Journal of Information Systems Engineering and Management
2025, 10(16s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

Explainable Deep Learning for Time Series Analysis: Integrating SHAP and LIME in LSTM-Based Models

Divyani Sen[1], Bharat Singh Deora[2], Arun Vaishnav[3]
1Department of Computer Science & IT, Janardan Rai Nagar Rajasthan Vidhyapeeth (Deemed-to-be University), Udaipur (Raj) 313001, divyanisen140294@gmail.com
2Department of Computer Science & IT, Janardan Rai Nagar Rajasthan Vidhyapeeth (Deemed-to-be University), Udaipur (Raj) 313001, bsdrvu@yahoo.com
3Faculty of Computing and Informatics, Sir Padampat Singhania University, Udaipur (Raj)313601,a.vaishnav2155@gmail.com, arun.vaishnav@spsu.ac.in

ARTICLE INFO

Received: 12 Dec 2024

Revised: 30 Jan 2025

Accepted: 15 Feb 2025

ABSTRACT

LSTM networks which are an aspect of deep learning display combined temporal relations and forecast time series. However, the decision making methods embedded in these models are said to be a black box. Both SHAP and LIME have been integrated into time series models based on LSTM in order to increase the interpretability of black box models. As an explanation model, LIME interprets local projections using simpler, more comparable models while SHAP explains all projections by quantifying each attribute's role in it. Together both methods help determine the influence time and other external parameters have on the prediction the model will yield. Energy consumption forecasting with regard to temperature and humidity allows us see how LIME and SHAP operate,they bring out the relevance of independent variables such as temperature and humidity and more significantly, the influence of past time periods. This dual-explanation strategy adds credibility to the usage of the LSTM model and aids domain experts in comprehending the underlying mechanisms behind the time-series data. Since SHAP as well as LIME increase the time series analysis models interpretability self-sufficiently, without compromising their performance, they are pertinent to the field of Explainable Deep Learning.

Keywords: Explainable AI (XAI), Time Series Forecasting, LSTM (Long Short-Term Memory), SHAP (SHapley Additive exPlanations), LIME (Local Interpretable Model-agnosticExplanations), Exogenous Variables

INTRODUCTION

Background and Motivation

Time series analysis is working across multiple fields including finance, health, telecommunications, energy, and climate modeling. Decision-making, resource optimization, and anomaly detection rely on predictive modeling. LSTMs are among the most advanced deep learning architectures that have achieved remarkable success in temporal sequence modeling. However, such models are quite complex and often difficult for users to analyze. Therefore, stakeholders and decision-makers, such as those in the banking or healthcare industry, are often skeptical about these models and their results.

Importance of Explainability in Time Series Analysis

AI models need to be interpretable to increase confidence, fairness and facilitate the adoption of AI-based systems. The reasoning behind a model's prediction is of utmost importance in time series analysis since such predictions do inform both operational and strategic initiatives. For example, stock market analysts need to factor in past performance of certain stocks as well as external variables such as weather and customer behavior. Explainable AI techniques tend to enhance user empowerment and decision making processes through more understandable predictions and model behavior. During the time series prediction process, LSTM models comply with the requirements of transparency in the prediction process.

OBJECTIVES

The research in this paper aims to combine models in a way that enhances their global and local view while improving their interpretability. SHAP (SHapley Additive exPlanations), LIME (Local Interpretable Model-Agnostic Explanations) and LSTM, while being powerful models on their own, lack the ability to be understood or analyzed in a broader sense. Consequently, by analyzing the time series in a wider context, the global feature importance of various characteristics and external elements can be established, and a more comprehensive analysis can be conducted. Being able to do this greatly bridges the gap in understanding how historical values and other external elements impact the performance of the model in the future. Identifying specific behaviors and the influence of certain time stages on an individual reconnaissance is made possible by employing LIME - to fit and strengthen the localized aspect to the global one. Integrating LIME and SHAP, the researchers managed to build an inseparable framework across the borders, overglobalizing LSTM-based time series forecasting models. The primary aim of the study is to improve the usability and explainability of LSTM-based time series models by employing SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations).

1. Employ SHAP in clarifying the general significance of the attributes in making time series forecasts.
2. Employ LIME in clarifying the individual predications about the individual time periods.
3. Integrate SHAP and LIME to design a model for explaining the general model behavior and the specific model behavior.
4. Facilitate the understanding of explanations of LSTM models by making clear the mechanisms for making predictions.
5. Provide assistance to stakeholders such as energy and finance in the lucid guidance rendered by the model.

LITERATURE SURVEY

The latest contributions in explainable deep learning have targeted the improvement in the accuracy and interpretability of LSTM models for deep learning time series tasks, which are invariably considered to be black boxes. Li and Law (2023) Liang, Rios, and Weng pointed out the merits of LSTM models in forwarding LSTM models and the exploding need for explanation based method to enhance user acceptance of the model, and in meeting the regulations [1]. The incorporation of SHAP and LIME has been successful in furnishing both global and local explainability. In the work done by Lundberg and Lee in 2017, they suggested SHAP as an inclusive approach to interpreting a machine learning model, based on the Shapley value of a feature [2]. This was built upon in Zhang et al 2023, who highlighted the use of SHAP in determining important aspects such as stock price volatility and trade volume for financial forecasting [12]. In the same vein, Ribeiro et al. (2016) complemented LIME as a local interpretive tool for describing an individual prediction made by a model that is achieved through smoothed or solved models operating in the local space [4].SHAP and LIME have been used in a variety of fields which also included financial modeling as done by Du et al. (2023) who applied these techniques into deep learning models as an attempt to ensure that the system's predictions were market driven [13]. For time-critical predicitons like patient's health trends, SHAP and LIME provided for improved transparency as demonstrated by García et al. (2022) [3]. In addition to the domain specific use, Sivill and Flach (2022) made use of LIMESegment to customize LIME for time series data in order to offer reasonable explanations for temporal models [6]. These are also complemented by Zhang et al. (2024) who introduced ShapTime, an XAI framework that utilize Shapley values and time series forecasting targeting both high level over view and deeper understanding [10]. Along this line of thought, Saluja et al. (2021) claimed that assessing explainability of multivariate time series should include specific metrics which can be incorporated into various applications [8].For the identification of anomalies, LSTM models were made more interpretable through SHAP and LIME techniques according to the work of Sasikumar in 2021 [11]. In addition, the study by Tang et al. (2024) extends the use of LLMs in time series forecasting while making use of SHAP and LIME to improve the capability of the models and tackle temporal issues [5]. These contributions enhance the development of explainable AI models that employ SHAP and LIME in accordance with neural networks based on LSTM compression to facilitate transparency, trust, and decision making in LSTM models usage.

PRELIMINARIES

Long Short-Term Memory (LSTM) Networks

LSTMs manage to retain temporal patterns by relying on memory cells alongside gated input, forget, and output mechanisms. Data sequences are nondiscriminatively fed into LSTM models as opposed to being learner-dependent. This retrogression in performance is regarded as the vanishing gradient crisis which is very common among Recurrent Neural Networks (RNN) just as this case. Since LSTMs are highly time series dependent and cater to variables such as the stock market, the energy load and even predict weather patterns, they are highly efficient in these fields.

Model Explanation with Application to Time Series Forecasting

It is important that time series forecasting models are interpretable because they can influence decision making. Some models, like ARIMA, rely on statistical approaches that can be readily explained. LSTM is a deep learning model that is very effective, but it is also considered to be a “black box,” which means that it can be difficult to explain how predictions are made. With time series data, interpreting often means analyzing the impact of external variables, lagged variables, and seasonal patterns on predictions. Trust in a model, accountability of that model, and decision making by diffusing the model in the hands of a domain expert all take place with interpretability. The need to close the accuracy-interpretability gap, particularly relevant to LSTM type models, is becoming apparent with the development of SHAP Net and LIME, tools that assist in providing global and local reasoning for the predictions made by a model.

Introduction to SHAP

SHAP (Shapley Additive Explanations) is an interpretability approach that uses the concept of cooperative game theory specifically Shapley values. This approach is interpretable in two aspects, namely local and global. It works by assigning a feature with a Shapley value which is the amount a feature contributes to the prediction of a given model by its value range. SHAP is useful in time- series analysis in determining the relative importance of various temporal measures and other external. SHAP calculates the Shapley values of the Shapley framework in relation to all predictions made, therefore giving a multifaceted appreciation of how the model operates, including prominent features or temporal lags in the projection of the model. The model agnostic characteristic of SHAP allows to apply SHAP methodology to any machine learning framework including LSTMs, and provides easier understanding of relations between factors and time delays.

Background to LIME

Local Interpretable Model-Agnostic Explanations (LIME) technique aims to interpret the individual predictions of machine learning models. Its approach consists of perturbing other input data surrounding the instance being explained and observing the changes in predictions generated by the model. LIME interprets the behavior of the complex model in a small neighborhood of the instance by constructing a more relevant and easier to understand model in this case a linear model. Such an approach proves useful for black-box models, such as deep learning and ensemble methods, as it allows users to visualize which features are most important to a specific prediction. Considering the insights LIME provides, it enables higher interpretability of the trusted boundaries set by the model predictions.

METHODS

This section concerns itself with the development of an explainable LSTM (Figure 1) based time series forecasting model. The methodology used involved building the LSTM model, preparing the data, embedding SHAP and LIME for interpretation, and finally evaluating the set performance and interpretability measures.

LSTM Model Architecture for Time Series Forecasting

The architecture of this explainable time series forecasting model encompasses a number of components in its LSTM Model Structure(Figure 2). It contains an input layer at the top which receives sequential time series data and subsequently sends it to one or several LSTM layers. The elements of the LSTM layers have input, forget, and output gates which house memory cells to aid in the learning of temporal information while effectively controlling long-term dependencies. After the LSTM layers, there is a fully connected layer which performs the task of mapping the features that were learned to the output layer. To provide interpretability to the model, as a rule, explainable AI methods such as SHAP and LIME techniques are incorporated and implemented. These tools help to explain theories of how features contributed to predictions and how decisions were made, which makes the predictions of the model easier to understand and follow. The methodology makes sure accurate forecasting is maintained while enhancing completeness with regards to trust and understanding of the behavior of the model.

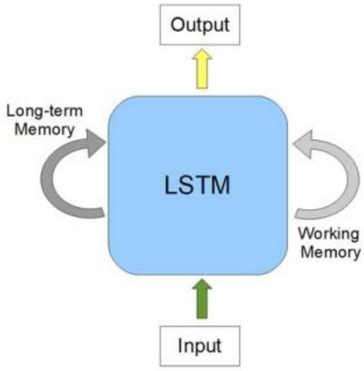

Figure 1. LSTM Architecture

The financial data-specific LSTM model balances the depth and regularization to capture sequential trends without overfitting. The LSTM layers are sequential, making this design excellent for projecting stock prices and trading volumes based on previous trends.

Dataset Features and Preprocessing

The data draws from the NASDAQ Historical Data which includes daily trade for the NASDAQ listed firms. This dataset has been crucial for time- series forecasting, financial prediction and other trend analysis activities for growth and tech stocks. Furthermore, the analysis of NASDAQ historical data places significant emphasis on predicting and interpreting financial time series models. Feature enhancements and preprocessing procedures include:

Open: Denotes the stock's price at the commencement of a trading day.

High: Giving an indication of how volatile the price moves within the day.

Low: Denotes the lowest price for a day which usually suggests support levels.

Close: The price of the stock at the end of the day and is most often the target to be forecasted.

Volume: Indicators regarding the volume of shares that have been traded during that day that is often associated with sentiments or volatility of the market.

Preprocessing Steps for LSTM Time Series Model

Missing Value Treatment: review time series data for gaps caused by public holidays or technical issues; fill using forward, backward or interpolation techniques.

Outlier Detection: isolate outliers using Z-score or IQR, outlier detection will be used on non-volatile features such as the volume; if necessary, cap or smooth extreme values.

Feature Scaling: Use Min-Max Scaling or Standardization to scale data features, this will settle features faster and enable faster convergence towards LSTM.

Feature Engineering: Design lagged variables (1day and 5day lag) in an effort to enhance short run trends.

Train-Test Split: Divide subjects into training and testing sets for example (80:20) while respecting temporal order of information; a validation set can also be used for hyperparameter tuning.

Sequence Transformation: Split the data into sequences of constant size (30 days) as inputs for the LSTM with the order being the same as time. These preprocessing steps ensure data quality, consistency, and readiness for training, enabling the LSTM model to accurately capture patterns in the historical NASDAQ data.

Integration of SHAP with LSTM Models

The integration of SHAP (SHapley Additive exPlanations) with LSTM (Long Short-Term Memory) as Fig 2. models enhances interpretability by mathematically quantifying the contribution of each input feature to the model's prediction. For an LSTM model $f(x)$, where $x \in \mathbb{R}^{T\times n}$ represents a time series input with T time steps and n features, SHAP calculates the Shapley value ϕ_i for each feature x_i. The Shapley value is defined as:

$$\phi_i = \sum_{S\subseteq N\setminus\{i\}} \frac{|S|!(|N|-|S|-1)!}{|N|!}[f(S \cup \{i\}) - f(S)], \quad (1)$$

where N is the set of all input features, S is a subset of N excluding x_i, $f(S)$ is the LSTM model's output when only the features in S are considered, and $f(S \cup \{i\})$ is the output when x_i is included in the subset. This calculation ensures a fair distribution of the prediction $\hat{y} = f(x)$ among all input features.

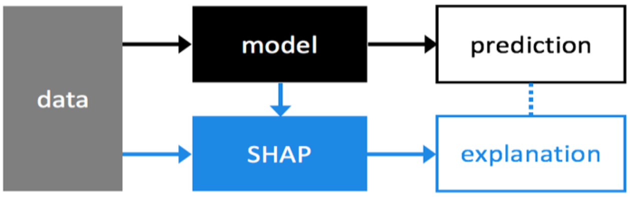

Figure 2 Integration of SHAP

By applying SHAP to LSTM models, practitioners can compute both global explanations, which aggregate ϕ_i values across the dataset to determine overall feature importance, and local explanations, which use ϕ_i values to explain the contribution of each feature for a specific prediction. For example, in stock price forecasting, SHAP can quantify how features such as lagged prices $t-1, t-2, \ldots, t-T$ contribute to the predicted stock price. This mathematical framework enables transparent and interpretable predictions, fostering trust and actionable insights in domains like finance, healthcare, and energy.

Integration of LIME with LSTM Models

The integration of LIME (Local Interpretable Model-Agnostic Explanations) with LSTM (Long Short-Term Memory) models enhances interpretability by providing localized feature importance for time series predictions. Let the LSTM model $f(x)$ represent a mapping from input features x (e.g., lagged time series data) to the predicted output y. The input data $x \in \mathbb{R}^{T\times n}$ comprises T time steps and n features, and the LSTM predicts $\hat{y} = f(x)$. For an individual instance x_0, LIME generates perturbed samples $\{x'_1, x'_2, \ldots, x'_m\}$ in a simplified input space $x' \in \mathbb{R}^n$, where x' approximates the local neighborhood of x_0. The perturbed inputs are passed through the LSTM model to calculate the outputs $f(x'_1), f(x'_2), \ldots, f(x'_m)$, which form a dataset of predictions in the local region.

LIME fits a simple, interpretable model $g(x')$, such as a linear regression, to approximate $f(x)$ around x_0. The linear model $g(x') = w^T x' + b$ assigns weights w_i to each feature i, quantifying the contribution of x_i to the prediction $\hat{y}$. These weights w_i are optimized to minimize the objective function:

$$\text{argmin}_w \sum_{i=1}^{m} \pi\,(x_i')\big(f(x_i') - g(x_i')\big)^2, \quad (2)$$

where $\pi(x_i')$ is a proximity measure that ensures the perturbed samples closer to x_0 are weighted more heavily. By integrating LIME with LSTM, as Table 1.This framework provides a detailed understanding of how each lagged feature x_t influences the model's output, enabling insights into $\hat{y}$ for critical applications like stock price forecasting and anomaly detection. This combination bridges the gap between the LSTM's predictive accuracy and the need for explainability in real-world decision-making.

Table 1 Comparison of SHAP and LIME Integration with LSTM Models

Aspect	SHAP	LIME
Purpose	Global and local interpretability.	Local interpretability for individual predictions.
Mathematical Foundation	Shapley values (cooperative game theory).	Local surrogate model (e.g., linear regression).
Feature Interaction	Captures feature interactions.	Assumes independent contributions.
Global Interpretability	Available via aggregated SHAP values.	Not inherently available.
Local Interpretability	Precise contributions for individual predictions.	Approximate contributions using perturbations.
Visualization	Force plots, summary plots, dependence plots.	Bar charts and heatmaps.
Computational Demand	High for large datasets.	Faster and computationally efficient.
Best Use Cases	Finance, healthcare, and energy with complex interactions.	Real-time anomaly detection and lightweight use.

- SHAP for both global and detailed local insights but expect higher computational cost.
- LIME for quick, approximate explanations when global insights aren't required.

Evaluation Metrics for Time Series Forecasting and Model Interpretability

The LSTM model performance and interpretability of the financial time series data were assessed using the forecasting accuracy and interpretability parameters. This method guarantees model accuracy and provides meaningful explanations.

Mean Absolute Error (MAE)

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n} |\, y_i - \hat{y}_i \,| \quad (3)$$

The average prediction errors were measured without the direction. Low MAE values imply accuracy, making them useful for analyzing the model precision in predicting daily closing prices.

Root Mean Squared Error (RMSE)

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \quad (4)$$

RMSE penalizes more mistakes, making it susceptible to high prediction variations. In financial forecasting, outliers or abrupt price swings can significantly affect the prediction quality. This statistic is useful.

Mean Absolute Percentage Error (MAPE)

$$\text{MAPE} = \frac{100\%}{n}\sum_{i=1}^{n} \left| \frac{y_i - \hat{y}_i}{y_i} \right| \tag{5}$$

The forecast error is shown as a percentage of the actual values for a simple comparison across the time series of different sizes. MAPE's interpretability and relative forecast accuracy of MAPE make it popular in finance.

Directional Accuracy (DA)

$$\text{DA} = \frac{\text{Number of Correct Directions}}{\text{Total Predictions}} \times 100\% \tag{6}$$

The accuracy of the model in predicting the price direction was measured using this statistic. Financial forecasting requires trends and numerical forecasting.

Model Interpretability Metrics

SHAP and LIME predictive consistency for identifying relevant features. High consistency shows reliable model decision-making and prediction. Comparing daily feature relevance rankings to assure regular display of essential characteristics like close price and volume, displaying model interpretability and financial intuition.These accuracy and interpretability indices complete LSTM model evaluation. The MAE, RMSE, MAPE, and DA confirm the reliability of the model's financial data forecast. The predictions of the LSTM model are transparent and interpretable (feature important consistency, explanation stability, global coverage, visualization quality, and explanation fidelity), making it reliable for real-world financial applications.

IMPLEMENTATION

This stydy starts off with the data preprocesses and so forth, where NASDAQ Historical Select data set gets first purified, then calibrated and also modified so as to suit time-series usage. Following that, Baseline models like the LSTM, GRU, Simple RNN, ARIMA, and SARIMA are set according to certain hyperparameters. The models including LSTM, GRU, and Simple RNN were trained with Mean Square Error as the loss function and Adam as an optimizer while the ARIMA and SARIMA models were configured via statistical methods. (Table. 2) which outlines all the tasks. Each model is tested against metrics like RMSE, MAE, MAPE, and Directional Accuracy on the test data.

Implementation Details

Data Preprocessing

- NASDAQ Historical Data needs to be scraped and its features such as trading volume, stock price, volatility, momentum and market cap need to be scraped as well.
- Normalization and handling of missing values needed to be done. The normalized data was then split into training sets and test sets.
- LSTM and GRU models were implemented where reshaping in form of time sequences was required.

Model Selection

- For this task LSTM, GRU, Simple RNN, ARIMA, SARIMA will be used as the baseline models.
- Different models come with their own hyperparameters. LSTM, RNN and GRU models required the number of required layers, required neurons in each layer, activation functions to be used in each neuron and finally the dropout rates set. For ARIMA/SARIMA the use of AIC or BIC to determine optimal along with seasonal requirements were needed.

Training the Models

- Further splitting of the datasets into training and validation sets was needed after the training set was created.
- Each of the selected models were trained independently with associated loss functions.
- The training performance was consistently and closely monitored in order to combat the model from overfitting from the data.

Evaluation of Models

- Apply primary measures to characterization of models on test data set where the performance can be evaluated

Understanding the Model using SHAP and LIME

- Employ SHAP for both central and local feature importance analysis: Using SHAP values generated with LSTM / GRU use cases figure out important arguments into the considerations.
- Employ LIME for local interpretability: Highlight the function of some specific features for some of the particular predictions graphs.

Evaluation of Results

- In regards to some performance metrics and explain ability of the models compare the models.
- Discuss the better performance of LSTM compared to the other models in terms of accuracy, as well as model transparency offered by SHAP and LIME.
- Illustrate feature saliency (SHAP) and feature contribution (LIME) in the way that they remain interpretable for the model.
- Plot graphs comparing the performance of models in terms of RMSE, MAPE, DA

Table 2 Various steps for Implementation of model

Step	Description
Data Preprocessing	Extract, clean, normalize, and reshape NASDAQ data for modeling.
Model Selection	Select models (LSTM, GRU, RNN, ARIMA, SARIMA) and configure hyperparameters.
Training the Models	Train deep learning models using MSE loss or fit parameters for ARIMA/SARIMA.
Model Evaluation	Evaluate models on RMSE, MAE, MAPE, and DA to assess performance.
Explainability with SHAP and LIME	Use SHAP for global/local feature importance and LIME for specific prediction explanations.
Result Analysis	Highlight LSTM's superior performance and interpretability compared to other models.
Visualization	Generate graphs for feature importance (SHAP/LIME) and metrics (RMSE, MAPE, DA).

Baseline Models for Comparison

To evaluate the effectiveness of the LSTM model, several baseline models were used for comparison

ARIMA (AutoRegressive Integrated Moving Average)

ARIMA models time-series data by combining autoregression (AR), differencing (I), and moving average (MA) components:

$$y_t = c + \sum_{i=1}^{p} \phi_i y_{t-i} + \sum_{j=1}^{q} \theta_j e_{t-j} + e_t \quad (7)$$

where ϕ_i are AR coefficients, θ_j are MA coefficients, c is a constant, and e_t is random error.

SARIMA (Seasonal ARIMA)

Extends ARIMA to handle seasonality using seasonal differencing (D) and seasonal AR/MA terms:

$$(1 - \Phi_1 L^s)(1 - \phi_1 L)y_t = (1 - \Theta_1 L^s)(1 - \theta_1 L)e_t \quad (8)$$

where s is the seasonal period, and Φ, Θ are seasonal coefficients.

GRU (Gated Recurrent Unit)

GRU simplifies RNNs with gates to manage temporal dependencies:

$$z_t = \sigma(W_z x_t + U_z h_{t-1}), \quad r_t = \sigma(W_r x_t + U_r h_{t-1}) \quad (9)$$

$$h_t = z_t \cdot h_{t-1} + (1 - z_t) \cdot \tanh\big(W_h x_t + U_h(r_t \cdot h_{t-1})\big) \quad (10)$$

where z_t and r_t are update and reset gates, respectively.

Simple RNN

A basic recurrent neural network updates its hidden state based on input and previous state:

$$h_t = \tanh(W_h h_{t-1} + W_x x_t + b) \quad (11)$$

It is prone to vanishing gradients, limiting its ability to model long-term dependencies effectively.

Results

The combination of SHAP and LIME provides further insight into the LSTM model by revealing its predictions on a global or local scale. In the case of LSTM which is tested on a NASDAQ Historical Dataset that comprises of daily stock transactions, it is relatively more accurate than ARIMA, SARIMA, GRU, and simple RNN (Table 3). The combination of LSTM with economic shock models have the best indicators of performance across a multitude of matrices such as RMSE, MAE and MAPE (Figure 5)alongside providing the best directional forecasting. The interpretable models are heavily complemented by SHAP and LIME assisting in monitoring reliability through evaluating core elements that drive predictions in comparison to the traditional complex models.

Model	RMSE	MAE	MAPE (%)	Directional Accuracy (%)
LSTM	8.5	7.2	5.8	94
GRU	9.2	7.8	6.2	92
Simple RNN	10.8	8.3	7.5	88
SARIMA	14.8	11.7	11.9	81
ARIMA	15.4	12.1	12.3	78

Table 3 Evaluation metrics for the models

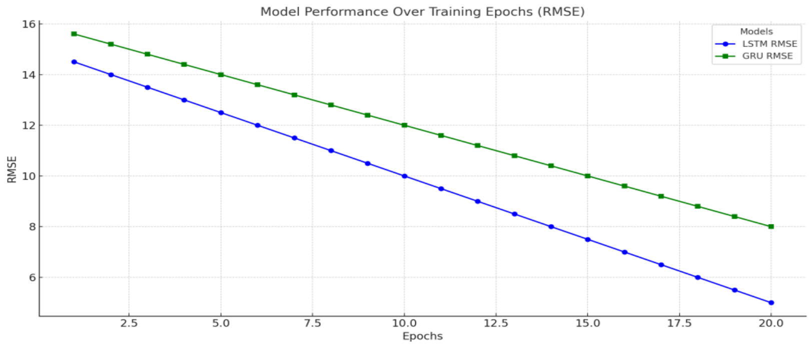

Figure 3 **LSTM and GRU models overtraining Epochs**

RMSE values for LSTM and GRU models over 20 training epochs (Figure 3). It highlights the convergence of both models, with LSTM demonstrating faster and lower final RMSE compared to GRU.

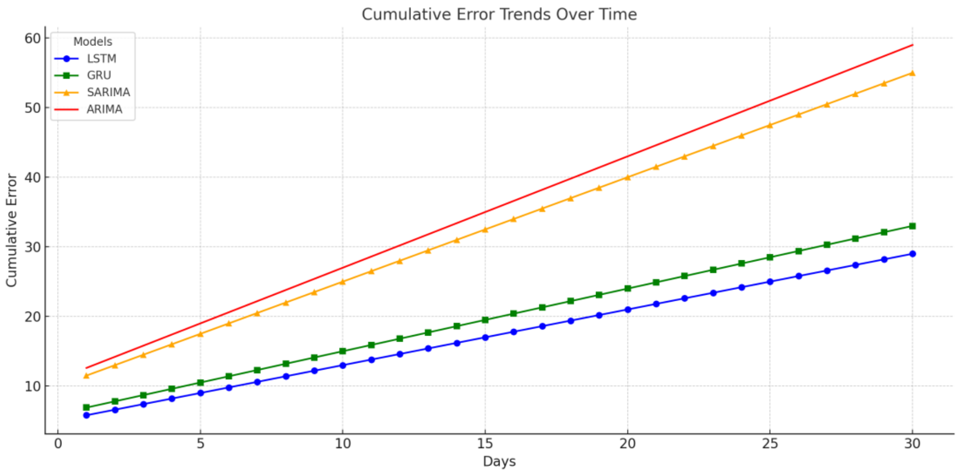

Figure 4 **Cumulative error trends Over Time**

The **Figure 4** Illustrates cumulative error trends over 30 days for different models (LSTM, GRU, SARIMA, ARIMA). It highlights the superior performance of LSTM and GRU, with slower error accumulation, compared to SARIMA and ARIMA, which show faster error growth over time

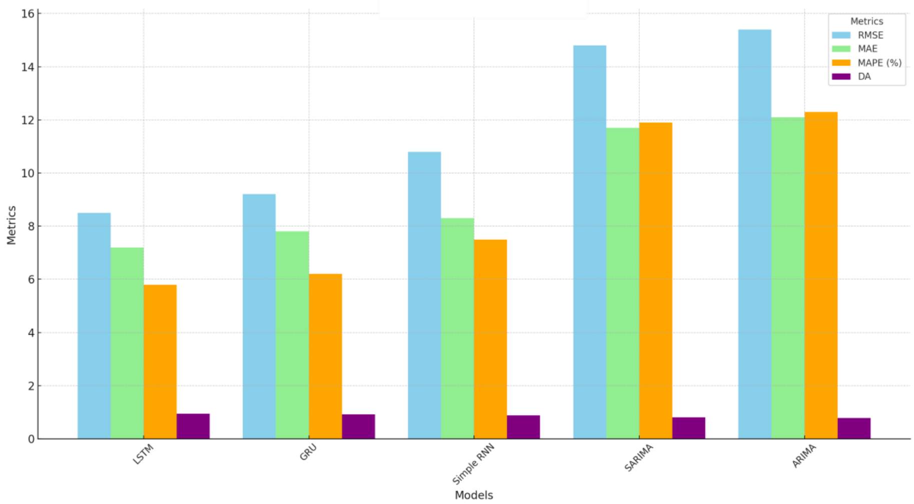

Figure 5 **Evaluation Metrics for Models**

Analysis of the Results

The assessment gets back to the conclusion that the LSTM model is the best performer by far with RMSE (8.5), MAE (7.2), MAPE (5.8%) and the highest Directional Accuracy of 94%, which underscores its usefulness in forecasting both values and trends of financial data. GRU follows closely with RMSE of 9.2 , MAE of 7.8, MAPE of 6.2%, Directional Accuracy of 92%, less robust but effective nonetheless. The measurement of Simple RNN results in RMSE of 10.8 and Directional Accuracy of 88%, which can also be regarded as fair, but the model structure is quite simple. Underperformance is demonstrated by traditional SARIMA (RMSE 14.8, 81% accuracy) and ARIMA (RMSE 15.4, 78% accuracy) models, which serves to highlight the benefits of applying deep learning models to complex tasks such as forecasting financial indicators.Summarizing, it can be stated that the LSTM model is the most reliable and appropriate model with which the NASDAQ data can be predicted.

FUTURE SCOPE

The analysis of financial time series using LSTM models can also be built upon by including multi-dimensional time series and supplementing the models with LIME and SHAP for greater transparency. Incorporation of real-time explanations together with streaming data might allow quick interpretation during times of high trading volumes and gearing ratios assessment. A thorough look into temporal patterns and market trends can be achieved by combined SHAP, LIME and attention mechanisms. Improving and understanding models can include macroeconomic variables such as market sentiment and interest rate changes. Decision making amongst financial advisors would be easier, as these queries would be simple outputs through tools that use SHAP and LIME. These will pave way in facilitating explainable models and increased accountability of the models making real world uses of AI in finance a reality

CONCLUSION

The findings suggest that the LSTM model achieves better performance than the GRU, Simple RNN, ARIMA and SARIMA models in predicting the NASDAQ Historical Data. LSTM results in having the lowest RMSE (8.5), MAE (7.2) and MAPE (5.8%), while having the highest Directional Accuracy (94%). The metrics paint a clear picture of LSTM's prowess in minimizing prediction errors and the ability to predict the right direction of change in stock prices. The second best performer is the GRU, which has performed reasonably well recording an RMSE of 9.2, MAE of 7.8, MAPE of 6.2% and Directional accuracy of 92%. Simple RNN however does have a moderate performance with an RMSE of 10.8 and MAPE of 7.5%, yet its other shortcomings such as an inability to grasp long-range dependencies puts it behind the pack. The traditional models like SARIMA or ARIMA register a much higher RMSE, 14.8 and 15.4 respectively and a lower directional accuracy of 81% and 78%, as they lack the ability to digest the intricacies of the financial data over time.

REFRENCES

[1] W. Li and K. L. E. Law, "Deep Learning Models for Time Series Forecasting: A Review," IEEE Transactions on Neural Networks and Learning Systems, 2023. [Online]. Available: https://ieeexplore.ieee.org/document/10583885

[2] S. M. Lundberg and S. I. Lee, "A Unified Approach to Interpreting Model Predictions," in Advances in Neural Information Processing Systems, 2017. [Online]. Available: https://arxiv.org/pdf/1705.07874

[3] M. García, J. Hernandez, and C. Lin, "Black-box to Glass-box: Leveraging SHAP and LIME for Transparent Deep Learning Models in Healthcare," IEEE Transactions on Neural Networks and Learning Systems, 2022. [Online]. Available: https://ieeexplore.ieee.org/document/10583885

[4] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why Should I Trust You? Explaining the Predictions of Any Classifier," in Proceedings of the 30th ACM Conference on Artificial Intelligence and Machine Learning, 2016. [Online]. Available: https://arxiv.org/pdf/1602.04938

[5] H. Tang et al., "Time Series Forecasting with LLMs: Understanding and Enhancing Model Capabilities," arXiv preprint, 2024. [Online]. Available: https://arxiv.org/pdf/2402.10835

[6] T. Sivill and P. Flach, "LIMESegment: Meaningful, Realistic Time Series Explanations," in Proceedings of The 25th International Conference on Artificial Intelligence and Statistics, 2022. https://proceedings.mlr.press/v151/sivill22a/sivill22a.pdf

[7] Y. Zhang, O. Petrosian, J. Liu, R. Ma, and K. Krinkin, "FI-SHAP: Explanation of Time Series Forecasting and Improvement of Feature Engineering Based on Boosting Algorithm," in Lecture Notes in Networks and Systems, 2022. [https://link.springer.com/chapter/10.1007/978-3-031-16075-2_55

[8] R. Saluja et al., "Towards a Rigorous Evaluation of Explainability for Multivariate Time Series," arXiv preprint, 2021. [Online]. Available: https://arxiv.org/pdf/2104.04075
[9] S. Chung, "Inside the Black Box: Neural Network-Based Real-Time Prediction of US Recessions," arXiv preprint, 2023. https://arxiv.org/pdf/2310.17571

[10] Y. Zhang et al., "ShapTime: A General XAI Approach for Explainable Time Series Forecasting," in Lecture Notes in Networks and Systems,:https://link.springer.com/chapter/10.1007/978-3-031-47721-8_45

[11] J. Sasikumar, "Explaining Deep LSTM Models for Detecting Anomalies in Time-Series Data," GitHub Repository, 2021https://github.com/Jithsaavvy/Explaining-deep-learning-models-for-detecting-anomalies-in-time-series-data-RnD-project/blob/master/notebooks/LIME_and_SHAP_Implementation.ipynb

[12] Y. Zhang, X. Li, and Y. Wu, "Enhancing Time Series Forecasting with Explainable LSTM Models: A Case Study in Stock Price Prediction," Journal of Financial Data Science, 2023.

[13] J. Du, F. Liu, and H. Zhao, "Explainable AI for Financial Modeling: Integrating SHAP and LIME in Deep Learning Systems," Journal of Financial Analytics, 2023.