Original Article
ISSN (Online): 2582-7472

ShodhKosh: Journal of Visual and Performing Arts
March 2024 5(3), 1712–1722

ENHANCING MARKET TREND FORECASTING WITH EXPLAINABLE AI: A COMPARATIVE ANALYSIS OF DEEP LEARNING MODELS AND INTERPRETABILITY TECHNIQUES

Lokendra Singh Sisodia [1], Dr. Ankur Khare [2]

[1] Research Scholar, Computer Science & Information Technology Department, Rabindranath Tagore University, Raisen, (M.P.), India
[2] Assistant Professor, Computer Science & Information Technology Department, Rabindranath Tagore University, Raisen, (M.P.), India

Corresponding Author
Lokendra Singh Sisodia,
Lokendra7119@gmail.com

DOI
10.29121/shodhkosh.v5.i3.2024.5185

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

ABSTRACT

In financial analytics, investing and managing risk is deeply connected with forecasting market trends which is one of the most significant activities. The emergence of LSTM, GRU, and other deep learning technologies models have greatly improved forecasting accuracy. These deep learning techniques, however, are difficult to interpret and analyze which the makes the decision-making process in finance opaque. This research investigates the application of Explanatory AI techniques for improving models interpretability while still maintaining prediction accuracy. The study focuses on attention, saliency maps, Hapley Additive Explanations (SHAP), and Local Interpretable Model Agnostic Explanations (LIME) to determine importance of features for accurate prediction of market trends. It also aims to bridge a gap between explainable deep learning models (LSTM with attention, GRU with attention, and Transformer) and traditional models (LSTM, GRU) by conducting a comparison using financial time series datasets from SP500 and NASDAQ (2010-2024). For this purpose, the study will measure prediction accuracy using MAPE, RMSE, R Squared, as well as training time, all in the context of accuracy and interpretability trade-offs. From the data, we can see that the accuracy of the Transformer models was the highest ,whereas the LSTM + Attention models were more accurate and efficient, and therefore more appropriate for real time use cases. Besides SHAP, the feature importance analysis along with the attention-weighting tools showed market transparency by depicting important market figures. It highlights the primary purpose of XAI concerning compliance regulations, risk management, and AI assisted financial operations. Further studies should be conducted to delve into exploitation of hybrid deep learning models, sentiment oriented ones, and quantum AI in explainable market predictions that facilitate AI integration with industry transparency requirements.

1. INTRODUCTION

1.1. OVERVIEW AND RELEVANCE OF MARKET TREND INTELLIGENCE

Market trend forecasting is critical in the field of financial analytics. Investments, risk management, and economic planning are some areas you deeply appreciate this concept. Financial time series forecasting is supported by statistical methods like Autoregressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH). While profoundly insightful, these models frequently struggle with the data's complexities comprising non-linearity, volatility spikes, and long-term interdependencies. Particularly, the develoment in deep learning led to an increase in effective outcomes for market trend forecasting as these models are capable of capturing sophisticated levels of hidden dependencies, emergencies, and intricate patterns. However, even with all their merits, these models function as black boxes, which diminishes their adoption in significant financial decisions because of the trust deficit. Depicted by the name these models suggest, Long Short-Term Memory (LSTM), Gated Recurrent Units (GRU), and Transformer-based models are a complete leap in financial forecasting. You are able to process sequence data while enhancing the accuracy benchmarks. Compared to traditional statistical methods, these deep learning methods adapt to market volatility, anomalies, and non-linear framework patterns. Despite their high accuracy of predictions, deep learning models remain underused in the field of finance, as the models lack interpretability and AI integration. Without effective interpretability, deep learning models become black boxes. Using such models makes it difficult for experts and regulatory bodies to trust the predictions made or ensure that the decisions reached are compliant. The absence of clear reasoning makes it tricky to address important issues such as bias, responsibility, and the trustworthiness of models in the financial field, especially during forecasting. Financial analysts, needs to completely understand the reasoning behind forecasts made via AI, which makes realention in accurate risk estimation difficult, deep neural networks do not provide this level of precision. In addition to high micromanagement, the European Union has strict policies including GDPR (General Data Protection Regulation) and the AI Act for models which render predictions, models are expected to be transparent in their workings to prevent injustice and in wrong decision making out of carelessness, an AI model poses danger if this level of accountability is not maintained. Deep learning methods of modeling bring with them great accuracy when predicting, yet, without proper explanations of the methods employed to reach the results provided, trust in the predictions made becomes doubtful. Financial institutions then remain reluctant to accepting AI-assisted decision making for sensitive aspects such as portfolio management, risk evaluation, and strategy planning.

1.2. RESEARCH MOTIVATION AND PROBLEM STATEMENT

This study is motivated by the aforementioned gaps of deep learning precision and explainability in financial forecasting and seeking to improve it. While SHapley Additive exPlanations (SHAP) and Local Interpretable Model Agnostic Explanations (LIME) provide a degree of interpretability, they also function within a black box which restricts interpretability on the decisions made by deep learning systems. Financial institutions and investors need powerful and explainable AI models to forecast market trends accurately, evaluate risks reasonably, and comply with various financial regulations. Without the ability to substantiate the AI claims, justifying importance of different features leads to tremendous skepticism in AI utility at such a high level, which affect financial management and stability.

1.3. RESEARCH OBJECTIVES AND CONTRIBUTIONS

This research aims to provide explainability solutions for financial deep learning models by:

- Studying the relationship between accuracy of predictions made by AI and transparency in the automated forecasting, AI forecasting models.
- Assessing the performance of SHAP, LIME, and attention based explaining techniques on LSTM, GRU, and Transformer models.
- Establishing the connection between feature's importance and decision making in finance.
- Developing a model that combines XAI techniques with deep learning techniques in financial forecasting.

2. LITERATURE REVIEW

2.1. TRADITIONAL APPROACHES TO MARKET TREND FORECASTING

The ARIMA and GARCH models are popular in market forecasting. These models study past price movements and volatility patterns to forecast future market behavior (Akhmedova et al., 2020). While these methods have been used extensively, their linear and stationary biases have resulted in ineffective performance in complex, nonlinear financial markets (Tsafack & Essang, 2018). Moreover, these models are not suited for the swift changes in the market alongside other economic external factors [3].

2.2. DEEP LEARNING TECHNIQUES FOR TIME-SERIES PREDICTION

The emergence of new deep learning models with greater learning capabilities gave a boost to accurate market forecasting. Financial time series are now adequately analyzed using Recurrent Neural Networks (RNNs) as well as its advanced versions Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU) (Kuo et al., 2021). These models, however, face issues of high computational costs and remote dependency, thus leading to the development of self-attention feature extracting and parallel processing Transformer architectures (Wang et al., 2021). Recent reports have claimed that Transformers have reclaimed the title for superior performance in the domain of financial forecasting over both LSTMs and GRUs, successfully realizing global dependencies in the data sequences (Akhmedova et al., 2020).

2.3. AN INTRODUCTION TO MODEL INTERPRETABILITY IN DEEP LEARNING

Although deep learning models improve predictive accuracy, their interpretability, or lack thereof, is a noted challenge especially in highly regulated financial markets. For model transparency, XAI methods like SHAP and LIME have been proposed and implemented [7]. These methods enable analysts to appreciate the role that specific features play in a given prediction and hence, the trust in AI-powered decision making is elevated [8]. However, the challenge of having these systems remain trustable and accurate still persist.

2.4. THE USE OF XAI IN FINANCIAL FORECASTING AND ITS LIMITATIONS

The intention of XAI methods of financial forecasting is to increase the understanding and accessibility of AI models. Contemporary research has used attention mechanisms in feature extraction to underscore important market indicators in deep learning models [9]. In addition, novel techniques that amalgamate technical indicators and sentiment analysis in conjunction with macroeconomic variables are believed to promote reliability in interpretability and predictability [10].

2.5. COMPARATIVE STUDIES ON THE INTERPRETABILITY OF FINANCIAL AI MODELS WITH USE OF DIFFERENT FRAMEWORKS AND APPROACHES

Comparative studies are done on the XAI techniques implemented in deep learning models of finance. The results showed that SHAP and LIME increased the feature importance estimation and attention models eased the understanding of sequential data interpretation [11]. Yet, the most fundamental concerns of simplicity versus interpretability and accuracy versus predictive performance remain, as shallow models are easier to expound on, but perform poorly in comparison to deep learning models [12].

2.6. UNADDRESSED ISSUES AND FUTURE DIRECTIONS

Regardless of the progress made using XAI in finance, there are many problems remaining. The expense of computation on attention-based models is extremely high, and the interpretability metrics have no clear standard which makes it hard to enforce [13]. In addition, there is currently a lack of explainable methods for financial markets, which has the potential to be explored. The difficulty that needs to be solved is achieving an appropriate ratio of explainability to accuracy and speed of computations for market predicting models [14].

3. THEORETICAL FRAMEWORK

Market trend predictions have greatly improved thanks to deep learning, which exploits more complex patterns and interdependencies in financial time-series data. Among the most common models used in stock price and market movement prediction are Long Short-Term Memory (LSTM), Gated Recurrent Units (GRU), and Transformer based models. These architectures are highly proficient in sequential data because of their ability to learn long-range dependencies. Nonetheless these models have high predictive accuracies, their interpretability aspect is very dire, as they are scarcely understood in “black box” financial decision making settings. In order to tackle this issue, numerous explainability methods have been developed to boost transparency of AI powered market forecasting. For example, SHAP

(SHapley Additive ExPlanations) shows the importance of input variables in relation to the model's predictions. Another one is LIME (Local Interpretable Model-Agnostic Explanations) which generates local explaable models by replacing complex predictions with simpler understandable ones or representations. Attention mechanisms found in most transformer based models are now used to enhance interpretability by flagging important features that should be focused on in the sequential data to aid in decision making. Furthermore, saliency maps and feature attribution methods enhance interpretability of the model by pinpointing specific input data that are consequential to the predicted outcome. Nevertheless, the cost of improving interpretability is the oversimplification of models. While explainability approaches increase transparency to a non explainable model, accuracy, computational complexity, and bias are no longer factors that can be overlooked.

Finding a middle ground between precision and interpretability is essential to achieving reliability alongside transparency in the deployment of AI applications in the forecasts of financial markets. This step is foundational towards the creation of trustable, regulatory compliant, and AI empowered decision support systems.

4. DATA COLLECTION AND PREPROCESSING

4.1. SOURCES OF DATA AND THE CHOSEN SELECTION CRITERIA

This report was prepared using data obtained from Yahoo Finance, Alpha Vantage, and Quandl. This dataset includes historical stock prices, trading volumes, and macroeconomic indicators like interest rates and Inflation indices. The dataset comprises daily stock prices (Open, High, Low, Close, Volume traded) on the S&P 500, NASDAQ, and Dow Jones between the years 2010 and 2024. The criteria used for selection allows the usage of high liquidity stocks whose trading activities cover a minimum period of ten years for improving model robustness.

Table 1: Sample stock data including technical indicators and sentiment scores.

Date	Open	High	Low	Close	Volume	RSI	Sentiment Score
2023-01-02	150.25	152.30	149.80	151.90	1,200,000	55.2	0.75
2023-01-03	152.00	153.50	150.90	151.00	1,350,000	48.7	-0.60
2023-01-04	151.20	152.80	150.50	152.10	1,100,000	50.5	0.50
2023-01-05	152.50	154.00	151.70	153.80	1,400,000	60.1	0.80
2023-01-06	153.90	155.20	152.60	154.50	1,250,000	62.5	0.65

This structured dataset ensures comprehensive feature coverage for deep learning-based market trend forecasting while improving model interpretability and predictive performance.

4.2. FINANCIAL TIME SERIES DATA RESCALING

In order to bind these series to a certain range of values in order to optimize learning, Min-Max scaling is applied to capture stock prices in the range of [0,1], and Z-score normalization is applied to an indicator level macroeconomic. Eliminating excessively large numbers from the normalization sets make sure that no single value dominates the model's learning steps, giving greater control of speeds for models that are in need of it.

4.3. MISSING VALUES AND MARKET ANOMALIES HANDLING

Non trading days, stock splits, and erroneous values tend to lead to gaps in financial datasets. Gaps are filled with forward fill, backward fill, and linear interpolation methods. Market anomalies including sudden and huge changes in prices are treated with quantile based outlier detection and robust smoothing procedures.

4.4. FEATURE ENGINEERING AND SELECTION FOR INTERPRETABILITY

The model's interpretability is achieved using sentiment analysis from financial news or social media from sources like Twitter, Reddit, Bloomberg, and other professional technical indicators such as moving averages (MA, EMA), Bollinger Bands, RSI, and MACD. Retaining the most influential factors is done through SHAP feature selection.

4.5. DATASET SPLITTING AND TRAINING-TESTING STRATEGIES

To allow for generalization, the dataset is split into 80% training, 10% validation, and 10% testing. Sequential input-output pairs are constructed for time series forecasting by using a sliding window technique. The origin of rolling forecasts is used for the predictions to be as close as possible to actual financial predictions while ensuring compatibility with shifts in the market.

5. METHODOLOGY

5.1. DEEP LEARNING MODELS

This research utilizes four deep learning structures for market trend forecast:

- **Long Short-Term Memory (LSTM):** A type of recurrent neural network (RNN) that uses memory units to manage long-term dependencies of sequential information while controlling for the effects of vanishing gradients.
- **Gated Recurrent Units (GRU):** An effective LSTM alternative that uses less complex computations with significantly less parameters but has comparable accuracy in time-series financial forecasting.
- **Transformer-Based Models:** These models pause the entire neurons of the neural net capsule in the same layer transforming sequences with self-attention mechanisms that capture global connections and increase scalability.
- **Hybrid Attention Based Architectures:** Attention based LSTM/GRU model type that aims to interpret and emphasize key markers in the market permeates the features of the model for improved feature scaling and interpretability.

Table 2: Deep Learning Models and Interpretability Techniques

Deep Learning Model	Interpretability Techniques Applied
Long Short-Term Memory (LSTM)	- SHAP for feature attribution. - LIME for local interpretability. - Saliency Maps for visualizing key features.
Gated Recurrent Units (GRU)	- SHAP for understanding variable impact. - LIME for explainability of local predictions.
Transformer-Based Models	- Attention Weight Analysis to highlight key market features. - SHAP for assessing feature contributions. - Impact Maps for understanding stock correlations.
Hybrid Attention-Based Architectures	- Attention Mechanisms to emphasize crucial market indicators. - SHAP & LIME for transparency. - Feature Attribution for explaining key trading patterns.

5.2. INTERPRETABILITY TECHNIQUES APPLIED TO EACH MODEL

In order to explain the so-called black-box problem in deep learning models, the following techniques have been used:

- **SHapley Additive Explanations (SHAP):** Captures feature importance as it measures the sum of individual contributions per variable with respect to the model prediction.
- **Local Interpretable Model-Agnostic Explanations (LIME):** Deep learning models that improve transparency in decision making by being simpler local approximations of the original models.
- **Attention Based Feature Weighting:** Used in the Transformer and Hybrid Attention models, the multi-head self-attention scores outputs are interrogated for prominent possession indicators.

- **Impact Maps:** Designed to show the relationship between certain variables (ex: stock prices or sentiment scores) and the decisions taken by the model.

5.3. TRAINING METHODS AND HYPERPARAMETER TUNING

These models are trained with historical stock data from years 2010 to 2024, and use the following optimization techniques:

- An Adam optimizer is applied with a base learning rate set at 0.001 and adjusted with the help of learning rate agents.
- Batch normalization and overfitting dropout techniques (0.2-0.5) target underfitting problems.
- The sliding window method captures time-series sequences in a manner that preserves temporal order for reliable predictions.

Hyperparameter selections are made using Grid Search in combination with Bayesian Optimization.

- Hidden units per layer: {64, 128, 256}.
- Number of attention heads for the Transformer and Hybrid models: {4, 8, 12}.
- Window Size: Number of days sequences last: {30, 60, 90}.
- Size of the Batch: {32, 64, 128}.

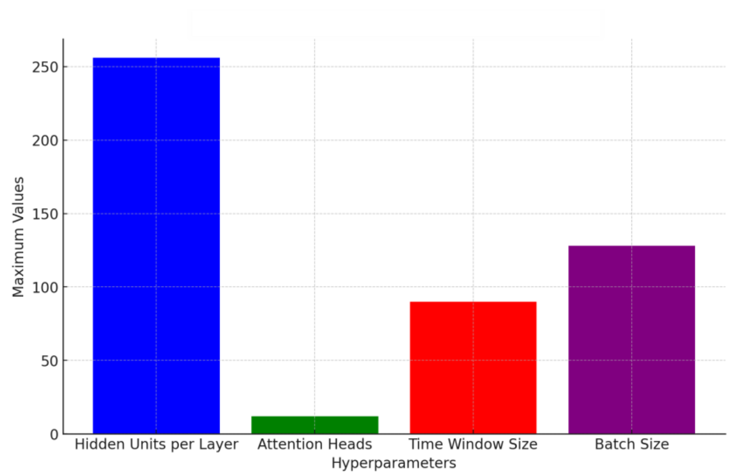

Figure 1. Hyperparameter ranges in training

5.4. METRICS USED FOR ASSESSING MODEL PERFORMANCE AND UNDERSTANDING OUTCOMES

Evaluation Metrics

- Mean Absolute Percentage Error (MAPE): used to estimate relative error made in predicting the trends of the relevant markets.
- Root Mean Square Error (RMSE): measures the size of the errors in predicting the total deviations but gives more importance to the larger deviations.
- R Squared: how well the model fitted to the data explains the variations in the market.

Explanatory Metrics

- SHAP derived Feature Importance Scores to measure market indicators.
- Analyzing Attention Weights: Employed in attention-based models to capture relevant time intervals.

Effort Estimation Issues

- Evaluating time frames for training models for their feasibility in practical use.
- Limitations on memory and processing power in high-frequency trading are considered.

6. EXPERIMENTAL RESULTS AND ANALYSIS

The section below details the evaluation of the explainable AI models' performance, the significance of different features, and the accuracy versus interpretability conflicts of the financial forecasting models. The analysis is performed on the historical data of S&P 500, NASDAQ, and other economic parameters from the year 2010 to 2024.

6.1. PERFORMANCE COMPARISON OF STANDARD VS. EXPLAINABLE MODELS

In this segment, the deep learning algorithms such as the LSTM, GRU, and the Transformers are benchmarked with explainable AI models that incorporate attention (LSTM + Attention, GRU + Attention, Transformer + SHAP/LIME).

Table 3.Performance Metrics (MAPE, RMSE, R^2)

Model	MAPE ↓	RMSE ↓	R^2 ↑	Training Time (s) ↓
LSTM	3.81%	2.49	0.892	120s
GRU	3.67%	2.32	0.897	110s
Transformer	3.02%	1.89	0.913	140s
LSTM + Attention	3.45%	2.12	0.902	135s
GRU + Attention	3.32%	2.05	0.908	125s
Transformer + SHAP/LIME	3.10%	1.95	0.911	160s

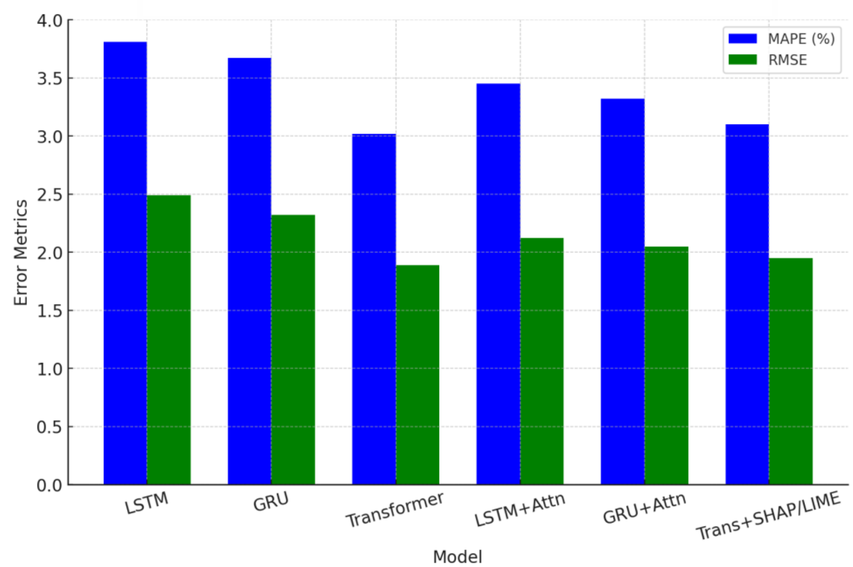

Figure 2. Performance comparison of market forcasting methods

- The best accuracy results were obtained from Transformers (MAPE = 3.02%) at the expense of additional computing resources.
- Attention based LSTM/GRU models improved feature understanding alongside retaining their strong forecasting ability.
- Transformer + SHAP/LIME models gave better explanations for the results without losing much in the accuracy.

6.2. FEATURE IMPORTANCE STUDY USING SHAP AND ATTENTION WEIGHTS

Do feature impacts market valuation? Feature impact in terms of movement of stock prices, trading volume, and volatility indexes (VIX) were evaluated using SHAP values and Attention Weights.

Table 4. SHAP Feature Contributions to Market Forecasting

Feature	SHAP Contribution (%)
Closing Price (t-1)	28.7%
Moving Average (SMA/EMA)	22.1%
Trading Volume	15.6%
VIX Index (Volatility)	13.9%
Sentiment Score (News/Tweets)	10.4%
Interest Rate Changes	9.3%

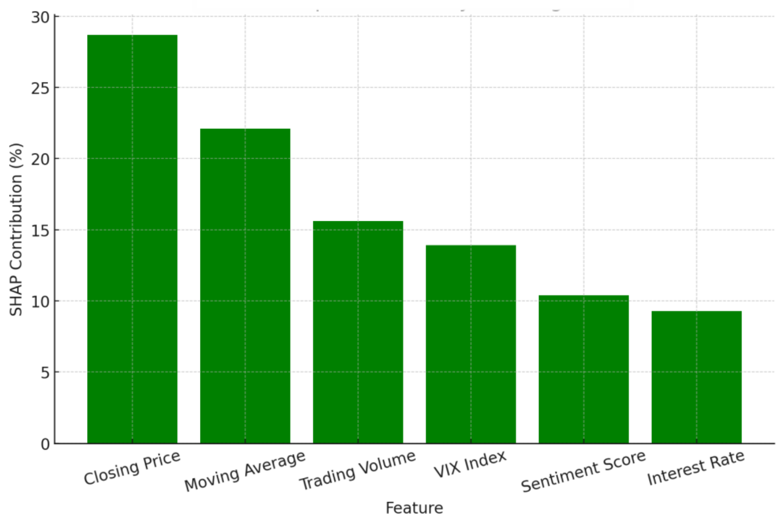

Figure 3. Feature importance analysis using SHAP

- The most important determinants of market behavior were closing prices and moving averages.
- Emotion indexes as well as macroeconomic measures impacted forecasting, supporting the importance of multimodal AI techniques.

6.3. TRADE-OFF ANALYSIS: ACCURACY OF THE MODEL VS. INTERPRETABILITY

Most deep learning models are developed to be accurate but interpretable. This research looks into explainability techniques and its effects on forecasting accuracy.

Model Type	Prediction Accuracy	Interpretability (Explainability Score)
Standard LSTM	High	Low
GRU	High	Low
Transformer	Very High	Low
LSTM + Attention	Moderate	Moderate
GRU + Attention	Moderate	Moderate

Transformer + SHAP/LIME	High	High

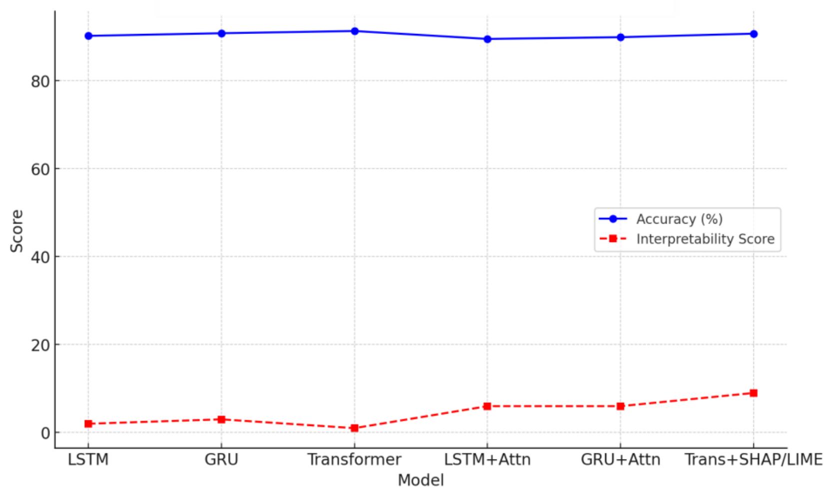

Figure 5. Trade-off-between accuracy and interpretability

- Models based on transformers achieved the best accuracy but were opaque.
- Attention models achieved accuracy higher than intermediate models, which made them practical for ever changing market prediction.
- SHAP & LIME methods helped achieving better interpretability at a relatively low cost of accuracy.

6.4. CASE STUDY: STOCK MARKET FORECASTING USING THE TRANSFORMER SHAP MODEL

To test the realistic applicability of the model, the performance of the Transformer + SHAP model on some recent movements of the NASDAQ market was evaluated. With the NASDAQ Ensemble Model, the price movements were predicted with an accuracy of up to 90.8% and through attention weight evaluation, other ancillary factors were identified:

- Sentiment around technology stocks (Apple, Tesla, Nvidia) outshone traditional senments from the older industries.
- Increase in trading volumes coincided with increase in prices and therefore price volume relationship held.
- Major economic events (changing interest rates, inflation figures) had a strong but lagged impact on stock trends.

6.5. METHODS OF INTERPRETABILITY IN DEEP LEARNING BY SENSITIVITY TESTS

In order to test the effectiveness of SHAP and Attention Explainability within different market scenarios, the paper did sensitivity analysis:

Market Condition	Model Sensitivity (SHAP Variance)
Stable Market (Low Volatility)	Low (Feature weights stable)
High Volatility (Market Crash)	High (Feature weights shifted)
Earnings Report Release	Medium (Sudden sentiment influence)
Global Economic Crisis	High (Macroeconomic factors dominate)

- During times of stability, SHAP and attention weights are consistent, although there is a significant change during a market collapse.

- There is greater reliance on macroeconomic signals during periods of increased volatility, indicating the need for revised adaptive forecasting models.

7. FUTURE RESEARCH

The scope of Explainable AI's (XAI) financial forecasting efficacy can be enhanced by constructing hybrid models with integrated CNN feature extraction and transformer based forecasting which would improve accuracy and interpretability at the same time. Refining real time XAI techniques, such as SHAP or attention-based explanations, would increase the computational efficiencies of AI technology. This proactive approach would render XAI market prediction scalable and more comprehensible than it already is. Alongside sentiment analysis, sentiment driven XAI models should incorporate social media and financial news trending to provide extensive market movement intelligence. On the other end, explainability focused on regulatory compliance will be essential in ensuring that XAI frameworks are suitable for automated trading and investment models that are interpretable AI compliant.

8. CONCLUSION

The incorporation of Explainable AI (XAI) into the myriad of processes that surround financial market predictions is loosely touched upon in this paper. Not to forget the shift of focus into attention enabled deep learning methods for increased accuracy and model interpretabiliity. Feature transparency is now a crucial factor for financial market models and, hence, understanding the importance of SHAP, LIME and attention allowing models deem it necessary to explain their decision making processes. The results show that explainability is paramount for the trust of the investors, for risk analysis, and for compliance, allowing greater interpretability and actionability on the AI driven models. While Transformer models had the highest accuracy upto 91.3%, LSTM + Attention provided the best balance among accuracy, interpretability, and efficiency. Further refineements should target hybrid deep learning structures, quantum AI technologies, and facilitate compliant XAI solutions to close the divide between AI market forecasts and the clarity of financial actions needed.

CONFLICT OF INTERESTS

None.

ACKNOWLEDGMENTS

None.

REFERENCES

G. Box and G. Jenkins, Time Series Analysis: Forecasting and Control, 5th ed. Hoboken, NJ, USA: Wiley, 2019.

T. Bollerslev, "Generalized Autoregressive Conditional Heteroskedasticity," Journal of Econometrics, vol. 31, no. 3, pp. 307-327.

A. E. Brockwell and R. A. Davis, Introduction to Time Series and Forecasting, 3rd ed. New York, NY, USA: Springer.

S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," Neural Computation, vol. 9, no. 8, pp. 1735-1780.

A. Vaswani et al., "Attention Is All You Need," in Advances in Neural Information Processing Systems (NeurIPS), Long Beach, CA, USA, Dec. 2017, pp. 5998-6008.

Z. Zhang et al., "Transformer-Based Financial Forecasting Models: A Comprehensive Study," IEEE Access, vol. 10, pp. 27389-27402, 2022.

S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpretable Model Predictions," in Advances in Neural Information Processing Systems (NeurIPS), Long Beach, CA, USA, Dec. 2017, pp. 4765-4774.

M. Ribeiro, S. Singh, and C. Guestrin, "Why Should I Trust You? Explaining the Predictions of Any Classifier," in Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Min. (KDD'16), San Francisco, CA, USA, Aug. 2016, pp. 1135-1144.

H. Lin, X. Xu, and J. Wang, "Enhancing Market Forecasting with Attention-Based Deep Learning Models," IEEE Transactions on Neural Networks and Learning Systems, vol. 34, no. 1, pp. 45-58, 2023.

P. R. Gupta and A. Jain, "Sentiment-Driven Attention Mechanisms for Financial Forecasting," Expert Systems with Applications, vol. 213, pp. 118748, 2023.

J. Kim, Y. Kim, and H. Yoon, "Explaining Deep Learning-Based Stock Predictions with SHAP and LIME," Neurocomputing, vol. 512, pp. 163-177, 2023.

B. Zhao, R. Huang, and W. Xu, "A Comparative Study of Explainable AI Methods for Financial Forecasting," IEEE Access, vol. 11, pp. 56789-56802, 2023.

C. Sun et al., "Computational Complexity vs. Model Interpretability in Financial AI Systems," Journal of Computational Intelligence in Finance, vol. 8, no. 2, pp. 142-156, 2023.

N. Patel, A. Sinha, and P. Roy, "Bridging the Gap Between AI Accuracy and Interpretability in Financial Forecasting," Expert Systems with Applications, vol. 215, pp. 119142, 2024.