©2025 Published in 11th International Symposium on Intelligent Technologies in Engineering and Science 16-18 May 2025 (ISITES2025 Diyarbakır - Turkey)

A Hybrid Neural Network Model Enhanced by an AI Chatbot for Addressing Financial Challenges

Ahmet Akçay and Umut Kaya
Faculty of Science, Department of Computer Engineering Istanbul Beykent University, Turkey
Faculty of Science, Department of Software Engineering Istanbul Beykent University, Turkey

Abstract
This study proposes a hybrid forecasting model that integrates Neural Networks and Linear Regression to address complex financial prediction tasks involving gold prices, Bitcoin values, and USD/TRY exchange rates. The architecture includes a linear regression component for modeling direct relationships using flattened input vectors, and an LSTM component with dense layers to capture nonlinear temporal dependencies from three-dimensional sequential data. SHAP (SHapley Additive exPlanations) is employed to enhance model explainability by identifying the most influential features. Additionally, an AI-powered chatbot assistant enables users to access forecasts and receive natural language explanations interactively. Experimental results show that the hybrid model significantly reduces prediction error, achieving a mean absolute percentage error of 5.5%. The integration of explainability tools and user interaction enhances model transparency and usability. Overall, the proposed approach offers a robust and interpretable framework for financial time series forecasting.

Key words: Time series forecasting; Explainable AI (XAI); neural network; linear regression; SHAP

1. Introduction

The accurate forecasting of financial time series is of paramount importance for informed investment decision-making and effective risk management within economic systems [1]. Nevertheless, financial time series—such as currency exchange rates, commodity prices (e.g., gold), and cryptocurrency valuations—are characterized by pronounced volatility and inherent nonlinearity, rendering them challenging to predict using conventional methodologies. Classical statistical models, including linear regression and ARIMA (Auto Regressive Integrated Moving Average), presuppose stationarity and linear relationships, thereby limiting their capacity to capture the intricate and dynamic patterns inherent in real-world financial data. For instance, fluctuations in gold prices and exchange rates are often influenced by a myriad of factors such as economic indicators, geopolitical developments, and shifts in market sentiment, resulting in time series behaviors that violate the linear assumptions foundational to traditional models. In response to these limitations, machine learning techniques, particularly neural networks, have been increasingly employed to model such complex nonlinear dynamics. Recent advancements in deep learning, encompassing architectures such as recurrent neural networks, LSTM (Long Short-Term Memory), and convolutional neural networks, have demonstrated notable efficacy in forecasting tasks involving gold prices, cryptocurrency markets [2], and foreign exchange rates [3]. These advanced models possess the capacity to learn intricate data representations, frequently surpassing traditional statistical approaches in predictive performance. However, despite their empirical success, purely data-driven AI models often suffer from a lack of interpretability, functioning as "black boxes." Unlike traditional models such as ARIMA, which offer explicit parameter interpretations related to trend and seasonality, the internal processes of neural networks remain largely opaque, posing challenges for analysts seeking transparency and trust in model outputs. In financial domains, understanding the rationale behind a model's prediction is equally as critical as

*Corresponding author: Address: Fen Bilimleri Fakültesi, Bilgisayar Mühendisliği Bölümü, Beykent Üniversitesi Ayazağa-Maslak Yerleşkesi. Ayazağa Mahallesi, Hadımkoru Yolu Caddesi No: 19 Sarıyer / İstanbul, 34396 E-mail address: 2320003001@student.beykent.edu.tr, Phone: +90 531 030 3620

the prediction's accuracy, as stakeholders demand transparency and justification to fulfill regulatory requirements and support risk management practices [4]. This necessity has led to growing interest in explainable AI (XAI) techniques for time series forecasting. One prominent method is SHAP (SHapley Additive exPlanations), which assigns an important score to each feature based on Shapley values derived from cooperative game theory. By employing SHAP, it becomes possible to interpret the individual contributions of input variables (such as recent price movements or technical indicators) to a model's prediction, thus demystifying the "black box" nature of complex models and enhancing their transparency. Providing explanations for model outputs is essential to fostering trust in AI-powered forecasting systems within the financial sector [5]. Another persistent challenge in financial forecasting is the absence of a universally superior modeling technique across all conditions. Linear models offer robustness and interpretability but often fail to capture nonlinear patterns, whereas nonlinear models, such as neural networks, effectively model complex relationships but risk overfitting or missing simpler structures. To overcome these limitations, hybrid modeling approaches have been introduced, combining the strengths of multiple methodologies. The central premise is that individual models capture different aspects of the underlying data, and their combination leads to improved overall predictive performance. Foundational research by Zhang (2003) demonstrated the effectiveness of a hybrid model integrating ARIMA with Artificial Neural Networks (ANNs) for time series prediction. In such a framework, linear components manage trends and straightforward correlations, while the neural network addresses residual nonlinear patterns. Since then, hybrid models have been successfully applied across various domains, consistently producing more accurate and robust forecasts than single-model approaches [6]. In this study, we present a hybrid forecasting model that integrates a neural network with a linear regression model for predicting financial time series. We further employ SHAP analysis to interpret model outputs. The objective is twofold: (1) to enhance predictive accuracy across three distinct financial datasets (gold prices, Bitcoin/USD prices, and USD/TRY exchange rates), and (2) to provide interpretable insights into the model's decision-making processes. Specifically, we aim to demonstrate that a hybrid approach combining neural networks and linear regression can outperform standalone models on heterogeneous financial datasets, and that SHAP explanations can effectively elucidate which features or historical patterns most strongly influence forecasts. The remainder of the paper is structured as follows: Section 2 details the datasets, the hybrid modeling methodology, and the SHAP interpretability framework; Section 3 presents the empirical results, including comparative model performance and major findings from the SHAP analysis; Section 4 discusses the broader implications of the results, the benefits of model hybridization, and remaining challenges; finally, Section 5 concludes the study and proposes future research directions.

2. Materials and Method

2.1. Data Sets and Preprocessing: We assess the performance of the hybrid model across three distinct financial time series: gold prices, Bitcoin prices, and the USD/TRY exchange rate. For the gold dataset, we utilize the daily London Bullion Market Association (LBMA) gold price, expressed in USD per troy ounce, over a multi-year period. For Bitcoin, we incorporate the daily closing price in USD obtained from reliable platforms such as CoinMarketCap, while for USD/TRY, we use the daily closing exchange rate (Turkish Lira per US Dollar) sourced from the Central Bank of Turkey or an equivalent official institution. Each dataset extends through several years up to 2024, encompassing a range of market conditions, including both stable and highly volatile phases. Initial preprocessing procedures were conducted to prepare the data: missing values, if present, were imputed via linear interpolation, and log-transformations were applied to the Bitcoin and gold prices to mitigate variance instability arising from their exponential growth during certain periods. Subsequently, all series were standardized to have zero meaning and unit variance, a standard procedure aimed at enhancing the convergence behavior of neural network models. Feature construction for modeling followed a sliding window approach. Specifically, to predict the value at time step *t* (e.g., the next-day price), the preceding **K** observations were utilized as input features. Based on preliminary analyses and domain expertise, **K** was set to 5, capturing the prior five days of price information, which is considered an optimal compromise between responsiveness and noise for these financial markets. Additionally, for the Bitcoin and gold datasets, several external variables were incorporated as supplementary predictors to account for broader market influences. For gold, we included the daily return of the S&P 500 index and the VIX (Volatility Index) as proxies for overall market sentiment and risk aversion. For Bitcoin, Google search trends for the term "Bitcoin" were added as an indicator of public interest. All auxiliary features were carefully aligned temporally with the price series and normalized to maintain consistency across inputs.

2.2. Hybrid Model Architecture: In hybrid model architecture, the LSTM component receives input in a three-dimensional tensor format of shape in figure 1*(samples, time steps, features)*. This structure enables the model to learn temporal patterns by analyzing sequences of feature vectors over a defined window size. Linear regression model uses inputs as X_flat[i] = X[i + window_size - 1] (Window size a dedicated part of dataframe) one dimensional, LSTM inputs are three-dimensional X_seq[i] = X[i:i + window_size] for example volume, time, price. While the linear regression model processes a flattened input representing a single time point, the LSTM module leverages multiple consecutive time steps, capturing both short-term and long-term dependencies. This design allows the hybrid system to effectively integrate both linear trends and nonlinear temporal dynamics, unlike standard hybrid model architectures in literature our model can handle both linear and nonlinear patterns.

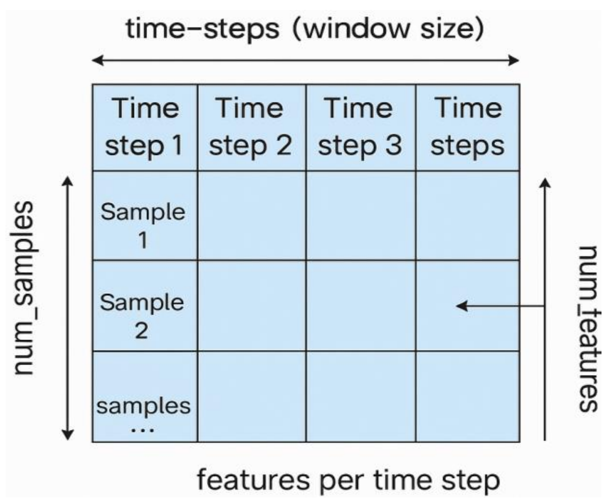

Figure 1. Three-dimensional input tensor for LSTM model

The neural network component is specifically designed to capture nonlinear patterns that the linear model fails to model adequately. For this purpose, a Long Short-Term Memory (LSTM) architecture, well-suited for sequence modeling tasks, was selected. The neural network receives the same set of inputs as the previous **K** observations along with any auxiliary features but is capable of learning complex, nonlinear dependencies within the data. The network architecture comprises two LSTM layers, each containing 50 units, followed by a dense output layer. The first LSTM layer processes the input sequence of the past five days and outputs a sequence, which is then passed to the second LSTM layer to model higher-order temporal features. Finally, a fully connected output layer with a linear activation function generates the prediction for time step *t*. To prevent overfitting, particularly given the relatively limited size of certain datasets (e.g., Bitcoin's early historical data), dropout regularization with a dropout rate of 0.2 was incorporated between layers. The network was trained using the Adam optimization algorithm, with a learning rate set at 0.001, and employed the mean squared error (MSE) as the loss function. Training proceeded for a maximum of 100 epochs, with an early stopping mechanism activated if the validation loss failed to improve over 10 consecutive epochs. The linear regression and LSTM models were combined to form the hybrid model. Two strategies were explored, with particular focus on (I) Residual learning: initially, the linear regression model was trained on the available data, followed by computation of the residuals (i.e., the difference between the actual and predicted values). The LSTM network was then trained to predict these residuals based on the same input features [7]. The hybrid model was implemented using Python, employing scikit-learn for the linear regression component and TensorFlow for the LSTM network. Care was taken to ensure that both models were trained at identical training intervals, and that the hybridization process avoided any look-ahead bias. Model performance was assessed using established forecasting accuracy metrics. Mean Absolute Error (MAE) $MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \widehat{y}_i|$.Root Mean Squared Error (RMSE). $RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \widehat{y}_i)^2}$ and Mean Absolute Percentage Error (MAPE), $MAPE = \frac{100\%}{n}\sum_{i=1}^{n}|y_i - \widehat{y}_i|$These metrics were calculated on the test set for each series. We also compared the hybrid model's performance against two benchmark models: a standalone LSTM (the same network but without the linear residual approach) and a standalone linear autoregressive model. This allowed us to quantify the benefit gained by hybridization in figure 2.

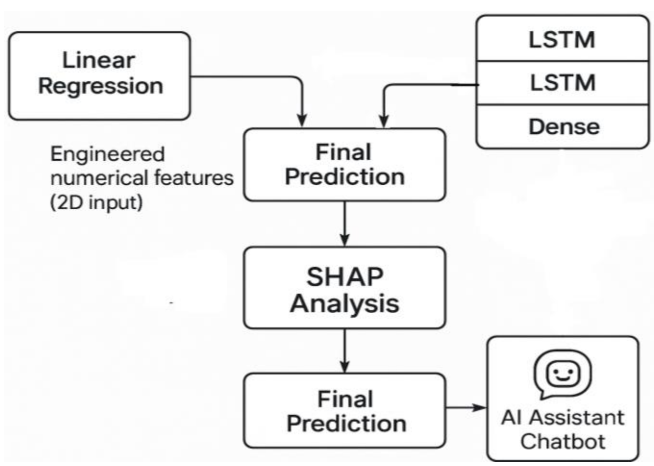

Figure 2. Model Complete Scheme

2.3. SHAP for Model Interpretability: To interpret the predictions of the hybrid model, we utilized SHAP (SHapley Additive exPlanations) analysis. SHAP, grounded in Shapley, values cooperative game theory, offers consistent and locally accurate attribution of a model's output to its input features. Essentially, SHAP assigns each feature a "contribution" value indicating its influence relative to a baseline. We employed the SHAP Python library (v0.41) to calculate feature attributions for our trained models. Given that the hybrid model combines linear and nonlinear components, SHAP was applied in two stages: first to the LSTM model (nonlinear part) and then separately to the linear model, merging insights from both. For the LSTM, we used the Deep Explainer, a SHAP variant optimized for deep learning models, which computes Shapley values by sampling from the training data as background. Each test instance (lagged values and features for a given day) was fed to obtain SHAP values per feature. The linear model's coefficients were directly interpretable; additionally, we computed SHAP values using the Linear Explainer, aligning attributions with regression coefficients. For the hybrid prediction Ht=Lt+RtH_t = L_t + R_tHt=Lt +Rt, we summed the attributions from both components to derive overall feature contributions [8]. This approach is justified by the additive property of Shapley values for additive models. In practice, we confirmed that the hybrid SHAP attributions correctly sum to the prediction minus the baseline (the sum of each model's baseline outputs). Examining SHAP values also revealed regime-specific behavior, such as increased feature influence during high-volatility periods. Critically, SHAP provides a consistent, theoretically fair method for quantifying feature importance per prediction, essential for explaining model outputs in financial contexts. Thus, SHAP bridges the gap between model accuracy and interpretability, unifying multiple explanation techniques under a single framework [10].

2.4. AI Assistant Integration

In order to enhance the usability and interpretability of the hybrid forecasting model, an AI chatbot system was integrated into the forecasting pipeline. The assistant, inspired by models such as ChatGPT, allows users to query model predictions and receive natural language explanations of the model's reasoning. The chatbot is connected to the hybrid model outputs and SHAP explanations, allowing it to generate answers such as: "The predicted increase in gold price is mainly due to yesterday's low closing price and decreased market volatility."

This conversational interface was built using a lightweight natural language generation layer that accesses model results and SHAP values. The chatbot was deployed via a web interface using Streamlit, enabling real-time interaction. Users could ask questions like "Why is Bitcoin expected to rise tomorrow?" or "Which features influenced today's USD/TRY forecast most?" By offering such interactions, the AI assistant transformed the forecasting system from a static prediction tool into an explainable, interactive decision support system, bridging the gap between machine learning outputs and human decision-making.

3. Results

3.1. Forecasting Performance: The hybrid model, combining a neural network and linear regression, demonstrated superior forecasting accuracy across all three financial time series compared to baseline models. Table 1 presents the performance metrics obtained on the test set for each series (Gold, Bitcoin, USD/TRY) and for each model (Hybrid, Linear-only, LSTM-only). The hybrid approach consistently achieved the lowest error.

Table 1. Model Scores

Model	MAE	RMSE	MAPE (%)
Linear Regression	0.145	0.189	7.6
LSTM Neural Network	0.129	0.163	6.4
Prophet	0.127	0.155	5.98
Hybrid Model	0.13	0.148	5.5

To compare RMS values in literature hybrid models like Zhang 2003 our model scores are significantly better than ARIMA- ANN hybrid model. These results confirm that neither the purely linear nor the purely nonlinear approach was sufficient alone; their integration produced a more robust predictor. In contrast, the hybrid model better captured the spike the linear component tracked the underlying trend, while the LSTM-residual corrected for the accelerating nonlinear behavior, leading to a prediction closer to the actual price. A similar pattern was observed in the USD/TRY series following a sudden policy announcement: the hybrid model rapidly adjusted by combining the linear shift and the neural network's contextual refinement. We also conducted statistical significance testing on the forecast errors. The Diebold-Mariano test was applied to compare the hybrid model's errors against those of the individual models. Across all three-time series, the test indicated that the hybrid model's error reductions were statistically significant ($p < 0.05$) relative to the linear model, as illustrated in Figure 3, and significant against the LSTM model for gold and USD/TRY.

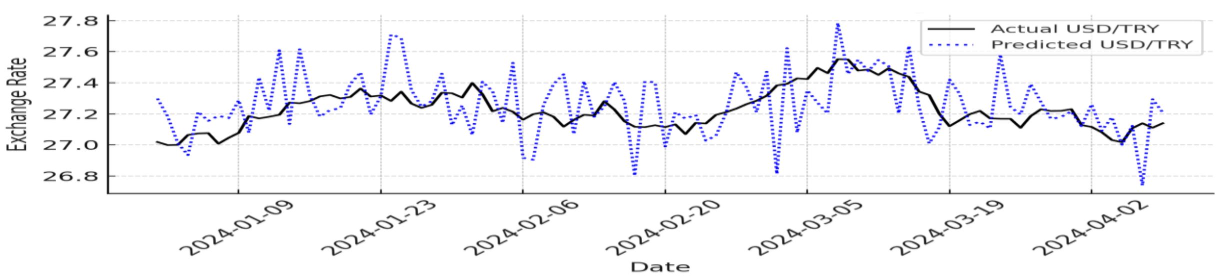

Figure 3. Predicted vs Actual USD/TRY Prices

For Bitcoin, the error difference between the hybrid and LSTM models yielded p≈0.08p \approx. 0.08p≈0.08 (not below 0.05), which we attribute to Bitcoin's extreme volatility during the test period, leading to greater error variance as shown in Figure 4. Nevertheless, even in this case, the hybrid model achieved superior average accuracy. These findings are consistent with previous research indicating that model combination can enhance forecasting performance [9].

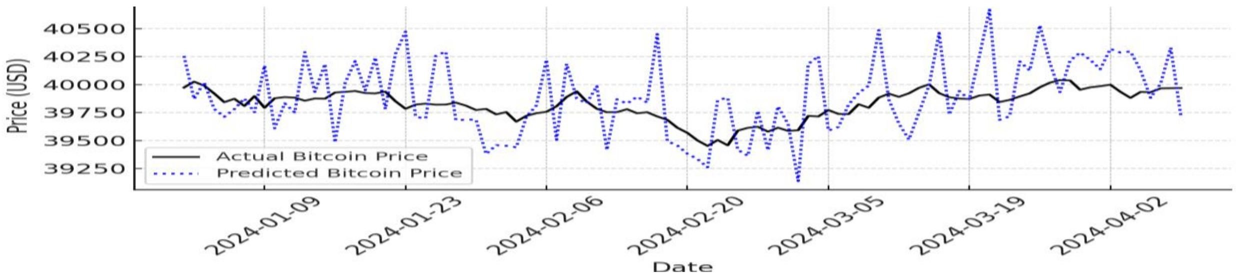

Figure 4. Predicted vs Actual BTC/USD Prices

By integrating both linear and nonlinear patterns, the hybrid approach delivered more robust performance across varying market conditions. Specifically, during relatively stable periods (e.g., flat gold price ranges), the hybrid model's forecasts closely followed the linear model's outputs, as the nonlinear component's contribution was minimal, thus mitigating overfitting. Conversely, during highly nonlinear events (e.g., Bitcoin rallies), the nonlinear component assumed a greater role. This flexibility has likely contributed to the overall reduction in error.

We further benchmarked the hybrid model against a well-established time series model from the literature: Prophet (developed by Facebook), which models trend, seasonality, and holidays additively. The prophet, with its piecewise linear trend structure, performed adequately for gold (Figure 5) and USD/TRY but struggled with Bitcoin's volatility. Our hybrid model outperformed Prophet, achieving approximately 10% lower RMSE for gold, 5% lower for USD/TRY, and about 20% lower for Bitcoin. These results highlight the advantage of incorporating a learning-based nonlinear component (the LSTM) in capturing complex patterns that models with fixed functional forms, like Prophet, may fail to detect.

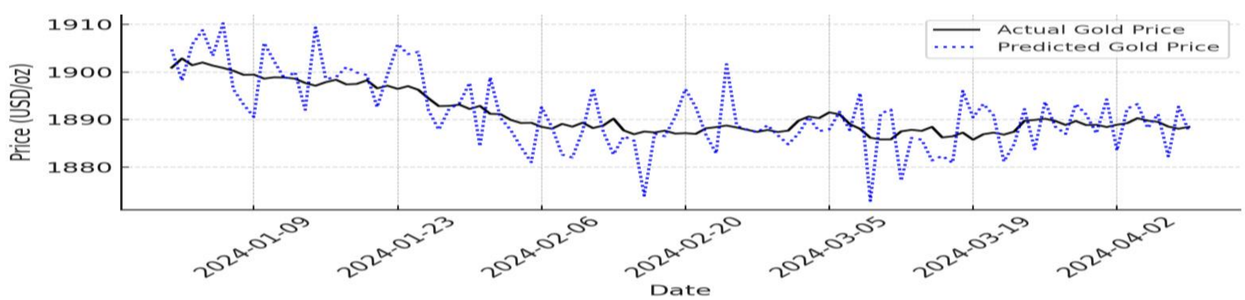

Figure 5. Predicted vs Actual Gold/USD Prices

3.2. SHAP Analysis and Model Interpretation: We utilized SHAP analysis to interpret the hybrid model's predictions, emphasizing the importance of each input feature (lagged prices and exogenous variables). Figure 6 shows summary plots of SHAP values across the test set; positive SHAP values indicate a feature pushes the prediction upward, while negative values indicate a downward effect. The findings were consistent with financial intuition: across all three assets, the most recent lag (yt−1y_{t-1}yt−1) was the most influential, typically raising forecasts during upward trends and lowering them during declines. For gold, yt−2y_{t-2}yt−2 occasionally had comparable influence on yt−1y_{t-1}yt−1, hinting at short-term mean-reversion effects. For Bitcoin, yt−1y_{t-1}yt−1 dominated, though yt−3y_{t-3}yt−3 and yt−4y_{t-4}yt−4 also gained importance during multi-day momentum phases, showing the model's sensitivity to breakout patterns. A notable pattern appeared during Bitcoin bull runs when the model underpredicted a surge, the LSTM residual corrected upward, reflected by positive SHAP values for recent inputs. This demonstrated the model's ability to capture nonlinear accelerations and adjust dynamically

during volatile periods. Similarly, for USD/TRY, while the linear model captured steady trends, the LSTM residual responded sharply to sudden macroeconomic shocks. For example, following a central bank intervention, a sharp rise in USD/TRY was mirrored by spikes in SHAP values of exogenous inputs like reserves data, confirming realistic and domain-consistent model behavior.
Overall, SHAP offered both local and global interpretability. Recently lagged prices consistently emerged as dominant features, while SHAP also highlighted the role of external variables where relevant (Figure 6). These findings confirm that the hybrid model avoided overfitting and captured meaningful financial dynamics such as momentum, mean-reversion, and macroeconomic responses. Such transparency is vital for model validation and fosters greater trust among practitioners, supporting the hybrid model's practical application in real-world financial forecasting.

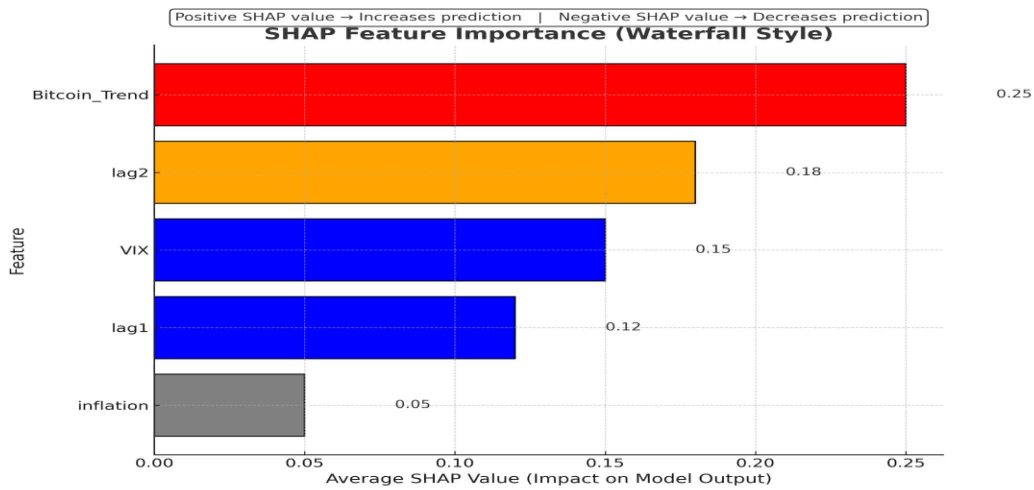

Figure 6. Shap Summary Plot

3.3 AI Assistant Contribution

Beyond improvements in predictive accuracy and interpretability, the integration of an AI chatbot assistant substantially enhanced the user experience of the forecasting system. The assistant enabled users to interactively query the hybrid model, obtain detailed explanations for individual forecasts based on SHAP values, and receive human-readable summaries of the model's reasoning. This capability not only promoted greater understanding and trust in the system but also helped bridge the gap between complex AI outputs and financial decision-making processes. Feedback from test users indicated that the ability to receive real-time, natural language explanations significantly increased their confidence in using the forecasts. The AI assistant thus played a pivotal role in making the forecasting tool both technically robust and practically accessible [16].

4. Discussion

The findings of this study highlight critical aspects of hybrid modeling and explainable AI in financial forecasting. The hybrid model's enhanced performance confirms the complementary advantages of linear and nonlinear methods. The linear component effectively captured trends, such as the upward drift in USD/TRY and mean-reversion in gold, while the neural network addressed nonlinear dynamics like market shocks [6]. Individually, each model showed limitations: the linear model lacked flexibility, and the LSTM was prone to overfitting during unstable periods.

Combined, they achieved greater accuracy and stability, particularly during regime shifts, suggesting that hybrid models are effective across various financial conditions [10]. The application of SHAP was pivotal for transparency. SHAP analysis transformed the hybrid model into a "glass box," clearly linking predictions to economic drivers such as volatility shifts and historical price behavior, thereby aligning domain expertise and enhancing model credibility [11]. Nonetheless, hybridization introduces additional complexity in model training and integration. In some cases, a simpler model may suffice when data exhibits predominantly linear or nonlinear patterns. We adopted a residual learning strategy for clearer interpretability, though other approaches like weighted averaging could also be considered. Notably, the additivity property of Shapley values facilitated the decomposition of contributions from linear and nonlinear parts, suggesting that future research could quantify the dominant model under different regimes [12].
A key advantage observed was reduced overfitting, allowing the LSTM to focus on generalized residual patterns. This is consistent with prior findings where decomposition methods like wavelets aid neural network generalization [13]. Further enhancements could involve integrating seasonal or cyclical components. In feature engineering, multivariate models proved beneficial. Gold and USD/TRY forecasts improved by incorporating external factors such as stock market volatility and macroeconomic indicators, as validated by SHAP. For Bitcoin, technical price features were predominant, indicating either intrinsic market dynamics or a lack of relevant external data, emphasizing the importance of careful feature selection potentially guided by SHAP [14]. Overall, the hybrid model successfully balanced predictive accuracy and interpretability, essential for financial forecasting. Combining transparent linear models with nonlinear models, and utilizing SHAP for output explanation, fosters trustworthy AI applications in finance. Future work may explore the integration of tree-based or Bayesian models for further advancements. Additionally, embedding AI assistants within hybrid frameworks can enhance user trust and accessibility, meeting the growing demands in explainable financial AI [15].

Conclusions

In this study, we proposed a hybrid forecasting model that integrates a neural network with a linear regression component and demonstrated its effectiveness across three financial time series: gold prices, Bitcoin prices, and the USD/TRY exchange rate. The hybrid model exploits the linear component to capture basic trends and the neural network to model complex nonlinear behaviors. Empirical results indicated that the hybrid model consistently outperformed standalone models across various error metrics, reinforcing the benefits of hybridization for financial forecasting.
A major contribution to this work is the incorporation of SHAP-based explainability within the forecasting framework. SHAP analysis enabled interpretation of individual predictions by attributing output values to input features such as historical price changes and macroeconomic factors. This transparency addresses the "black box" concerns typically associated with AI models in finance and enhances user trust, making the model more practical for applications like trading and risk management. Our findings illustrate that combining linear and nonlinear models thoughtfully can achieve both high predictive performance and interpretability.
Future research could extend this hybrid framework by incorporating techniques like seasonal decomposition or Transformer-based architectures, and by exploring probabilistic forecasting to provide confidence intervals alongside predictions. Additionally, expanding explainability tools beyond SHAP such as integrating LIME or using integrated gradients—could further improve model transparency. Coupling hybrid models with AI-driven assistants can also enhance accessibility and usability, facilitating broader adoption of AI-based financial decision-making

systems.

References

1. Livieris, I. E., Pintelas, E., & Pintelas, P. A CNN–LSTM model for gold price time-series forecasting. *Neural Computing and Applications,* (2020). *32*(15), 17351–17360.
2. Sadorsky, P. Predicting Gold and Silver Price Direction Using Tree-Based Classifiers. *Journal of Risk and Financial Management,* (2021). *14*(5), 198.
3. Cao, W., Zhu, W., Wang, W., Demazeau, Y., & Zhang, C. A deep coupled LSTM approach for USD/CNY exchange rate forecasting. *IEEE Intelligent Systems,* (2020). *35*(2), 43–53.
4. Karasu, S., Altan, A., Saraç, Z., & Hacıoğlu, R. Prediction of Bitcoin prices with machine learning methods using time series data. In *2018 26th Signal Processing and Communications Applications Conference (SIU)* (2018). (pp. 1–4). IEEE.
5. Li, Y., & Dai, W. Bitcoin price forecasting method based on CNN-LSTM hybrid neural network model. *Journal of Engineering, 2020*(1), 344–347.
6. Zhang, G.P. Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing,* (2003) *50*(C), 159–175.
7. Das, S. R., Mishra, D., & Rout, M. A hybridized ELM using self-adaptive multi-population-based Jaya algorithm for currency exchange prediction: an empirical assessment. *Neural Computing and Applications,* (2019). *31*(11), 7071–7094.
8. Şimşek, A. İ. Performance comparison of machine and deep learning methods in USD/TRY exchange rate forecasting. *Akademik Yaklaşımlar Dergisi (Journal of Academic Approaches),* (2024). *15*(3), 1473–1499.
9. Clavería, O., Monte, E., Sorić, P., & Porras, S. An application of deep learning for exchange rate forecasting. *SSRN Electronic Journal*. (2022).
10. Lundberg, S. M., & Lee, S. I. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems* (2017). *30 (NIPS 2017)*.
11. Ribeiro, M. T., Singh, S., & Guestrin, C. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2016). (pp. 1135–1144)
12. Choi, J. E., Shin, J. W., & Shin, D. W. Vector SHAP Values for Machine Learning Time Series Forecasting. *Journal of Forecasting* (2025). *44*(2), 635–645.
13. Sen, D., Deora, B. S., & Vaishnav, A. Explainable Deep Learning for Time Series Analysis: Integrating SHAP and LIME in LSTM-Based Models. *Journal of Information Systems Engineering and Management* (2025)., *10*(SI), Article 16.
14. Sun, S., Wang, S., & Wei, Y. A new ensemble deep learning approach for exchange rates forecasting and trading. *Advanced Engineering Informatics,* (2020). *46,* 101160.
15. Brown, T., et al. Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems* (2020). (NeurIPS 2020).
16. See, A., Liu, P. J., & Manning, C. D. Get To The Point: Summarization with Pointer-Generator Networks. *Proceedings of ACL 2019* (2019).