis

Future studies can examine whether machine learning models perform differently during rising markets, falling markets, and periods of high volatility.

This would help determine whether an algorithm's predictive ability is stable or depends strongly on market conditions. Such analysis is particularly important because financial relationships may change over time.

10.4 Explainable and reliable forecasting

Research can further investigate how feature importance and prediction explanations vary across test periods. Stable explanations may provide useful evidence that a model relies on recurring relationships, although they do not establish causation.

Combining interpretability with chronological evaluation could improve the transparency of financial prediction systems.

10.5 Application to different markets

Further research is needed to determine how well predictive methods transfer across countries, indices, and individual stocks.

For example, researchers could compare models across Indian equities, US stocks, and other markets using clearly documented datasets and consistent evaluation procedures. The objective would be to establish which findings are market-specific and which appear more general.

10.6 Realistic financial evaluation

Future studies should evaluate not only forecast errors but also the performance of predefined trading strategies after transaction costs and realistic execution assumptions. Reporting risk measures, drawdowns, turnover, and net returns would help distinguish statistical improvement from practical investment value.

10.7 Reproducibility

Reproducible research requires clear documentation of data sources, preprocessing, features, model configurations, validation procedures, and performance metrics.

Where possible, researchers should provide code and sufficient methodological detail to enable independent replication. This would improve confidence in published findings and facilitate meaningful comparison across studies.

11. Conclusion

Machine learning has become a prominent tool in stock market forecasting due to its flexible techniques used to analyze the past prices, technical indicators, fundamental features, macroeconomic factors, and sentiment of investors in the financial market. Some examples of machine learning algorithms include Support Vector Machines, Random Forest, Gradient Boosting, Artificial Neural Networks, and Long Short Term Memory networks.

The literature presented in this paper shows that machine learning techniques are capable of being applied in financial prediction, but the effectiveness of these algorithms is highly dependent on the dataset characteristics, prediction horizon, chosen features, and evaluation technique. There is no one particular technique which could be called the best for all cases and markets.

This review emphasizes some methodological problems that should be solved in order to properly interpret results obtained while using machine learning algorithms, such as noise in the dataset, changes in market, overfitting, data leakage, limited interpretability, and the difference between statistics and profitable trading.

Interpretable AI will provide researchers with extra methods for exploring what drives the predictions made by the models. If used together with validation and proper benchmarking, it can help researchers better understand how machine learning algorithms behave.

The future research should concentrate on standardized comparison of models, use of multiple sources of financial data, market regimes analysis, interpretability of predictions, and financial evaluation under realistic assumptions. Such studies will enable us to create forecasting systems that are both sophisticated and reliable.

In general, machine learning is a tool for studying financial uncertainty, not an instrument for getting rid of it. The utility of a forecasting system lies in its predictive relationships and their applicability outside of the data that were used to estimate the model.

References

Apley, D. W., & Zhu, J. (2020). Visualizing the effects of predictor variables in black box supervised learning models. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 82(4), 1059–1086.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. The Journal of Finance, 25(2), 383–417.

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.

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. The Review of Financial Studies, 33(5), 2223–2273.

Kang, Y., Ryu, D., & Webb, R. I. (2025). How well do machine learning models in finance work? Financial Innovation, 11, Article 128.

Kumbure, M. M., Lohrmann, C., Luukka, P., & Porras, J. (2022). Machine learning techniques and data for stock market forecasting: A literature review. Expert Systems with Applications, 197, Article 116659.

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems 30 (pp. 4765–4774).

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.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144).

Sonkavde, G., Dharrao, D. S., Bongale, A. M., Deokate, S. T., Doreswamy, D., & Bhat, S. K. (2023). Forecasting stock market prices using machine learning and deep learning models: A systematic review, performance analysis and discussion of implications. International Journal of Financial Studies, 11(3), Article 94.