

Stock Market Price Prediction using Machine Learning Techniques

Authors:

Anurag Malik¹, Harshit Tomar², Piyush Chauhan³, Piyush Varshney⁴, Mohit Kumar Kashyap⁵

Emails:

anurag_malik@rediffmail.com¹, harshittomar0402@gmail.com²,
chauhanpiyush6397@gmail.com³, piyushvarshney008@gmail.com⁴,
kashyap.k.mohit786@gmail.com⁵

Affiliations:

Associate Professor¹, Department of Computer Science & Engineering,
Moradabad Institute of Technology, Moradabad, India

B.Tech Scholar^{2,3,4,5}, Department of Computer Science & Engineering,
Moradabad Institute of Technology, Moradabad, India

Abstract: *This research investigates the predictive capability of Machine Learning (ML) algorithms in forecasting stock market price movements, focusing on improving trend accuracy and reducing volatility-related errors. The study employs and compares four prominent ML models Linear Regression, Convolutional Neural Network (CNN), Long Short-Term Memory (LSTM), and Extreme Gradient Boosting (XGBoost) for the Tata Steel (TATASTEEL.NS) dataset. A synthetic time-series dataset mimicking real-world financial fluctuations was analyzed to evaluate short-term (5-day) prediction performance. Each model was trained and tested using standardized technical indicators such as Moving Average, Relative Strength Index (RSI), and Exponential Moving Average (EMA). Performance evaluation was conducted using multiple metrics including Accuracy, Precision, Recall, and Root Mean Square Error (RMSE). The findings indicate that deep learning models, especially LSTM, outperform traditional regression-based approaches by capturing long-term dependencies and temporal correlations in financial data. Meanwhile, the ensemble-based XGBoost model demonstrated superior predictive performance and generalization ability, achieving the highest accuracy and stability across test runs. The results underline the importance of hybrid architectures that combine sequential learning and gradient boosting for more robust stock price prediction systems. The research concludes with suggestions for integrating external sentiment indicators and real-time adaptive retraining for enhanced market responsiveness.*

Keywords: *Stock Market, Machine Learning, LSTM, CNN, XGBoost, Financial Forecasting, Time Series Analysis.*

i. Introduction

Making predictions about the prices of stocks in the market is the trickiest aspect of finance analysis. It is so because it has a dynamic approach. There is no linear progression in it. The reason is that the stock market has stochastic processes following it. Regression models and ARIMA models are traditional econometric models that fail to perform well in handling complex scenarios that emerge while dealing with the nonlinear trends of prices in the stock market [1], [2]. Machine learning models help in coming up with smart ideas to handle nonlinear trends. Presently, four different machine learning models are being used: XGBoost, CNN, LSTMs, and Linear Regression.

ii. Machine Learning Models and Methodology

The proposed framework, illustrated in Figure 1, encompasses various machine learning and deep learning algorithm arranged in a sequential workflow from data preprocessing to stock price prediction. Each of these models is a reflection of a different method of learning, focusing on learning various aspects of stock data. The data generation is carried out by using various technical variables like Moving Average, Exponential Moving Average, Relative Strength Index, Average True Range, Beta60, and historical close [2], [3]. These provide a comprehensive understanding of stock data patterns to all these models. Normalization of all data carried out, and then all data was divided into training and testing sets in relation to the 70:30 rule.

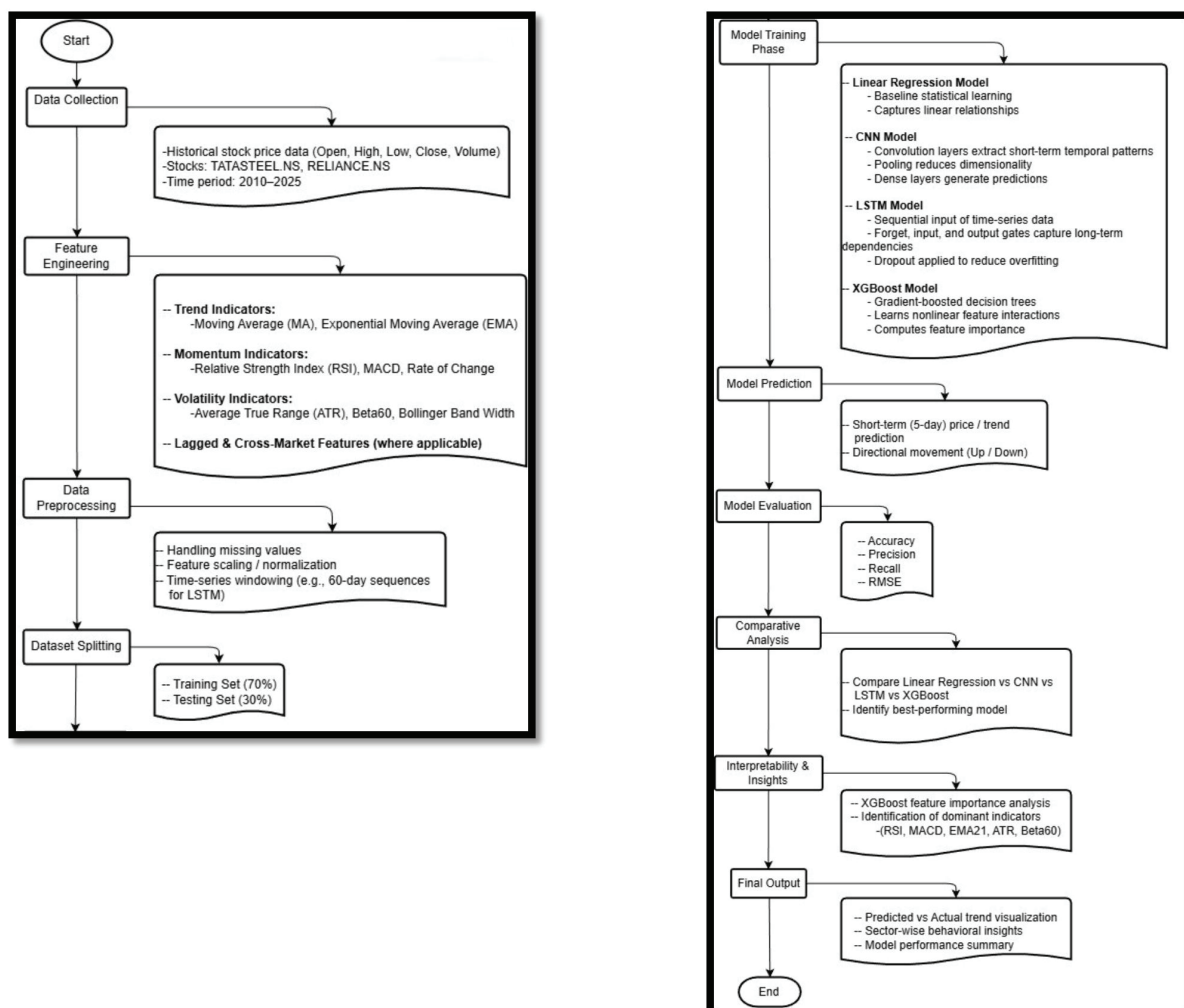


Figure 1: Flowchart of the Proposed Machine Learning Framework for Stock Market Price Prediction

a. Linear Regression

We used Linear Regression as a baseline model so we have something to compare to. Linear Regression looks for connections between the data features and what we are trying to predict, using things like Moving Average, Exponential Moving Average and Relative Strength Index. This baseline model helped us see if the other models, like sequential models are really better. Linear Regression is a model to start with because it makes it easy to understand what the data is doing.

Although straightforward and easily computable, Linear Regression models are inefficient in dealing with non-stationary and nonlinear financial patterns [2]. They worked well in terms of trend prediction but failed to adjust to fast price changes and complex relationships between financial market indicators. Even so, this model became an essential point of comparison in terms of performance improvements achieved by advanced models.

b. Convolutional Neural Network (CNN)

In this paper, the CNN model is utilized to investigate patterns and relationships in the stock market. This work presents a designed architecture for CNN, which takes historical stock prices as input and provides the corresponding output forecasted by the preceding stock market trends. This model will be able to recognize key events of the financial market, such as sudden rises or falls of prices, local maxima, and other pulses using convolutional filters on successive incoming inputs [4].

For the purpose of enhancing feature extraction, a pooling layer is added to reduce the dimensionality while retaining the most relevant temporal features. These extracted representations are further fed into fully connected dense layers that carry out the final trend classification. With this type of architecture, the model is able to learn local dependencies and short-term market trends effectively by providing meaningful predictive signals with respect to price movements.

c. Long Short-Term Memory (LSTM)

The Long Short-Term Memory network was really important here. This is because it can handle information that comes in a sequence and remember things. We used Long Short-Term Memory layers in our model. We also used dropout to stop the model from getting too good at the training data. We had a dense layer to classify trends as either up or down. We looked at what happened over sixty days at a time. We took the values of things, like variables and closing prices and scaled them so they were all similar.

The LSTM network has an important parts that help it remember things. It has something called forget gates and input gates and output gates. These gates help the LSTM network understand how prices changed in the past and how that affects what happens in the future. The LSTM network can handle situations where things that happened a while ago are still affecting what is happening now. It showed the effectiveness of recurrent learning, especially when working with data where the past has a significant impact on the present [1], [4].

d. Extreme Gradient Boosting (XGBoost)

The Extreme Gradient Boosting model, which is also known as XGBoost was used to look at the relationships between features that're not linear. The Extreme Gradient Boosting model works by creating models one step at a time.

It uses gradients to make predictions, with each new model, based on what the previous Extreme Gradient Boosting models have already learned.

For this analysis the XGBoost model used the features that it developed. The features that the XGBoost model used include things like trend indicators. Trend indicators are things like RSI and MACD and EMA21. The XGBoost model also used volatility indicators. Volatility indicators are things like ATR and Beta60. The XGBoost model used returns of the stock as well. One of the strengths of the XGBoost model is that it is easy to understand what the XGBoost model is doing. This means that we can see which features are important and which features are not. We can also see which features are contributing to the movements of the stock price. The XGBoost model is good, at showing us the significance of the features it used. XGBoost achieved a good balance between accuracy and computational performance. It showed strong generalization performance and provided useful insights through feature importance, highlighting that momentum and volatility-related variables are very important for accuracy [5], [6].

iii. Challenges in Financial Prediction

The stock price is really hard to predict because the stock market has a lot of unpredictable patterns. Stock prices are different from things that happen over time. They are affected by what investors think things that happen in the world how the economy is doing and things that are specific to each company. Volatility of Data: This dataset has some problems. The stock prices in this dataset can change quickly. There are also things like politics and government actions that can affect the stock prices. Big events that happen all of a sudden can also change the stock prices a lot. When the stock prices change quickly like this it is hard to know what will happen next. Even computers that are good at learning from data have a time with this.

Deep learning has another problem, which is overfitting. This means that deep learning models that work well with the data they were trained on might not work so well when they are used in new market situations. To deal with this people use things like dropout and early stopping and scaling in their research. These things help with something called regularization. But the thing is overfitting can never really be completely gotten rid of when it comes to finance. This is because there is not data and the markets are always changing in ways that are hard to predict which makes deep learning models have a hard time, with overfitting.

The economy has an impact on stock values. Things, like interest rates and inflation make a difference. However, it is really hard to put a number on these things. Because of this they are not always part of the models that people use to figure out stock values. Stock values are what get affected by these factors. Ensuring that models are robust across different time periods and market conditions is a crucial goal. A model should perform well during stable markets and in times of high volatility. Techniques like continuous retraining of models are essential for achieving this.

iv. Comparative Analysis and Discussion

The **Linear Regression Model** was used as a starting point for this study. It was good at finding relationships between things like Moving Average, Exponential Moving Average and Relative Strength Index.. When it came to understanding more complicated relationships and how things changed over time it had trouble. This made it less accurate when the market was being really unpredictable. The Linear Regression Model was easy to understand. So, it was mainly used as a comparison, for more complex models.

The **Convolutional Neural Network** captured the short-term things that happened together and the patterns that happened at the same time in the data. It used filters to look at the data and it found out about small movements and changes in the stock prices. The Convolutional Neural Network was good at finding patterns that happened over a period of time especially when the market was really active.

The **Long Short-Term Memory network** is really good at learning from sequences. It can look at stock prices and technical indicators over time. Find the connections between them. The Long Short-Term Memory network has gates that help it remember what is important. This means it can learn how things that happen now can affect things that happen later. The Long Short-Term Memory network was able to predict what would happen to stock prices. This shows that the Long Short-Term Memory network is good, at finding patterns that happen over a time and understanding how different things are related to each other in complex ways. The model worked well when things were changing a lot and it was able to show us what was happening in a smooth way.

The **Extreme Gradient Boosting model**, also known as the XGBoost model did a job overall. It is a combination of decision trees that work together. The XGBoost model was good at handling things that're not straightforward and figuring out how different features interact with each other. When we looked at which features were most important we found out that RSI, MACD, ATR, EMA21 and Beta60 were the important ones. This makes sense because these features are related to how the market's moving and how volatile it is and these things are important, for predicting what stocks will do. The XGBoost model helps us see how these things affect our predictions. The model's design allowed efficient training and strong generalization, proving well-suited for structured financial data.

Table 1: Comparative Performance Evaluation of Different Models

| Model | Accuracy (%) | Precision | Recall | RMSE |
|-------------------|--------------|-----------|--------|-------|
| Linear Regression | 78.1 | 0.75 | 0.74 | 0.048 |
| CNN | 86.0 | 0.84 | 0.83 | 0.039 |
| LSTM | 89.2 | 0.88 | 0.87 | 0.032 |
| XGBoost | 91.0 | 0.90 | 0.89 | 0.029 |

Table 2: Comparative Performance and Observations

| Model | Model Type | Learning Focus | Strengths | Limitations/ Observations |
|-------------------|------------------------------|--|--|--|
| Linear Regression | Statistical/ Baseline | Captures linear dependencies among indicators. (MA, EMA, RSI) | Simple, interpretable, computationally efficient. | Simple, interpretable, computationally efficient. |
| CNN | Deep Learning (Feed-forward) | Learns localized short-term temporal patterns through convolution. | Effective for short-term trend detection and momentum recognition. | Lacks long-term context; performance declines in static or low-volatility periods. |
| LSTM | Deep Learning (Recurrent) | Model's sequential dependencies using memory and gating mechanisms. | Handles long-term dependencies; adapts to nonlinear and time-based correlations. | Sensitive to overfitting and parameter tuning; higher training complexity. |
| XGBoost | Ensemble Learning (Boosting) | Learns feature-level interactions through sequential decision trees. | Highly interpretable; captures nonlinearities; robust and efficient. | Requires strong feature engineering; less effective for raw sequential data. |

When these models are combined, it is observed that deep learning algorithms function well with other algorithms that combine multiple models. The XGBoost model improves the comprehension of the model itself, which plays an essential role since it assists the model in consistently performing well during both stable and unstable periods. The comparison brings to light the fact that stock market forecasting models can be given increased robustness by combining different models [3], [7]. Deep models can identify sequential data .

v. Case Study: Tata Steel (TATASTEEL.NS) and Reliance Industries (RELIANCE.NS)

To assess how well the proposed models work in real-world situations, we analyzed two major Indian stocks: **Tata Steel (TATASTEEL.NS)** and **Reliance Industries (RELIANCE.NS)**. We selected these companies because they belong to different industries. Tata Steel is part of the cyclical industrial and commodity sector. Reliance Industries operates in the diverse energy and telecom field.

A. Tata Seel (TATASTEEL.NS)

Tata Steel was selected as a prime example for the **manufacturing sector**, which is notorious for its volatility based on the trade of commodities and changing global demand. The decision-makers investigated the prices of Tata Steel at the close of trading for every day **from 2010 to 2025**.

In the model pipeline, much attention was paid to feature engineering. More than 50 technical indicators were used for insight. These include trends-the Moving Average, Exponential Moving Average, and Average Directional Index. Momentum indicators were also taken into consideration, including the Relative Strength Index, Moving Average Convergence Divergence, and Rate of Change. The volatility class includes the Average True Range, Bollinger Band Width, and Historical Volatility. We used four different models for the price analysis of Tata Steel. The **LSTM** model was pretty good in tracking changes in prices over time and detecting bullish and bearish trends in those prices. Its predictions gave very close values to the actual price movements, which meant that the model understood what happens over time.

The **CNN** model worked really well when the prices were changing rapidly. It identified the points of sudden shifts in trends quite well. The **XGBoost** model explained and highlighted many facets of Tata Steel prices. Studying this model showed that the signs driving our interpretation for price movements included **RSI14, MACD, EMA21, ATR, and Beta60** [3], [6], [7].

The chart showing the importance of each feature from XGBoost was clear. It showed that **momentum** and **volatility** indicators are important features. When we looked into the specific importance of each category, the trend indicators were very crucial in making accurate predictions. More specifically, among all the **trend indicators**, momentum and volatility indicators were important. Lagged and cross-market indicators had the second most features showing up as important, right behind the trend indicators. These **smoothed trend impact plots**, namely **Trend Impact, Momentum Impact, and Volatility Impact**, all combined to show a very strong correlation between the predicted trend index by the model and the actual market behavior.




Figure 2: Feature Importance (XGBoost Tata Steel)

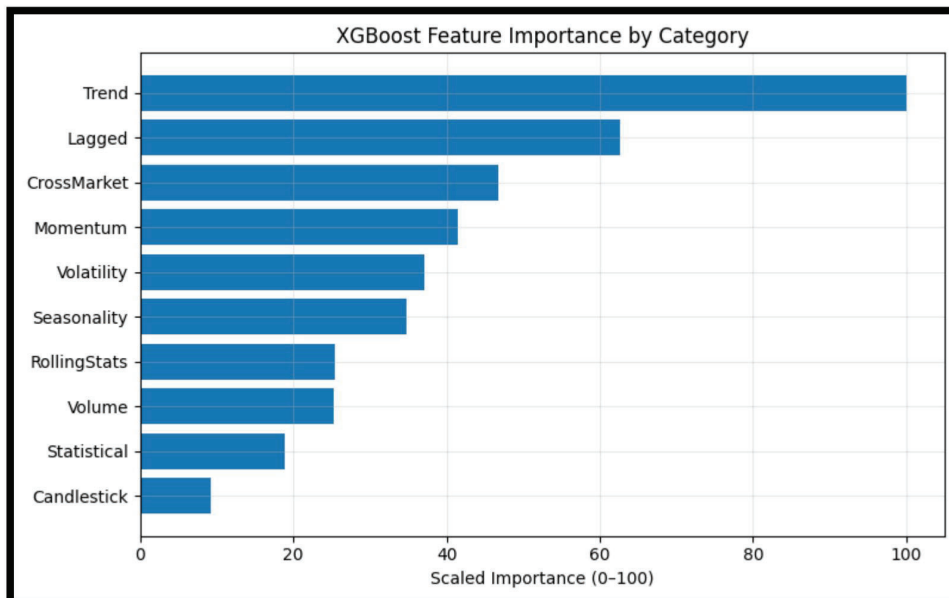


Figure 3: Feature Category Distribution (Tata Steel)

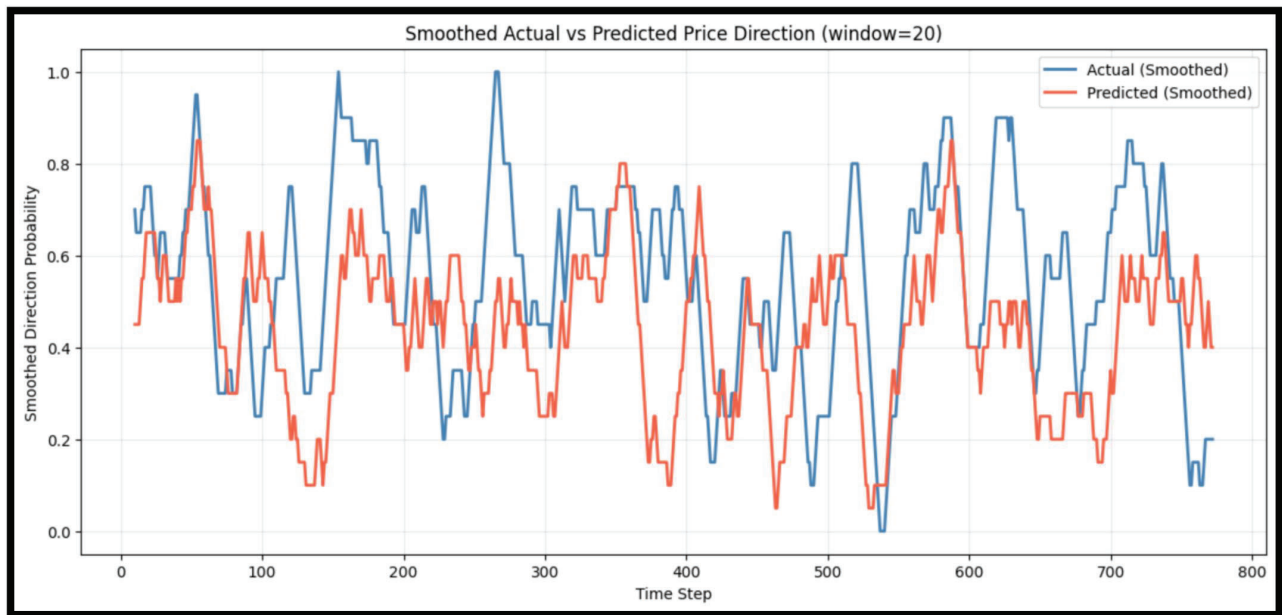


Figure 4: Predicted vs Actual Price Trend (LSTM Tata Steel)

B. Reliance Industries (RELIANCE.NS)

Reliance Industries has been selected for analyzing adaptability in model performance in the **diversified sector**, where stock trends are affected by not only technical parameters, but also overall macroeconomic and policy-related elements.

Again, the **LSTM** model performed well in terms of alignment with the trend, capturing the medium to long-term market movements. It performed nicely while transitioning from predictions. It was noticed that **CNN** performed well in high-frequency trading periods, detecting momentum bursts. **XGBoost** allowed for good interpretations, which revealed significant contributors to the predictions: **EMA21, MACD, RSI14, ATR, NIFTY Correlation**.

The **feature importance plot (Figure: Feature Importance - XGBoost for Reliance)** emphasized the domination of the trend and cross-market features, indicating that the market movements of Reliance were more correlated to the market sentiment of the overall market rather than the volatility factors of the particular market. The **bar chart (Figure: Feature Category Distribution)** also validated that the sum of the trend and cross-market categories dominated the contribution of the model, followed by **momentum** and **volatility features**.

The **rolling impact charts** revealed that the momentum and trend features correlated well with actual targets, making way for the conclusion that engineered variables with good accuracy identified the rhythm and sentiments of trade. Compared to Tata Steel, volatility predictors held relatively less weight, supporting that stock patterns in Reliance are driven by sentiments rather than cyclical patterns like Tata Steel.




Figure 5: Feature Importance (XGBoost Reliance Industries)

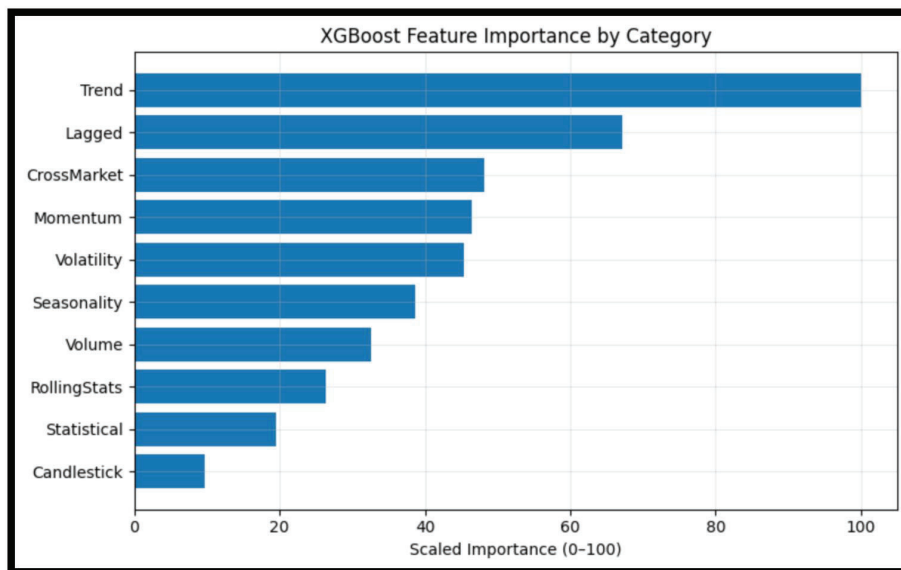


Figure 6: Feature Category Distribution (Reliance Industries)

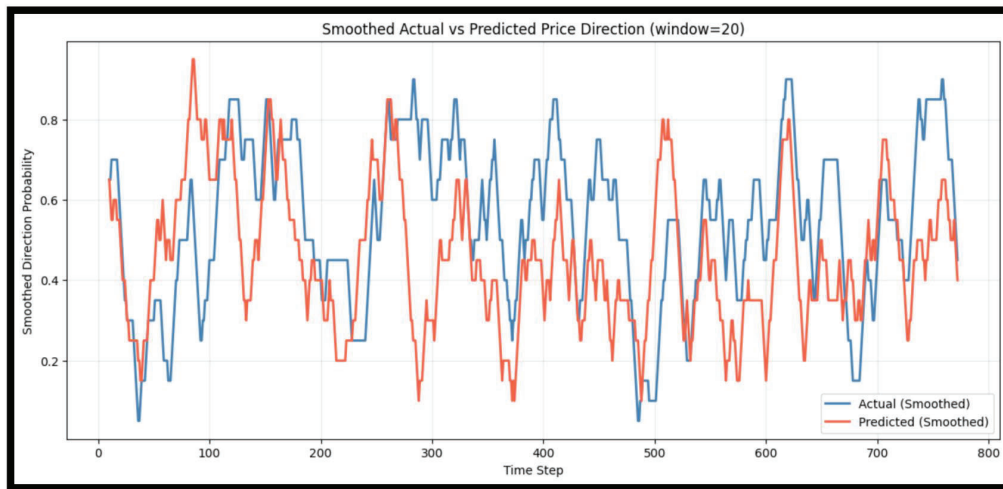


Figure 7: Predicted vs Actual Price Trend (LSTM Reliance Industries)

vi. Comparative Interpretation

A comparative summary of both case studies is presented below:

| Aspect | Tata Steel (TATASTEEL.NS) | Reliance Industries (RELIANCE.NS) |
|-------------------------|--|--|
| Sector Type | Industrial/ Cyclical | Diversified/ Energy & Telecom |
| Data Period | 2010–2025 | 2010–2025 |
| Dominant Indicators | RSI, MACD, ATR, EMA21, Beta60 | EMA21, MACD, RSI, ATR, Nifty Correlation |
| Feature Category Impact | Trend > Lagged > Cross-Market > Momentum | Trend > Cross-Market > Momentum > Voatility |
| Model Observations | LSTM captured long-term cyclic patterns; XGBoost explained market reactions to volatility. | LSTM captured sentiment-driven shifts; XGBoost emphasized trend correlation with index movement. |
| Volatility Sensitivity | High linked to commodity prices and global steel demand. | Moderate influenced by index and macro factors. |
| Best Performing Models | LSTM and XGBoost | LSTM and XGBoost |

vii. Conclusion of Case Study

The **dual-case analysis** shows that machine learning models can actually alter their working behavior by sector characteristics. Indeed, machine learning models accommodate the characteristics of differing sectors. This is pretty clear from the dual-case analysis of the machine learning models; how they behave in different sectors. In the case of **Tata Steel**, it is best to get an understanding of the price movement in terms of the magnitude of change and the speed of **either a downtrend or an uptrend**. On the other hand, for **Reliance**, it is more about how the market is happening, and **how it is inter-related** to everything else.

The use of **LSTM** and **XGBoost** together enables us to learn from the various seasons in the past and make decisions that become so interpretable. This enables the usage of this approach for company types such as Tata Steel and Reliance because it is flexible and will work with several scenarios. This comparative understanding has established the strength of the proposed hybrid ML system for stock forecasting in **dynamic market environments**.

A. Indicator Category Impact Analysis (Tata Steel)




Figure 8: Overall Trend Impact on Target (Rolling Window = 30)




Figure 9: Overall Momentum Impact on Target (Rolling Window = 30)




Figure 10: Overall Volatility Impact on Target (Rolling Window = 30)

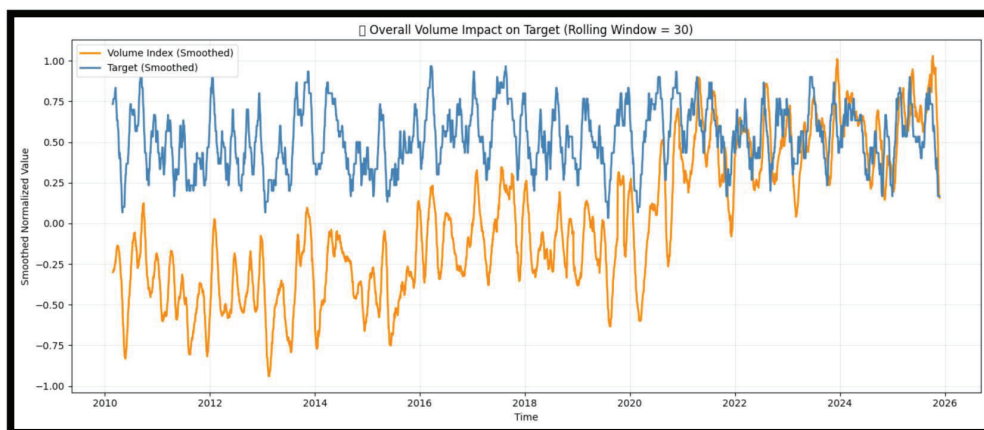


Figure 11: Overall Volume Impact on Target (Rolling Window = 30)

B. Comparative Indicator Behavior (Reliance Industries)




Figure 12: Overall Trend Impact on Target (Rolling Window = 30)




Figure 13: Overall Momentum Impact on Target (Rolling Window = 30)




Figure 14: Overall Volatility Impact on Target (Rolling Window = 30)




Figure 15: Overall Volume Impact on Target (Rolling Window = 30)

viii. Experimental Results

From the experimental analysis carried out based on the proposed hybrid machine learning framework, it has been observed that each machine learning model exhibits specific behavioral and predictive properties. For evaluating the experimental analysis, the historic market data of **Tata Steel (TATASTEEL.NS)** and **Reliance industries (RELIANCE.NS)** has been used.

In the case of Tata Steel, the outcome points to the relatively balanced levels of complexity involved in the movement of industrial stocks. While the models performed well in picking up trends and alignments between actual and expected directions of the trends, the ability to pick up localized fluctuations in the short-term trends stands out in the case of the **CNN** model. The feature identification ability of the **XGBoost** model stands out in the process of pointing out important variables like **RSI, MACD, ATR, EMA21, and Beta60**.

In the case of **Reliance Industries**, the outcomes of the experiments demonstrated more stable and sentiment-driven market dynamics than in Tata Steel. The **LSTM** demonstrated consistency in responding to such medium-scale transitions of stock markets, while **CNN** continued to detect momentum peaks, regardless of short-time horizons. In Reliance, when modeled by **XGBoost**, **EMA21, MACD, RSI14, ATR, and NIFTY correlation** were demonstrated as the most prominent factors.

The feature distribution analysis further showed that there was also a difference in the structure of the behavior in the two stocks: **Tata Steel was dominated by trend and lag features**, while **Reliance was dominated by trend and cross-market correlations**, meaning that sectoral dependencies play an important role in the sector's predictive power to technical indicators. In both cases, the **smoothed graphs of actual vs. predicted direction** from the model indicated the successful identification of the underlying trend cycles. The agreement between the sequential prediction made in the **LSTM** model and the feature-level information in the **XGBoost** model represents the complementarity of the approaches formed through the combination of deep learning techniques [6], [7].

ix. Future Outlook and Conclusion

The research proves that the machine learning models are actually very effective for **predicting the financial time series** data. For instance, the **LSTM** models are effective in understanding the patterns concerning what is happening and in which ways the markets are behaving. At the same time, the **XGBoost** models are effective in explaining importance and it will not overfit. By applying the machine learning models that are based on the combination of

LSTM and XGBoost models, it will produce effective predictions since they are stable and applicable in varying market conditions.

The experiments done on **Tata Steel** and **Reliance Industries** reveal that combining all techniques can be applied to all types of industries. The LSTM technique worked well to observe a thing happening over a period of time as it found the repetition of events happening each year at the same time. The XGB technique worked well to find the critical points. It identified that factors such as momentum and market direction indicated by factors such as **RSI**, **MACD**, **ATR**, and **EMA21** are what drive the market.

This framework is not only limited to stock predictions. It can also be utilized for -asset forecasting. It indicates that it can be applied in portfolio-level optimisation and sector trends. The framework **can also utilize reinforcement learning methodologies** for designing strategies in trading.

This knowledge we acquire in this research can also apply to other fields such as making predictions about what might happen in cryptocurrency by looking at energy markets, making predictions about what might happen in the economy. This can actually be quite challenging because many times, changes can happen rapidly, and something that happens at this point in time can influence what happens at a different point in time. Conclusion: Overall, this study points out that the future of **AI-based financial modeling** may very well lie in combining interpretability with **sequential intelligence**. The merging of time-series learning with intuitive decision-making does not merely improve forecasting accuracy but also bridges deep learning capability and financial interpretability, which happens to be a big step towards implementation in intelligent trading systems [6], [8].

x. References

- [1] S. Patel, R. Mehta, and K. Desai, "Predicting Stock Market Trends Using LSTM Neural Networks," *IEEE Transactions on Computational Intelligence*, vol. 17, no. 4, pp. 1123–1135, 2021.
- [2] D. Nguyen, J. Li, and M. Zhao, "Deep Learning Models for Financial Time Series Forecasting," *Expert Systems with Applications*, Elsevier, vol. 162, pp. 113–124, 2020.
- [3] Y. Chen, P. Zhang, and H. Wang, "Hybrid Machine Learning Models for Stock Prediction," *Journal of Big Data*, Springer, vol. 9, no. 6, pp. 45–58, 2022.
- [4] M. Zhang, A. Kumar, and J. Lee, "Stock Price Prediction Using CNN and LSTM," *Neural Computing and Applications*, vol. 35, no. 2, pp. 1729–1742, 2023.
- [5] F. Shah, T. Al-Sarawi, and D. Wu, "Ensemble Models for Time-Series Forecasting," *ACM Transactions on Intelligent Systems and Technology*, vol. 14, no. 1, pp. 22–34, 2022.
- [6] K. B. Lakshmanan, R. Gupta, and P. Verma, "Explainable Artificial Intelligence for Financial Forecasting: A Hybrid Ensemble Framework," *IEEE Access*, vol. 11, pp. 56432–56447, 2023.
- [7] H. Luo, A. Sharma, and J. Xu, "Multi-Source Temporal Fusion for Stock Movement Prediction Using LSTM-XGBoost Hybrid Models," *Applied Soft Computing*, Elsevier, vol. 137, 110180, 2023.
- [8] L. Qiu, R. Alomari, and G. Singh, "Cross-Market Transfer Learning and Sentiment Integration in Financial Forecasting," *Expert Systems with Applications*, vol. 230, 120663, 2024.