

Stock Price Predictions using Explainable AI

A. R. Pradnyavant¹, Anushka Prakash Kuchakar², Sanika Dattatray Jadhav³, Rasita Ramesh Chavan⁴ and Sanika Shrikant Maske⁵

¹⁻⁵Computer Science and Engineering, ADCET, Ashta, India

Emails: arp.cse@adcet.in, kuchakaranushka@gmail.com, sanikajadhav02004@gmail.com, chavanrasita@gmail.com, sanikamaske2607@gmail.com

Abstract— Predicting stock prices is not an easy job because the market keeps changing a lot and many outside things affect it. In this project, we are trying to build an Explainable AI (XAI) model that not only predicts stock prices but also explains how it came to those predictions. For this, we are using old stock data, market patterns, and some outside economic factors to train a special deep learning model that mixes LSTM (Long Short-Term Memory) and TFT (Temporal Fusion Transformer). Along with that, we are also using LIME (Local Interpretable Model-agnostic Explanations) to make the model's thinking more clear by showing which factors are affecting the predictions. The results show that our model is giving accurate predictions and is also transparent in how it works. This explainable part helps investors and financial people understand how things like market trends, trading activity, and world events affect stock prices. We plan to make this model work inside a website or a desktop app, where users can see both the predictions and the reasons in a simple way. By adding explainability, we are making the system more trustworthy and useful for investors who want clear, data-based guidance.

Index Terms— Stock Price Prediction, Explainable AI, LSTM, Temporal Fusion Transformer, Market Trends, Financial Forecasting, Feature Importance.

I. INTRODUCTION

Artificial Intelligence (AI) is growing really fast and it's changing how many industries work — finance is no exception. These days, in online stock trading, AI models are being used a lot to handle large amounts of data and predict how stock prices might move. But the problem is, most traditional AI models work like a “black box” — they just give you results without telling you how they got them. This makes it hard for users to trust the system or understand the logic behind it.

That's where Explainable AI (XAI) comes in — it helps us understand why the AI is making a certain prediction. This kind of transparency is very important in stock trading, because better understanding means better decisions, less risk, and possibly better returns.

In our project, we are building an XAI-based model that not only predicts stock prices accurately but also explains the logic behind its predictions in real-time. With these explainable methods, we want to give investors more confidence and reduce the need to blindly trust what the AI says. Our system will use past data as well as live market trends and news to make predictions, and show both the results and the reasons in a simple web or desktop app. Basically, the goal is to make advanced AI tools easier to understand and more useful for regular investors — helping them make smarter, well-informed decisions in the stock market.

II. ABBREVIATIONS AND ACRONYMS

XAI – Explainable Artificial Intelligence

This basically means AI systems that not only give you results but also tell you why they made that prediction. It helps make things more transparent and builds trust — especially useful when you're dealing with something sensitive like stock market predictions

AI – Artificial Intelligence

This is when machines are made smart enough to act like humans — learning, thinking, and making decisions. In finance, it helps a lot with things like forecasting stock prices.

ML – Machine Learning

This is a part of AI where machines learn from data using algorithms. It's used a lot for predicting stock prices by learning patterns from past data.

LSTM – Long Short-Term Memory

This is a special kind of neural network that's good at handling time-based data. It can remember trends from the past and use that to forecast stock price movements.

TFT – Temporal Fusion Transformer

TFT is a powerful deep learning model made specially for time-series forecasting. It uses advanced techniques like attention mechanisms and variable selection to learn both short-term and long-term patterns in data — perfect for something like predicting stocks over time.

TFT and LSTM Hybrid Model – In this project, we are combining both TFT and LSTM models to get the best of both worlds — LSTM's memory and TFT's powerful attention-based forecasting — for more accurate stock price predictions.

LIME – Local Interpretable Model-Agnostic Explanations

LIME helps explain what the AI model is doing. It works by building a simpler model around each prediction to show which features (like news, volume, trend) influenced the result the most.

SHAP – SHapley Additive Explanations

This one comes from game theory. It gives each input feature a "score" to show how much it contributed to a prediction — helping you understand which factors are most important.

API – Application Programming Interface

Think of this like a bridge that lets two different software systems talk to each other. It helps connect the stock prediction model with websites or mobile apps.

RMSE – Root Mean Squared Error

This is a number we use to check how far off our model's predictions are from the real stock prices. Lower is better — it means our predictions are more accurate.

MAE – Mean Absolute Error

Another accuracy metric. This one shows how much the model is wrong on average. Again, lower means better predictions.

R^2 – R-squared

This tells us how well our model fits the data. If the value is close to 1, it means the model is doing a great job in explaining the variations in stock prices.

GUI – Graphical User Interface

This is the front-end or visual part of our system — the buttons, charts, and input fields that users interact with. It makes the tool easy to use for everyone.

JSON – JavaScript Object Notation

This is a lightweight format we use to send and receive data, especially in web apps. It's easy to read and works well with most programming languages.

F1-Score

This metric is used to check how good the model is, especially when the data is imbalanced. It balances precision and recall and gives one final score to tell how well the model is performing.

III LITERATURE SURVEY/ RELATED WORK

A. The Importance of Explainable AI in Financial Forecasting

Saranya and Subhashini [1] conducted a systematic review of Explainable Artificial Intelligence (XAI) models and applications, highlighting their increasing relevance in finance. Their study distinguishes between knowledge-driven and data-driven XAI approaches and emphasizes their growing adoption in sectors such as

healthcare and finance. Our work builds on this by implementing XAI techniques specifically for stock price prediction.

B. Hybrid Approaches for Explainable AI

De et al. [2] proposed a hybrid XAI approach that combines local and global explanation methods to enhance the interpretability of deep learning models. Their study focuses on human-interpretable explanations, a principle we integrate into our stock prediction model by utilizing techniques like SHAP and LIME to enhance transparency.

C. Deep Learning in Stock Market Prediction

Olorunnimbe and Viktor [3] provided a comprehensive survey on deep learning applications in stock market forecasting. Their research explores time-series forecasting, sentiment analysis, and portfolio management, but highlights the challenges of model interpretability. Our approach addresses this gap by integrating XAI methods to improve model transparency.

D. Evaluating Fund Performance with Explainable AI

Kovvuri et al. [4] examined the role of XAI in financial decision-making, particularly in investment fund evaluations. Their work emphasizes the importance of transparent decision-making in finance, which aligns with our objective of providing interpretable stock price predictions to assist investors in making informed choices.

E. Prediction-Explanation Networks for Stock Forecasting

Li et al. [5] introduced a Prediction-Explanation Network (PEN) for stock price forecasting, enhancing the explainability of AI-driven financial models. Their study informs our methodology, as we incorporate similar principles to improve the interpretability of our LSTM and Temporal Fusion Transformer (TFT)-based models.

F. Deep Learning and Explainability for Market Trends

Muhammad et al. [6] explored the combination of deep learning and XAI for stock market trend prediction, demonstrating the effectiveness of interpretable models in financial applications. We extend their findings by implementing a hybrid AI approach for stock price forecasting with real-time explainability.

G. Integrating AI and XAI for Stock Price Prediction

Marey et al. [7] conducted an empirical study integrating deep learning and XAI for stock price prediction. Their work underscores the importance of explainable models for financial decision-making, a principle we adopt in our system by providing detailed, human-interpretable explanations alongside stock price forecasts.

IV. METHODOLOGY

In the project “Stock Price Predictions Using Explainable AI,” we’ve made stock market prediction easier for everyone to understand by clearly explaining how each result is generated. This is done through the following steps:

A. Motivation for Hybrid Model

To get better prediction results, we combined two powerful models – LSTM and TFT – because each has its own strengths. LSTM Strengths:

- It’s very good at understanding data that comes in a sequence, like daily stock prices.
- It can remember and learn from long-term patterns in time-series data.
- It works well with raw historical stock data and learns from it

TFT Strengths:

- It can predict multiple future steps at once (multi-horizon forecasting).
- Uses attention mechanisms to identify important time steps and features.
- It uses an “attention mechanism” to automatically figure out which time periods and features (like news or volume) are important for making better predictions.

Hybrid Approach Goal

- Combine LSTM’s strong temporal pattern learning with TFT’s explainability and attention mechanisms to achieve better accuracy and interpretability.

B. Model Architecture

The proposed hybrid model consists of below major components:

- **Input Data Representation.** The input dataset consists of historical stock price data along with relevant technical indicators and external covariates. The input features are categorized as follows:
 - Past Observations: Open, High, Low, Close, Volume, Moving Averages (100 days, 200 days).
 - Future Known Inputs: Economic indicators, trading holidays, corporate earnings reports.
 - Static Covariates: Sector classification, stock ticker information.

Each input is normalized and structured into lookback sequences for time-series forecasting.

- **LSTM-Based Temporal Feature Extractor.** The LSTM component is responsible for capturing sequential dependencies within stock price time-series. It consists of:
 - Multiple LSTM layers to model temporal correlations.
 - Dropout layers to prevent overfitting.
 - Dense layers to transform LSTM embeddings into a meaningful feature representation.

Mathematically, the LSTM cell at time step t is defined as:

$$h_t = f(W_h h_{t-1} + W_x x_t + b) \quad (1)$$

where h_t is the hidden state, x_t is the input at time t , and W_h, W_x, b are the weight matrices and bias term.

- **TFT-Based Attention and Feature Selection Module.** The Temporal Fusion Transformer (TFT) module dynamically selects relevant features and captures long-term dependencies. It consists of:
 - Variable Selection Network (VSN): Learns the importance of different input features at each timestep.
 - Multi-Head Self-Attention: Captures long-range dependencies and contextual interactions.
 - Gated Residual Network (GRN): Controls information flow using gating mechanisms.

The self-attention mechanism in TFT is given by:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (2)$$

where Q, K, V are the query, key, and value matrices, and d_k is the dimensionality of keys.

- **Fusion and Output Layer.** The results from both the LSTM and TFT models are combined together to create one single feature set. This combined information then goes through the following steps:
 - Fully Connected Dense Layer: This layer takes the combined features and processes them to make the final stock price prediction.
 - Linear Activation Function: This makes sure that the final output comes out as a proper continuous value, like an actual stock price.

The final prediction is given by:

$$\hat{y}_t = W_f [h_t^{LSTM} \oplus h_t^{TFT}] + b_f \quad (3)$$

where $\oplus$ denotes concatenation, and W_f, b_f are learnable parameters.

C. Training Process

- **Data Collection**
 - Retrieve the stock price dataset using Python's YFinance library.
- **Data Preprocessing**
 - Normalize time-series data.
 - Convert sequences into input-output pairs.
 - Separate static, past, and future known variables for TFT.
- **Model Training**
 - Train LSTM and TFT independently on historical data.
 - Fine-tune the hybrid model using combined loss optimization.
- **Loss Function and Optimizer**
 - Loss Function: Mean Squared Error (MSE).
 - Optimizer Adam with an adaptive learning rate.
 - Batch Size: 64
- **Evaluation Metrics**
 - MAE (Mean Absolute Error)
 - MSE (Mean Squared Error)
 - R^2 Score (goodness-of-fit measure) - Root Mean Squared Error (RMSE).

D. Explainable AI (XAI)

- Perturb the Instance: For an instance x , generate perturbed instances x' by sampling around x .
- Prediction with the Complex Model: Obtain predictions $f(x')$ for each perturbed instance x' .
- Weight Perturbations: Assign weights to the perturbed instances based on their similarity to x , using a kernel function $\pi(x, x')$ (e.g., Gaussian):

$$\pi(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2\sigma^2}\right) \quad (4)$$

where σ is a scaling parameter.

- Train the Surrogate Model: Fit a simple, interpretable model $g(x')$ (e.g., linear regression) to the perturbed instances x' , weighted by $\pi(x, x')$ to approximate the behavior of the complex model $f(x)$ locally. $g(x') = \arg\min_{\pi(x, x_i)} \sum L(f(x_i), g(x_i))$ (5) where L is a loss function (e.g., squared error).
- Interpretation: The surrogate model $g(x')$ provides the explanation of the prediction for the original instance x .

E. Model Evaluation

Explainable AI (XAI) is crucial in stock prediction, as it is important not only to achieve accurate forecasts but also to comprehend the reasoning behind the model's predictions. This is where XAI techniques play a significant role. To incorporate Explainable AI into our model, we can utilize various algorithms such as LIME, SHAP, and Counterfactual Explanations to enhance transparency [1].

Purpose: Assess model performance on unseen data. Metrics:


- Root Mean Squared Error (RMSE):

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

- Mean Absolute Error (MAE):

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

- R-squared (R^2): Measures variance explained by the model.



```

StockPricePrediction_HybridModel(LSTM-TFT).ipynb
File Edit View Insert Runtime Tools Help

In[ ]: + Code + Text

# Calculate performance metrics
mae = mean_absolute_error(y_test, y_pred)
mse = mean_squared_error(y_test, y_pred)
rmse = np.sqrt(mse)
r2 = r2_score(y_test, y_pred)

# Display performance metrics
print(f"Model Performance Metrics:")
print(f"Mean Absolute Error (MAE): {mae:.4f}")
print(f"Mean Squared Error (MSE): {mse:.4f}")
print(f"Root Mean Squared Error (RMSE): {rmse:.4f}")
print(f"R2 Score: {r2:.4f}")

Model Performance Metrics:
Mean Absolute Error (MAE): 0.0060
Mean Squared Error (MSE): 0.0001
Root Mean Squared Error (RMSE): 0.0106
R2 Score: 0.9984

```

Fig 1. Model Evaluation

The above image presents the performance metrics of the Stock Price Prediction Model using the Hybrid LSTM-TFT approach. We have trained our model on 'AAPL' stock ticker of the Apple Inc. The model demonstrates exceptional accuracy with a Mean Absolute Error (MAE) of 0.0060, indicating minimal deviation between actual and predicted values. The Mean Squared Error (MSE) of 0.0001 and Root Mean Squared Error (RMSE) of 0.0106 further confirm the low prediction error. Most notably, the R^2 score of 0.9984 signifies that the model explains 99.84% of the variance in stock prices, proving its robustness. These results highlight the model's high precision and reliability.

Visual Evaluation: Below are the graphs comparing actual stock prices with predicted stock prices using models such as LSTM, TFT, and a Hybrid model of LSTM & TFT. For these graphs, I have selected the 'AAPL' stock ticker, representing Apple Inc.

➤ Stock Price Prediction using LSTM Model




Fig 2. Actual vs Predicted Price - LSTM Model

The graph represents stock price predictions using an LSTM model. The blue line indicates actual stock prices, while the red line represents predicted prices. The model captures trends well but shows some deviations.

➤ Stock Price Prediction using TFT Model




Fig 3. Actual vs Predicted Stock Price - TFT Model

The figure compares actual stock prices (blue line) with the predicted values (green line) obtained from the Temporal Fusion Transformer model. The predicted trajectory aligns well with the actual price trend, demonstrating the TFT model's strength in capturing temporal dependencies and seasonal variations. Although minor prediction fluctuations are observed, the overall direction and magnitude remain consistent, validating the model's forecasting capability.

➤ Stock Price Prediction using Hybrid Model

The graph illustrates the model's prediction performance over time. The blue line represents actual stock prices, while the orange line shows predictions generated by the hybrid model. The predicted values closely follow the actual trend, demonstrating the model's ability to learn and generalize temporal stock patterns effectively. Minor fluctuations in the predicted line reflect the model's sensitivity to short-term variations while maintaining overall accuracy.




Fig 4. Actual vs Predicted Stock Price - Hybrid Model

- Analyze prediction accuracy over time.

After training and evaluating multiple models, the best-performing one is selected based on metrics like accuracy, interpretability, and suitability for the problem.

F. Model Deployment: Once the best model is selected:

- Real-time Predictions: The model can be deployed to provide real-time stock predictions using livestock data feeds.
- API Integration: It can be integrated into a web or mobile application via APIs, where users can input stock data and get predictions along with explanations.

G. Hybrid LSTM-TFT Model: Working Principle:

The Hybrid Model integrates Long Short-Term Memory (LSTM) and Temporal Fusion Transformer (TFT) to improve time-series forecasting accuracy. Below is the step-by-step working principle along with mathematical formulas that describe how this model processes data.

- Problem Formulation:

Given a time-series dataset with historical observations X_t and known future covariates Z_t , the goal is to predict future values Y_t :

$$Y_{t+1:t+n} = f(X_{t-n:t}, Z_{t+1:t+n}) \quad (6)$$

Where:

X_t = Historical observed values (e.g., stock prices).

Z_t = Known future covariates (e.g., moving averages, economic indicators). Y_t = True target values. $\hat{Y}_t$ = Predicted values.

The hybrid model is defined as:

$$\hat{Y}_t = g(\text{LSTM}(X_t), \text{TFT}(X_t, Z_t)) \quad (7)$$

where g is a fusion function that combines LSTM and TFT outputs.

- LSTM Working Principle

LSTM is used to extract long-term dependencies from time-series data. It follows these key operations:

Forget Gate:

Decides which past information to keep or forget:

$$f_t = \sigma(W_f[h_{t-1}, x_t] + b_f) \quad (8)$$

Input Gate:

Updates the memory cell with new information:

$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i) \quad (9)$$

New candidate cell state:

$$\tilde{C}_t = \tanh(W_C[h_{t-1}, x_t] + b_C) \quad (10)$$

Cell State Update:

The memory cell is updated as:

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \quad (11)$$

Output Gate and Hidden State:

The new hidden state is computed as:

$$o_t = \sigma(W_o[h_{t-1}, x_t] + b_o) \quad (12)$$

$$h_t = o_t \odot \tanh(C_t) \quad (13)$$

Final Output from LSTM:

$$HLSTM = h_t \quad (14)$$

➤ TFT Working Principle

TFT is designed for multi-horizon forecasting and interpretable feature selection. It consists of: Variable Selection Network (VSN): TFT selects the most relevant features:

$$\alpha_j = \text{softmax}(W\alpha X_j + b\alpha) \quad (15)$$

$$X_j' = \alpha_j \odot X_j$$

where α_j is the attention weight for input feature X_j .

Gated Residual Network (GRN):

This module processes inputs through a gated skip connection: (16)

$$HGRN = \text{LayerNorm}(W_1X + b_1) + \text{Gating}(X) \quad (17)$$

where W_1 is a learnable weight matrix.

Temporal Attention Mechanism:

TFT applies multi-head attention to focus on important time steps:

$$A_t = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (18)$$

where:

Q, K, V = Query, Key, and Value matrices. d_k = Dimensionality of key vectors.

Final Output from TFT:

$$HTFT = A_t \quad (19)$$

➤ Fusion of LSTM and TFT Outputs

The outputs from LSTM and TFT are concatenated and passed through a fusion layer:

$$H_{\text{fusion}} = \sigma(W_f[HLSTM, HTFT] + b_f) \quad (20)$$

A fully connected layer is applied for final prediction:

$$\hat{Y}_t = W_o H_{\text{fusion}} + b_o \quad (21)$$

➤ Loss Function and Optimization

The model minimizes the Mean Squared Error (MSE):

$$L = \frac{1}{N} \sum_{i=1}^N (Y_i - \hat{Y}_i)^2 \quad (22)$$

Optimization is performed using AdamW optimizer:

$$\theta^{t+1} = \theta^t - \eta \left(\frac{\nabla L}{\sqrt{v_t} + \epsilon} \right) \quad (23)$$

where η is the learning rate and v_t is the adaptive moment estimate.

V. RESULTS AND DISCUSSION

Below are the outputs demonstrating how the models functionality works.




Fig 5. Stock price predictions for AAPL Company

The figure shows the predicted and actual stock price of Apple Inc. (AAPL) over the past month using the hybrid LSTM-TFT model. The actual latest price is 200.85, while the model predicted 204.97 dollars, showing a small difference of 4.12. This indicates the model's ability to closely follow real stock price movements. Such accurate predictions can help investors make better-informed decisions.




Fig 6. Model Performance

The figure presents the performance of the LSTM model in predicting stock prices, including key evaluation metrics and a comparison plot. The model's performance is summarized with an MAE of 18.83, MSE of 527.93, RMSE of 22.98, and a relatively low R^2 score of 0.098, indicating limited correlation between predicted and actual values. The line graph visually compares the actual stock prices with those predicted by the LSTM model, showing that while the model captures the overall trend, there are noticeable deviations in some regions. This highlights the model's ability to follow general price movements but also suggests room for improvement in prediction accuracy.



Fig 7.TFT Model Performance

The image shows the performance of the TFT (Temporal Fusion Transformer) model in predicting stock prices, along with key evaluation metrics. The model has a high Mean Absolute Error (MAE) of 30.91, Mean Squared Error (MSE) of 1161.57, and Root Mean Squared Error (RMSE) of 34.08, indicating significant prediction errors. The negative R^2 value of -0.9846 suggests that the model performs worse than a simple average-based prediction. The line graph clearly illustrates that the TFT model's predicted values deviate noticeably from the actual stock prices, showing it struggles to capture the real trends in the data.




Fig 8: Hybrid Model Performance.

The image presents the performance metrics and prediction results of a hybrid Ridge Regression model used for stock price prediction. The model achieved a Mean Absolute Error (MAE) of 6.93, Mean Squared Error (MSE) of 81.52, Root Mean Squared Error (RMSE) of 9.03, and an R-squared (R^2) value of 0.8607. These values indicate that the model fits the data well, with a high level of accuracy and low prediction error. The line graph compares the actual stock prices with the predicted values generated by the hybrid model. The close alignment between the two curves confirms the model's ability to effectively follow the trend and fluctuations of real stock prices over time.




Fig 9: Model Performance Metrics

The image compares the performance of three different models—LSTM, TFT, and a Hybrid Ridge Regression model—used for stock price prediction. The evaluation metrics show that the Hybrid model performed the best, with the lowest errors (MAE: 6.93, MSE: 81.52, RMSE: 9.03) and the highest R^2 score of 0.8607, indicating strong predictive accuracy. In contrast, the LSTM and TFT models showed higher error values and poor fit, especially the TFT model, which had a negative R^2 score. The line graphs below the metrics visually compare actual and predicted values, where the Hybrid model closely follows the actual trend, while LSTM and TFT show noticeable deviations, confirming the superior performance of the Hybrid approach.




Fig 10: Feature Impact On Predicted Price.

The image shows a LIME (Local Interpretable Model-Agnostic Explanations) interpretation, which explains how the past 10 days' closing stock prices influenced the model's prediction. The horizontal bar graph displays both positive (green) and negative (red) contributions of each day's closing price to the predicted stock price. Notably, the closing prices on Day -1 and Day -5 had the most significant positive impact (0.12 and 0.10, respectively), while Day -2 and Day -4 had the most negative impact. This visualization helps us understand which specific past days the model considered most important when making its prediction, improving transparency and interpretability.

LIME Explanation Table: Closing Price Impact (Past 10 Days)

Day	Feature	Impact on Predicted Price	Influence Type
1 Day -1	Closing Price (Day -1)	0.120000	Positive
2 Day -2	Closing Price (Day -2)	-0.060000	Negative
3 Day -3	Closing Price (Day -3)	0.050000	Positive
4 Day -4	Closing Price (Day -4)	-0.060000	Negative
5 Day -5	Closing Price (Day -5)	0.100000	Positive
6 Day -6	Closing Price (Day -6)	-0.040000	Negative
7 Day -7	Closing Price (Day -7)	0.040000	Positive
8 Day -8	Closing Price (Day -8)	-0.030000	Negative
9 Day -9	Closing Price (Day -9)	0.020000	Positive
10 Day -10	Closing Price (Day -10)	-0.010000	Negative

Fig 11: Impact Table.

The image shows a LIME explanation table that highlights how the past 10 days' closing stock prices influenced the predicted stock price. Each row represents one of the last 10 days, labeled from Day -1 (most recent) to Day -10. The "Impact on Predicted Price" column shows how much each day's closing price affected the prediction—positive values (in green) boosted the prediction, while negative values (in red) lowered it. The "Influence Type" column uses emojis and color coding to make it easier to understand whether the impact was helpful (Positive) or harmful (Negative) to the prediction. This table helps users interpret which recent prices most influenced the model's decision.

VI. CONCLUSION

This study presents an explainable AI-based approach for stock price prediction, integrating LSTM and Temporal Fusion Transformer (TFT) models with explainability techniques like LIME to provide transparency in decision-making. The system enables investors to make informed financial decisions by offering interpretable stock forecasts. Despite achieving promising results, challenges such as data volatility, external market influences, and model interpretability limitations remain. Future work will focus on improving real-time prediction accuracy, incorporating more financial indicators, and enhancing the explainability framework with advanced techniques like LIME to provide deeper insights into stock market trends.

ACKNOWLEDGMENT

We would like to express our sincere gratitude to the Department of Computer Science and Engineering, ADCET Ashta, for providing the necessary resources and support to carry out this research.

We also extend our thanks to our guide, Mr. A. R. Pradnyavant, for his valuable guidance, encouragement, and insights throughout the development of this project.

Lastly, we thank all the team members for their hard work, collaboration, and dedication that made this project successful.

REFERENCES

- [1] A. Saranya and R. Subhashini, "A systematic review of Explainable Artificial Intelligence models and applications: Recent developments and future trends," *Decision Analytics Journal*, vol. 7, p. 100230, 2023.
- [2] T. De, P. Giri, A. Mevawala, R. Nemani, and A. Deo, "Explainable AI: A Hybrid Approach to Generate Human-Interpretable Explanation for Deep Learning Prediction," *Procedia Computer Science*, vol. 168, pp. 40–48, 2020.
- [3] K. Olorunnimbe and H. Viktor, "Deep learning in the stock market—a systematic survey of practice, backtesting, and applications," *Artificial Intelligence Review*, vol. 56, no. 3, pp. 2057-2109, 2023.
- [4] V. R. R. Kovvuri, et al., "Fund performance evaluation with explainable artificial intelligence," *Finance Research Letters*, vol. 58, p. 104419, 2023.
- [5] S. Li, X. Zhang, and P. Chen, "PEN: Prediction-Explanation Network to Forecast Stock Price Movement with Better Explainability," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 4, 2023.
- [6] D. Muhammad, et al., "An explainable deep learning approach for stock market trend prediction," *Heliyon*, vol. 10, 2024.
- [7] N. Marey, A. A. Abu-Musa, and M. Ganna, "Integrating Deep Learning and Explainable Artificial Intelligence Techniques for Stock Price Predictions: An Empirical Study Based on Time Series Big Data," 2024.
- [8] P. Kumar, et al., "Analysing Forecasting of Stock Prices: An Explainable AI Approach," *Procedia Computer Science*, vol. 235, pp. 2009-2016, 2024.
- [9] P. Venkatasubbu, et al., "Explainable AI for stock price prediction in stock market," 2023, pp. 347-366.
- [10] Yahoo Finance, "Historical data for AAPL stock," [Online]. Available: <https://finance.yahoo.com/quote/AAPL/history>. [Accessed: Jan. 2025].