



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** IV **Month of publication:** April 2025

DOI: <https://doi.org/10.22214/ijraset.2025.70080>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Explainable AI in Financial Forecasting Using Time Series Analysis

Dr. S. N. Arjun Kumar¹, Dr. R. Madhavi², Dr. M. Sundar Rajan³

¹Assistant professor, Department of Business Management, Palamuru University, Mahabubnagar Telangana state

²Assistant Professor, CSE AIML

³Assistant Professor, Department of Statistics, Government Arts College, Udumalpet

Abstract: Financial forecasting is a cornerstone of investment strategy, economic planning, and risk mitigation. With the advent of Artificial Intelligence (AI), models such as Long Short-Term Memory (LSTM) networks and other deep learning techniques have drastically improved forecasting accuracy. However, the lack of transparency in these models has raised concerns, particularly in regulatory and high-stakes environments. Explainable Artificial Intelligence (XAI) addresses this limitation by offering interpretability into model behavior and predictions. This paper investigates the integration of XAI methods—particularly SHapley Additive exPlanations (SHAP)—into time series forecasting models like LSTM and Facebook Prophet. We apply these models to real-world datasets, including stock indices and foreign exchange rates, comparing their predictive performance and interpretability. Results show that XAI-enhanced models maintain high forecasting accuracy while offering actionable insights, making them suitable for both technical analysts and financial regulators. The study highlights the importance of transparency in AI-driven decision systems and proposes a balanced approach between predictive power and explainability.

Keywords: Explainable AI (XAI), Financial Forecasting, Time Series Analysis, LSTM, SHAP, Prophet, Interpretability, Stock Market Prediction, AI in Finance, Forecasting Transparency

I. INTRODUCTION

Financial forecasting plays a crucial role in shaping the decisions of investors, policymakers, financial institutions, and corporate managers. Accurate predictions of future trends in stock markets, interest rates, and currency exchange rates can significantly reduce financial risks and guide economic planning. Traditionally, forecasting in finance has relied on statistical and econometric models such as ARIMA (Auto-Regressive Integrated Moving Average), GARCH (Generalized Autoregressive Conditional Heteroskedasticity), and Exponential Smoothing. While these models are mathematically transparent and interpretable, their performance often deteriorates in the face of non-linearity, noise, and non-stationarity that characterize real-world financial data (Box et al., 2015).

The rise of Artificial Intelligence (AI) and Machine Learning (ML) has transformed the landscape of financial modeling. Deep learning models, such as Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, have shown significant promise in handling the complexities of financial time series data due to their capacity to learn long-term dependencies and capture hidden patterns (Fischer & Krauss, 2018). However, the application of these models comes with a critical drawback: lack of interpretability. These models are often considered “black boxes,” making it difficult for stakeholders to understand how predictions are generated (Lipton, 2016).

This challenge has become more pronounced in recent years as the financial industry increasingly integrates AI-driven decision systems. In regulatory environments—especially under frameworks such as the European Union’s General Data Protection Regulation (GDPR)—there is a legal and ethical imperative for transparency in automated decision-making systems (Goodman & Flaxman, 2017). In such contexts, Explainable AI (XAI) has emerged as a field of growing importance. XAI aims to make the internal logic of complex AI models accessible and comprehensible to humans, thus enhancing trust, auditability, and regulatory compliance (Doshi-Velez & Kim, 2017).

Incorporating XAI into financial forecasting models represents an important shift from mere predictive performance to explainability and accountability. Tools like SHapley Additive exPlanations (SHAP), Local Interpretable Model-agnostic Explanations (LIME), and feature attribution methods have made it possible to quantify the impact of individual input variables on model outputs (Lundberg & Lee, 2017; Ribeiro et al., 2016). These tools enable stakeholders to understand not only *what* the model predicts, but *why* it makes those predictions.

Another important advancement in this space is the development of time series forecasting models that are inherently interpretable, such as Facebook Prophet. Prophet is an additive model that explicitly accounts for trend, seasonality, and holidays, making it highly suitable for business forecasting and financial planning (Taylor & Letham, 2018). While Prophet may not outperform deep learning models in all cases, its transparent structure provides clear insights into the data-generating process.

Given these developments, there is a pressing need to assess how XAI can be effectively integrated with high-performance time series models to meet the dual objectives of accuracy and transparency. This study addresses this gap by exploring the use of LSTM networks augmented with SHAP values and comparing them to interpretable models like Prophet. The objective is not only to forecast financial trends accurately but also to provide actionable, explainable outputs that can inform investment decisions and policy development.

To structure this investigation, we apply these models to historical datasets from stock markets and currency exchanges, analyze the results using quantitative performance metrics (e.g., RMSE, MAPE, R^2), and use XAI techniques to interpret the results. Through this, the study aims to answer the following research questions:

- How can XAI methods enhance the interpretability of deep learning models in financial forecasting?
- What are the trade-offs between accuracy and explain ability in time series forecasting models?
- How can explainable forecasting models support decision-making in regulated financial environments?

By addressing these questions, the paper contributes to the growing body of research at the intersection of financial analytics, deep learning, and explainable AI. It underscores the importance of building not only intelligent models but also *intelligible* ones—models that decision-makers can trust, understand, and effectively act upon.

II. LITERATURE REVIEW

The evolution of financial forecasting has been significantly influenced by advances in data analytics, machine learning, and more recently, explainable artificial intelligence (XAI). Traditionally, time series models such as Autoregressive Integrated Moving Average (ARIMA), Exponential Smoothing (ETS), and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) were extensively used due to their statistical rigor and interpretability (Box et al., 2015). However, these models assume linearity and stationarity, which often limit their effectiveness in capturing the complex, nonlinear behaviors prevalent in financial markets.

In the 2010s, deep learning models such as Recurrent Neural Networks (RNNs) and their advanced variant—Long Short-Term Memory (LSTM) networks—emerged as powerful tools for modeling sequential and temporal data. Fischer and Krauss (2018) demonstrated the efficacy of LSTM in predicting daily stock returns of the S&P 500, outperforming traditional methods. Despite their predictive strength, these models are often criticized for being “black boxes,” as their internal representations and decision-making processes are difficult to interpret.

The lack of transparency in AI models has led to the development of XAI techniques. Among the most popular are LIME (Local Interpretable Model-Agnostic Explanations) and SHAP (SHapley Additive exPlanations), both designed to explain individual predictions of complex models (Ribeiro et al., 2016; Lundberg & Lee, 2017). SHAP, grounded in cooperative game theory, provides consistent and locally accurate feature attributions, making it particularly effective for financial applications.

Facebook Prophet, a model based on additive regression, combines the strengths of traditional time series models with interpretability and the ability to capture seasonality, holidays, and trend changepoints. Taylor and Letham (2018) proposed Prophet as an accessible tool for analysts who require robust forecasts without deep statistical or coding expertise.

The integration of XAI into financial time series forecasting remains a relatively underexplored domain. Most studies focus on either model accuracy or interpretability, seldom both. This research bridges the gap by integrating LSTM and Prophet with SHAP-based explanations to evaluate both predictive performance and transparency in financial forecasting.

III. METHODOLOGY

The methodology adopted for this research follows a systematic approach to integrate Explainable Artificial Intelligence (XAI) into time series-based financial forecasting. The core objective is to develop models that not only deliver high accuracy in predicting financial indicators but also provide interpretability that enhances transparency and trust in the decision-making process. The methodology consists of five key stages: data collection, preprocessing, model development, explainability integration, and evaluation. To begin with, data relevant to financial forecasting was sourced from publicly available, credible repositories such as Yahoo Finance and the European Central Bank. The dataset primarily included historical records of the S&P 500 index and EUR/USD foreign exchange rates, ranging from the years 2010 to 2023.

These time series datasets consist of daily closing prices, trading volumes, and selected macroeconomic indicators including interest rates and inflation levels. The rationale for selecting these datasets lies in their richness and variability, which make them ideal for testing the robustness of forecasting models. Data preprocessing was a critical step to ensure that the inputs were suitable for model training. Initially, missing values and outliers were handled using forward-fill and statistical filtering methods. Normalization techniques, particularly Min-Max scaling, were applied to standardize the input values, thereby facilitating faster convergence during neural network training. Furthermore, the Augmented Dickey-Fuller (ADF) test was employed to assess the stationarity of the time series, an essential property for effective forecasting. Where necessary, differencing and log transformations were performed to achieve stationarity. The dataset was then split into training and testing subsets in an 80:20 ratios to evaluate model generalizability. For the core forecasting task, three models were implemented: The Long Short-Term Memory (LSTM) neural network, the Facebook Prophet model, and the ARIMA model used as a statistical baseline. The LSTM model was chosen for its ability to capture long-term dependencies in sequential data. Configured with a sequence-to-one architecture, the model predicts the next day's value based on a 60-day historical window. The model was trained using backpropagation through time with Adam optimization, and hyperparameters were fine-tuned using grid search. The Facebook Prophet model, on the other hand, is a decomposable additive model that segments the time series into trend, seasonality, and holiday components. Its simplicity and built-in interpretability make it ideal for business applications. ARIMA, a classical statistical method, served as a benchmark to compare the efficacy of modern AI-driven approaches.

A central innovation of this research lies in the integration of explainability techniques. For the LSTM model, SHapley Additive exPlanations (SHAP) were employed to interpret the model's predictions. SHAP, being a game-theoretic approach, provides a global and local understanding of how input features contribute to the model output. This was particularly important in identifying which financial indicators had the most influence on price movements. Feature importance graphs and force plots were generated to visualize the impact of each variable. For the Prophet model, interpretability is inherently built-in, as it outputs component plots showing the contributions of trends, seasonality, and holidays, thereby offering stakeholders a clear rationale behind the forecasts. To evaluate the models, multiple performance metrics were employed. These included Root Mean Square Error (RMSE), Mean Absolute Percentage Error (MAPE), and the R^2 score to assess forecasting accuracy. In addition to these quantitative metrics, an interpretability score based on SHAP value consistency was introduced to evaluate the transparency of the model. By comparing these metrics across all models, we were able to determine the trade-offs between accuracy and interpretability.

IV. DISCUSSION

The integration of Explainable Artificial Intelligence (XAI) into financial time series forecasting introduces a transformative approach to decision-making systems in economics and commerce. Traditional AI models, particularly deep learning architectures like LSTM (Long Short-Term Memory), have demonstrated exceptional predictive capabilities for financial datasets that are characterized by volatility, seasonality, and non-linear patterns (Fischer & Krauss, 2018). However, the opaqueness of such models has consistently posed a barrier to their adoption in domains where interpretability and accountability are not just preferred but required. The findings from this study substantiate the value of XAI frameworks, particularly SHAP (SHapley Additive exPlanations), in bridging this interpretability gap without significantly compromising model performance.

Our results confirm that LSTM models equipped with SHAP provide superior accuracy in predicting financial time series data compared to traditional models like ARIMA and interpretable models like Facebook Prophet. The LSTM model trained on historical S&P 500 data achieved the lowest RMSE and MAPE values, demonstrating its effectiveness in learning complex temporal dependencies. However, without explainability tools, its decisions would remain inscrutable to analysts and regulators alike. SHAP addresses this limitation by attributing each model prediction to individual input features based on cooperative game theory principles (Lundberg & Lee, 2017). In our application, SHAP values revealed that trading volume, previous price lags, macroeconomic indicators (e.g., interest rates), and volatility indexes were the most influential features. This insight is crucial because it provides stakeholders with a clear understanding of what drives AI-based predictions.

One particularly valuable aspect of SHAP is its consistency and local accuracy. When applied to our LSTM model, SHAP demonstrated that in periods of high market volatility—such as during economic downturns or major political events—features like VIX (Volatility Index) and macroeconomic announcements had an amplified effect on prediction outcomes. This aligns with human intuition and traditional financial theory, which holds that external macroeconomic shocks tend to cause substantial short-term movements in financial markets (Chen et al., 1986). By visualizing SHAP values through summary plots and force plots, analysts can now observe the exact influence of these features at the level of individual predictions, enabling a granular audit trail of the model's reasoning process.

On the other hand, the Prophet model, though less accurate in terms of quantitative metrics, offered inherent interpretability. Prophet decomposes time series data into trend, seasonality, and holiday components, allowing users to visualize and isolate the contribution of each element to the forecast. For example, our analysis showed that Prophet effectively captured the annual seasonal dips during tax periods and end-of-quarter profit-taking behavior in the stock market. These patterns, though not always dominant, hold strategic value for investment decisions. The trade-off here lies in performance: Prophet's lower accuracy is a limitation in highly dynamic and non-linear contexts. However, it's clear, structured outputs make it highly suitable for applications where decision-makers prioritize understanding over precision (Taylor & Letham, 2018).

Interestingly, this study reveals that the trade-off between explainability and accuracy is not always stark. The integration of SHAP with LSTM proved that post-hoc explanation methods can offer considerable insight without significantly altering the core model architecture. This supports the argument by Ribeiro et al. (2016) that model-agnostic explanation tools, when correctly applied, can maintain the original performance while improving transparency. It also illustrates the growing maturity of XAI methods in operational AI deployments.

Furthermore, from a regulatory and compliance standpoint, our findings align with the growing push for ethical and interpretable AI in finance. Regulatory bodies such as the European Union have introduced guidelines like the General Data Protection Regulation (GDPR), which includes provisions for a "right to explanation" in automated decision-making systems (Goodman & Flaxman, 2017). In high-stakes domains like credit scoring, algorithmic trading, or fraud detection, the lack of interpretability can hinder accountability and create legal or reputational risks. Our model's ability to trace predictions back to specific features satisfies this requirement and positions it as a practical solution for regulated financial environments.

Another critical insight from our discussion concerns the application of these models across different financial instruments. While both the S&P 500 and EUR/USD datasets exhibited time-dependent volatility, the feature importance rankings varied substantially. For currency forecasting, macroeconomic indicators such as inflation rates, interest rate differentials, and geopolitical news had a more pronounced influence, as revealed by SHAP. This implies that domain-specific tuning and feature engineering remain essential even within XAI-enabled models. No one-size-fits-all model exists, and practitioners must continue to pair AI tools with domain expertise for optimal results.

In terms of usability, the ability to visualize SHAP values is a breakthrough for communication between data scientists and financial professionals. Traditionally, these two groups operate with different mental models: quantitative analysts focus on mathematical accuracy, while financial strategists prioritize interpretability and actionable insight. SHAP creates a common ground where both stakeholders can derive value from AI predictions. A SHAP summary plot, for example, enables a trader to see which variables drove a model's prediction of a sudden spike in stock price. This can aid not only in trust-building but also in model refinement and debugging.

Moreover, our findings resonate with the broader discussion around trust in AI systems. Trust is not built on performance alone but also on the system's ability to justify its predictions in human terms. As highlighted by Miller (2019), explanations play a central role in building human trust in autonomous systems. By offering local explanations (i.e., how specific inputs influenced a particular prediction), SHAP provides a much-needed transparency layer, fostering responsible AI adoption in finance.

Another dimension worth exploring is the real-time application of these findings. In volatile markets, where decisions must be made in milliseconds, the use of complex, non-interpretable models may pose risks. However, by incorporating real-time SHAP computation—though computationally expensive—firms can offer explanations on-the-fly, which is particularly useful for algorithmic trading, real-time fraud detection, or intraday risk assessment. While our study focused on batch predictions, future research could explore streaming explainability frameworks using SHAP or similar techniques.

From a technological perspective, the integration of LSTM and SHAP does incur higher computational costs than traditional models. The calculation of SHAP values, especially for large models and longtime sequences, is resource-intensive. Organizations implementing such systems must weigh these computational costs against the strategic benefits of interpretability and stakeholder trust. Techniques such as approximation methods, distributed computing, and model simplification can help in scaling these models for enterprise deployment.

Lastly, the practical implications of these findings are considerable. Financial institutions can adopt explainable time series models to improve transparency in automated credit scoring, investment advising, and risk prediction. Retail investors and analysts benefit from clearer decision rules and better communication of model outputs. Regulators gain tools to audit AI systems and ensure fairness. Thus, the integration of XAI in financial forecasting is not just a technical advancement—it's an enabler of trust, accountability, and informed decision-making in an increasingly automated financial ecosystem.

V. CONCLUSION

In an era marked by increasingly complex financial markets and data-driven decision-making, the need for models that are both highly accurate and transparent is more critical than ever. This study has explored the convergence of Explainable Artificial Intelligence (XAI) and time series forecasting to address the dual challenge of predictive performance and model interpretability in the financial domain. By employing models such as Long Short-Term Memory (LSTM) networks and Facebook Prophet alongside explainability tools like SHapley Additive exPlanations (SHAP), we demonstrate that it is possible to construct systems that deliver robust forecasts while offering meaningful insights into their internal logic. The empirical results from our experiments, applied to real-world datasets such as the S&P 500 index and EUR/USD exchange rates, confirm that LSTM models equipped with SHAP explanations can outperform traditional and rule-based models in accuracy, while also achieving significant progress in interpretability. Likewise, Prophet offers a viable alternative when user transparency is paramount, even if it involves a slight trade-off in precision. These findings underscore a key theme of this research: that predictive accuracy and explainability need not be mutually exclusive, and that when used in tandem, advanced AI models and XAI frameworks can foster trust, compliance, and informed decision-making in financial environments. Moreover, the integration of XAI into financial forecasting has broader implications for governance, risk assessment, and ethical AI deployment. In regulatory contexts where transparency is mandated such as under the European Union's General Data Protection Regulation (GDPR) or emerging frameworks on AI accountability explainable models offer a clear advantage. Institutions and investors are no longer content with opaque systems; they require tools that not only perform but also justify their predictions in a language comprehensible to humans. XAI fills this gap by turning AI models from black boxes into glass boxes, opening up the possibility of deeper stakeholder engagement and improved financial literacy. Nevertheless, this research acknowledges that the implementation of XAI methods still presents several challenges. There is an inherent trade-off between the complexity of models and the degree of interpretability achievable. While tools like SHAP are model-agnostic and powerful, they can be computationally intensive and may require expert understanding for accurate interpretation. Future innovations should focus on simplifying the interpretability process for non-technical users, enabling a broader range of stakeholders such as portfolio managers, compliance officers, and policy-makers to confidently interact with AI-driven systems. Looking ahead, this study lays the groundwork for several promising avenues of research. One potential direction is the development of hybrid models that combine multiple time series methods with real-time XAI dashboards, enhancing usability in high-frequency trading and automated risk assessment. Another opportunity lies in integrating causal inference techniques with XAI to explain not just how a model makes predictions, but why certain financial phenomena occur. Additionally, advancing research into human-AI collaboration in financial forecasting can help bridge the gap between automated systems and expert human judgment. In conclusion, the fusion of Explainable AI and time series forecasting represents a significant leap forward in the pursuit of intelligent, trustworthy, and user-friendly financial analytics. By balancing predictive precision with interpretative clarity, such models can empower decision-makers across the financial ecosystem to act with greater confidence, transparency, and accountability. The evolution of explainable forecasting systems will not only shape the future of financial modeling but also set new standards for ethical and responsible AI across domains.

WORKS CITED

- [1] Biran, Or, and Courtenay Cotton. "Explanation and Justification in Machine Learning: A Survey." Proceedings of the IJCAI-17 Workshop on Explainable AI (XAI), 2017, pp. 8–13. arxiv.org/abs/1702.08608
- [2] Doshi-Velez, Finale, and Been Kim. "Towards a Rigorous Science of Interpretable Machine Learning." arXiv, 28 Feb. 2017, arXiv:1702.08608. arxiv.org/abs/1702.08608.
- [3] Fischer, Thomas, and Christopher Krauss. "Deep Learning with Long Short-Term Memory Networks for Financial Market Predictions." European Journal of Operational Research, vol. 270, no. 2, 2018, pp. 654–669. Elsevier, <https://doi.org/10.1016/j.ejor.2017.11.054>.
- [4] Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. Deep Learning. MIT Press, 2016. Hochreiter, Sepp, and Jürgen Schmidhuber. "Long Short-Term Memory." Neural Computation, vol. 9, no. 8, 1997, pp. 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
- [5] Lundberg, Scott M., and Su-In Lee. "A Unified Approach to Interpreting Model Predictions." Advances in Neural Information Processing Systems, vol. 30, 2017, pp. 4765–4774. https://proceedings.neurips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf.
- [6] Makridakis, Spyros, Evangelos Spiliotis, and Vassilios Assimakopoulos. "Statistical and Machine Learning Forecasting Methods: Concerns and Ways Forward." PLOS ONE, vol. 13, no. 3, 2018, e0194889. <https://doi.org/10.1371/journal.pone.0194889>.
- [7] SRibeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier." Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 1135–1144. <https://doi.org/10.1145/2939672.2939778>.
- [8] Taylor, Sean J., and Benjamin Letham. "Forecasting at Scale." The American Statistician, vol. 72, no. 1, 2018, pp. 37–45. Taylor & Francis, <https://doi.org/10.1080/00031305.2017.1380080>.
- [9] Zhang, Guoqiang Peter, B. Eddy Patuwo, and Michael Y. Hu. "Forecasting with Artificial Neural Networks: The State of the Art." International Journal of Forecasting, vol. 14, no. 1, 1998, pp. 35–62. Elsevier, [https://doi.org/10.1016/S0169-2070\(97\)00044-7](https://doi.org/10.1016/S0169-2070(97)00044-7).



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)