Journal of Data Analysis and Information Processing, 2026, 14(2), 215-246
https://www.scirp.org/journal/jdaip
ISSN Online: 2327-7203
ISSN Print: 2327-7211

Predicting Stock Price Movements Using Deep Learning Techniques

Dennis Kibet Kirui

Department of Mathematics and Statistics, Austin Peay State University, Clarksville, USA
Email: dkirui@students.apsu.edu, kiruidennis008@gmail.com

How to cite this paper: Kirui, D.K. (2026) Predicting Stock Price Movements Using Deep Learning Techniques. *Journal of Data Analysis and Information Processing*, **14**, 215-246.
https://doi.org/10.4236/jdaip.2026.142011

Received: May 23, 2026
Accepted: April 6, 2026
Published: April 9, 2026

Abstract

This project aimed to assess the performance of attention extended LSTM and Gated Recurrent Unit models in stock price movement forecasting. The traditional models suffer from the challenges due to noisiness of financial markets, leading to incomplete specification of time dependencies in financial time series. These limitations are addressed by integrating attention mechanisms in LSTM and GRU models to allow the models to focus on specific critical time steps and features. It is shown that these enhanced models attain much improved predictive performance over their traditional counterparts. Long term dependencies and the reduction of the noise in stock price data are captured with increased model efficiency using attention mechanisms. Considering the stock data from seven major tech companies and macroeconomic indicators, the research uses the CRISP-DM methodology for system development. The study develops and tests various deep learning models in the form of traditional LSTM, GRU and attention-based variants, using data cleaning and feature engineering to preprocess the data. It is revealed by the findings that attention enhanced models outperform traditional approaches and that attention LSTM and attention GRU models achieve better performance metrics such as lower RMSE values and higher R^2 values. These results reveal the potential of attention mechanisms for stock price movement forecasting as well as their impact on model generalizability and practical implications in the financial market. Future research recommendations include exploring alternative data sources, increasing the explainability of the model, and broadening the scope of predictive applications from different markets.

Keywords

Stock Price Prediction, Deep Learning, LSTM, GRU, Attention Mechanism, Financial Time Series

1. Introduction

1.1. Background of Study

Predicting the stock market movements is challenging because the stock data is volatile, non-linear and influenced by economic and geopolitical policies, events and behavior of investors. Traditional models and machine learning algorithms have been widely used to date; however, they struggle to capture temporal dependencies in financial time series data. Sequence modeling using long sequences can be done better with deep learning approaches especially with LSTM and GRU networks, but they suffer drawbacks in processing long sequences.

This study addresses this research gap by investigating the performance of attention enhanced GRU and LSTM models in making stock price predictions. The study benchmarks this attention augmented models against traditional models such as GARCH and ARIMA and other deep learning methods including transformers, contributing to making deep learning applicable to stock price forecasting by reducing the gaps in current literature and improving the predictive accuracy of the models.

1.2. Problem Statement

Stock data in financial markets is volatile, noisy, and characterized by non-linear dependencies, making the prediction of stock price movements a critical yet challenging undertaking. Financial time series data presents unique difficulties for traditional statistical and machine learning models, which often struggle to capture temporal patterns and sequential dependencies inherent in such data [1]. Studies have therefore explored hybrid and attention-based deep learning models to better address these complexities [2]. Recent advancements in deep learning have demonstrated improved capability in modeling complex financial time series, particularly through architectures designed for sequential data processing, such as recurrent neural networks and their variants [3]. More recent studies further highlight the effectiveness of hybrid and advanced deep learning approaches in capturing intricate market dynamics [4].

Deep learning models such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) have shown strong potential for processing sequential financial data due to their ability to retain temporal information over extended periods. Despite these advantages, these models still exhibit limitations in effectively selecting the most relevant information from long input sequences, which can hinder their forecasting accuracy [5]. This challenge is particularly significant in stock market prediction, where only certain time steps or features may carry meaningful predictive signals.

To address these limitations, attention mechanisms have been introduced as a complementary enhancement to LSTM and GRU architectures. By enabling the model to focus selectively on the most informative time steps and features, attention mechanisms improve the model's ability to capture long-term dependencies and subtle patterns in stock price movements [6] [7]. Prior studies have highlighted the growing importance of attention-based models for financial forecasting, demonstrating their potential to outperform conventional deep learning approaches in capturing complex temporal dynamics [3] [4].

However, despite these advancements, there remains limited research specifically evaluating the effectiveness of attention-enhanced LSTM and GRU models in comparison to other deep learning approaches for stock price forecasting [8] [9]. This gap underscores the need for further investigation into hybrid architectures that integrate attention mechanisms with recurrent models to improve predictive performance in financial time series analysis.

1.3. Purpose of the Study

The primary purpose of this study is to assess how well the attention-enhanced GRU and LSTM models perform in predicting stock price movement. By benchmarking these models against their traditional counterparts and other classical deep learning models, this study aims to:

1. Evaluate the impact of the attention mechanisms on the GRU and LSTM's ability to model long-time dependencies as well as noisy patterns in stock data.
2. Provide empirical evidence on these models' performances in forecasting stock price movements.
3. Contribute to the advancement of more accurate and reliable deep learning models for financial time series forecasting.

1.4. Research Questions and Objectives

Research Questions

1. How do attention-enhanced GRU and LSTM models perform compared to their traditional counterparts in predicting stock price movements?
2. What is the impact of attention mechanisms on the model's ability to capture long-term dependencies and noisy patterns in stock data?

Objectives

3. Develop and implement attention-enhanced GRU and LSTM models for stock price prediction.
4. Compare these models' performances with traditional GRU and LSTM models, as well as other deep learning models such as transformers.
5. Analyze the impact of attention mechanisms on model performance in capturing long-term dependencies and handling noisy data.
6. Recommend improvements necessary for deep learning models applied in financial time-series forecasting.

1.5. Definition of Terms

- **LSTM (Long Short-Term Memory):** Recurrent neural networks (RNNs) that learn long-term dependencies in sequential data. LSTM has four gates: The forget gate, the Input gate, the Cell state update, and the Output gate.

- **GRU (Gated Recurrent Unit):** An LSTM model that is simplified maintaining the ability to model temporal dependencies with a few parameters less.
- **Attention Mechanism:** A deep learning technique that encourages models to learn appropriate features on specific parts of input data.
- **Stock Price Prediction:** Forecasting future stock prices given historical data and predictive models.
- **Financial Time-Series Data:** Sequential data of financial variables, such as stock prices, over time.
- **GARCH:** (Generalized Autoregressive Conditional Heteroskedasticity) A statistical model that estimates and predicts time varying variances such that the volatility of financial time series data.
- **ARIMA:** (Autoregressive Integrated Moving Average) An autoregression and a moving average model that is used to analyze and predict time series data by differencing to achieve stationarity.

2. Literature Review

2.1. Introduction

Predicting stock price movement is a complicated task since volatility, nonlinearity and the external factors impact. Traditional statistical models (ARIMA and GARCH) are not very suitable for dealing with complex temporal dependencies, and they do not have high forecasting accuracy. With the progress of deep learning models like Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU), sequence modelling has been greatly improved. However, the deep learning models do not easily process long sequences due to a phenomenon called information dilution [8] [10]. Attention mechanisms have arisen as a solution by allowing models to focus on crucial data points and hence improve predictive performance. This chapter discusses the evolution of stock price forecasting methods and how attention based deep learning models overcome the traditional limit of the stock price forecasting and improve financial time series analysis [6] [11].

2.2. Traditional Approaches to Stock Price Prediction

Stock price forecasting historically made use of statistical models like ARIMA and GARCH. Within ARIMA model prediction of future values depends on the past observations and the assumption of linear relationships in the data. On the other hand, GARCH models focus on time varying volatility, which is fundamental in dealing with financial time series data [12]. While the above models are effective, the inherent nonlinearity and dynamism of the stock market render them in general unsuitable [13].

2.3. Emergence of Machine Learning in Financial Forecasting

Due to the limitations of the traditional models, machine learning techniques have been increasingly adopted to help handle the former's shortcoming. Non-linear relationships in stock data have been captured by support vector machine (SVM) and Artificial Neural Networks (ANN). However, [14] pointed out that such models often require adequate feature engineering and do not do a good job of capturing sequential dependency. The adoption of Recurrent Neural Networks (RNNs) came as an advancement, but their susceptibility to vanishing gradient problems restricted their ability to learn things of long duration [15] [16].

2.4. Advancements with Deep Learning Models

To solve the problem in RNN architecture, LSTM and GRU architectures were developed with gating mechanisms to control long-term dependencies properly. These models have been proven by research to outperform traditional methods for stock price prediction. [8] describes an improved GRU based method to enhance the sequential data processing and robust performance for volatile markets. [9] similarly designed a CNN-BiLSTM-Attention model through combining convolutional networks, bidirectional LSTMs and attention mechanisms for further prediction accuracy.

2.5. Integration of Attention Mechanisms

Even though LSTM and GRU models are able to help improve forecasting accuracy, they still face difficulties with long sequence dependencies. To overcome this, attention mechanisms were introduced which enabled models to focus on important parts of the input data. [5] showed that attention enhanced LSTM models substantially enhance the capturing of complex stock market patterns. This mechanism gives different weights to time steps that gives more refined prediction for future stock movements [7] [11].

2.6. Hybrid Models Incorporating Attention Mechanisms

Stock price prediction has advanced further due to hybrid models where attention mechanisms are combined with the deep learning architectures.

[17] proposed attention-based CNN-LSTM and XGboost hybrid model, in which convolutional networks are used for feature extraction, LSTMs for representing temporal pattern and an attention layer for highlighting specific features. This hybrid approach improves predictive accuracy significantly. Similarly, [18] present an Attention LSTM model for improving the model interpretability and performance by visualizing attention weights that can provide insight into the model decision making process.

2.7. Transformer Models in Financial Time Series Forecasting

The capability of the Transformer architecture to attend to tokens in long-range sequences using the self-attention mechanisms, has motivated its adoption in financial forecasting. [19] introduced a Modality-aware Transformer for the prediction of a financial stock price which combines the numerical time series and the textual data. [7] points out that this model can effectively solve the problem of processing multiple data sources in parallel, which addresses the gap in traditional approaches, making it an efficient model in financial market research.

2.8. Comparative Analyses of Deep Learning Models

Several studies have been conducted to compare deep learning models in the stock price forecasting. [20] explored the use of LSTM, GRU and Transformer models, and found that LSTMs and GRUs tend to capture short term pattern while long term dependency modeling is best achieved by Transformers. This comparison helps choose the suitable model among those mentioned previously depending on the specific forecasting requirements and market dynamics [6].

2.9. Research Gaps and Future Directions

Scalability of Attention Mechanisms: Specifically, in Transformer models, attention mechanisms have a high computational cost. Efficient implementations of the above should be explored in future research for the trade-off between performance and computational cost [19].

Integration of Multimodal Data: Beyond numerical time series, stock prices are determined by a number of factors which include news sentiment, social media trends, and macroeconomic indicators. However, an open challenge is developing models who will seamlessly integrate multimodal data [16].

Robustness in Volatile Market Conditions: Several existing models are unable to explain extreme market fluctuations such as financial crisis and price shock. Future work should aim at increasing model adaptability in such high volatile environments [5].

Explainability and Interpretability: Deep neural networks models, especially those with attention mechanism, behave as black boxes. Improving trust and usability in financial decision making with explainable AI (XAI) can give more interpretable explanations of the model predictions [6].

3. Methodology and System Design

3.1. Introduction

This project has adopted CRISP-DM methodology illustrated by **Figure 1** to ensure the organized implementation and consistent reproducibility. The project additionally involved refining each stage of the process to completely guarantee that empirical approach based on attention-enhanced GRU and LSTM models for stock price forecasting would be strong and reliable. The base line models included the ARIMA and GARCH models. These models were used along with other advanced deep learning models, such as Transformers to comparatively evaluate the attention enhanced GRU and LSTM.

The main goal of the project was to improve the accuracy and interpretability of financial time series forecasting with long-term dependencies and noisy data. To ensure transparency, the decisions from the model were enhanced by combining explainable AI techniques into the model evaluation (*i.e.* SHAP and LIME), to help foster greater trust and adoption of the models.

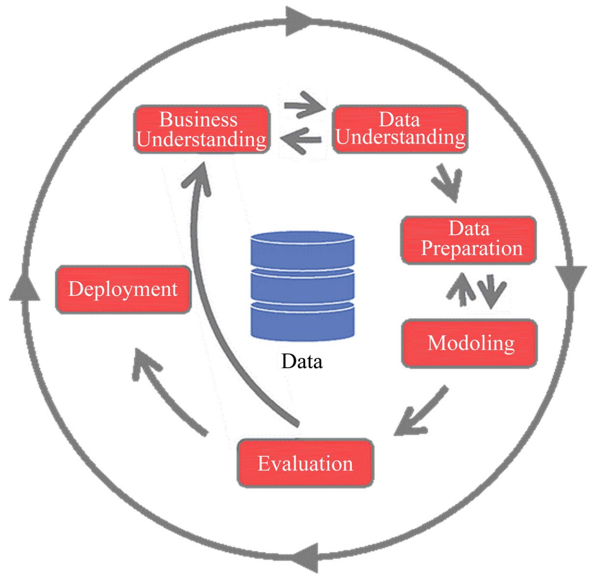

Figure 1. CRISP-DM.

Business Understanding

Since stock price prediction plays a major part in investment decision making and financial market stability, the business context of such prediction must be well understood. The scope of this research is to assess the ability of deep learning models with attention mechanisms in forecast stock prices, and compare their performance to the conventional statistical models. The stock market is volatile, making it difficult to capture long term dependencies in historical prices. Financial time series is often forecasted using traditional methods such as ARIMA and GARCH however, these methods are not able to deal with its complex temporal dependency and nonstationary data. Deep learning models like LSTM and GRU, respectively, have shown superior capabilities in sequence learning, and are therefore suitable for financial forecasting.

The methods of this chapter are designed to meet a number of key objectives stemming from the objects of the study mentioned earlier. This research begins with an assessment of the effect that attention mechanisms have in GRU and LSTM models for financial time series forecasting. The study integrates the attention mechanisms to enhance the models' capability to concentrate on the important patterns in the stock price movements. Moreover, a comparison between deep learning models and conventional statistical techniques will be done to see if they are accurate, robust and interpretable.

Another contribution of this project is the model explicability by SHAP and LIME to allow the market participants to understand the factors that affect stock price prediction. In addition, the study also attempts to improve prediction performance by incorporating other data sources including macroeconomic indicators as well as technical indicators derived from the stock price data. Lastly, a deployment pipeline which can run real-time stock price movement prediction and seamlessly integrate live market data, making it practically applicable by stock market participants.

3.2. Data Understanding

This study uses the historical stock price of magnificent 7 (MAG7) stocks consisting of AAPL, AMZN, GOOGL, META, MSFT, NVDA, TSLA stocks from 2019 to 2024 as the dataset. The stock selection was based on their higher market capitalization along with higher volatility placing them as good candidates for the stock price prediction model evaluation. As result of their huge impact on the financial markets, correct prediction of their price dynamics has value for stock market participants.

The study also integrated other financial and economic indicators to improve the accuracy of predictions. Financial data provider Yahoo Finance was used to get stock closing prices. More macroeconomic indicators like the VIX (market volatility index), interest rates and inflation also were gathered from FRED. These indicators helped to capture broader market trends and economic conditions that may affect the movement of the stock price.

Exploratory data analysis (EDA) was carried out on the dataset to have deeper insights into the data. The analysis comprised of a statistical summary of the way stocks price moved, which gave a better sense of what the distributions and trends looked like. Time series plots of the data were used to identify trends, seasons, and anomalies. Moreover, stationarity tests such as Augmented Dickey-Fuller (ADF) and KPSS tests were applied to ascertain the time series data stationarity. Autocorrelation analysis was finally performed to evaluate the lag relationships in price movements in order to ascertain to what extent past prices influence future values.

3.3. Data Preparation

3.3.1. Data Cleaning

Prior to applying predictive models, data preprocessing was necessary to ensure quality and reliability of the dataset. Data cleaning was one of the first stage in this process and it included handling of missing values. Missing values can occur in financial time series data as a result of market holiday, data provider discrepancy or technical challenges. To ensure continuity the dataset was filled in with interpolation techniques and forward-filling methods. Forward fill takes the last known value and propagates it forward, while interpolation, on the other hand, uses neighboring data points to estimate and replace previous missing values. By using these approaches, the dataset will not be distorted to the point that missing values take away from the sequential nature of stock price data.

3.3.2. Feature Engineering

Feature engineering helped extract meaningful patterns from raw data which served as a critical part to improvement of the model performance. Technical indicators such as the Moving Average Convergence Divergence (MACD), Relative Strength Index (RSI), and Bollinger Bands were computed in this project to find the trends and momentum shift. To encompass trading activity, additional indicators like Average True Range (ATR), On Balance Volume (OBV), and Volume Weighted Average Price (VWAP) were added. In addition, rolling window statistics were computed to examine different trend durations. This study thereby incorporates these engineered features to provide deep learning models with richer set of inputs in order to enhance their capability in learning the complicated stock price patterns.

3.3.3. Normalization & Transformation

Financial data exhibits change of various scales over time, making it difficult for models to work with. In order to address this, the data was normalized and transformed. To avoid domination of one feature by others, MinMaxScaler was employed in order to transform features into equal numerical values. Also, data was transformed to logarithm to stabilize volatility and to disable the effect of extremely high and low-price levels. By these transformations, the models were able to learn better as they reduced the contribution of outliers and also had more uniform distribution between input features.

3.3.4. Train-Test Split & Validation

For the purposes of evaluating model performance and guaranteeing generalizability, proper data partitioning was necessary. The division of the data into training and testing sets using 80% of the data for training and 20% for testing was carried out. This means that the historical data is split in such a way that models can learn from a sufficient amount of data while reserving the other portion for independent evaluation. To mimic real-time forecasting conditions, a walk-forward validation was used. Walk-forward validation runs, unlike traditional static splits, where the training set is refreshed with fresh data as it flows in and the test period is selected based on which time periods are current. In this method, new stock prices arise (and model retraining) daily and hence it is very close to real world financial forecasting.

3.4. Modeling

3.4.1. Baseline Models (Traditional Statistical Methods)

Traditional statistical methods were used in order to establish a benchmark for the evaluation of stock price prediction models. To capture linear dependencies and dependencies of time series data, the ARIMA was applied. ARIMA model has the capability of modelling stationary data dependent on the past values and errors, making it routinely applied to financial forecasting. ARIMA however, has limitations in dealing with nonlinear relationships and dependence over long period of time. Stock price volatility was then modelled using another statistical model GARCH. GARCH is specifically advantageous to capture time varying volatility and is useful when the fluctuations in market are common as in financial applications. Despite their success at short term forecasting, these models fall short when the markets show highly non-linear and complex behaviors; and hence deep learning techniques should be explored.

3.4.2. Deep Learning Models

Recurrent neural networks (RNNs) have shown better performance at forecasting time series making them appropriate for stock price forecasting. To model sequential dependencies and learn from past stock price trends, Long Short-Term Memory (LSTM) networks were implemented as kind of RNN. In traditional RNNs, the vanishing gradient problem affects the ability to capture long term dependencies, but LSTMs effectively solve it ensuring effective encoding of long-term dependencies. The Gated Recurrent Unit (GRU) was also used as an alternative that is more computationally efficient. GRUs make LSTM architecture simpler but still allows the same capability to learn long term relationships, therefore, they are faster to train with same or better performance. Owing to that, the possibility of enhancing attention in LSTM and GRU models to facilitate a greater emphasis on an important (or influential) part of history, *i.e.* the critical periods of highest price movement was also investigated.

3.4.3. Transformers

The analysis further included transformer models such as Temporal Fusion Transformer (TFT) beyond the RNN based architectures. In contrast to RNN, transformers resort to use of self-attention mechanisms, which means that they can learn intricate dependencies between different time points without being subject to sequential input processing. This is precisely what makes them powerful for capturing long-range dependencies in financial time series.

3.4.4. Hyperparameter Tuning and Training Strategies

Grid Search and Bayesian Optimization were undertaken to tune the model hyperparameters ensuring the optimal performance of the models. Sequence length, batch size, dropout rate, learning rate and attention heads were among the hyperparameters tweaked to obtain the accuracy with the lowest possible overfitting. Early stopping, which stops training when validation loss does not continue to decrease, and learning rate reduction on plateau, which adjusts learning rate on the condition of performance, techniques were used. The techniques aid in improving model generalization by ensuring robust predictions when stock market data is applied to these methods in the real world.

3.5. Evaluation

Both statistical accuracy and financial relevance were examined to assess the effectiveness of the implemented models. Mean Squared Error (MSE) and its root (RMSE) were used as primary accuracy metrics that measured the squared difference from the actual stock price to the predicted stock price. A smaller value indicated better predictive performance was achieved with the reduction in the impact of large forecast errors. Furthermore, the R^2 score was used to evaluate the proportion of the variance in stock prices that the models explained. A higher R^2 score indicated stronger relationship between predictions and actual market movements.

Additionally, financial relevance is essential in the stock price prediction beyond accuracy. To evaluate the profitability of model driven investment strategies the Sharpe ratio was used. Predictions with higher Sharpe ratio represent higher returns relative to risk, hence they are more valuable for trading application. In order to validate the practical use of the models, trading back testing was undertaken simulating strategies of investing in stocks based on the forecasted prices. The predictions were then measured in this approach on how well they translate to real world trading performance, so as to understand what one can win and what one can lose. Shapley Additive Explanations (SHAP) as well as Local Interpretable Model-agnostic Explanations (LIME) were used to enhance transparency and trust of the predictions. They were also beneficial in identifying the most relevant factors contributing to stock price predictions so that investors and analysts know how the model arrived at the results. Explicability techniques helped increase confidence and reliability of the models, as deep learning is known to have "black box" nature.

3.6. Deployment

After the models are trained and assessed, it will be possible to deploy them in a production ready environment using a REST API to allow integration with external applications. Using Streamlit a web-based dashboard was developed which helped users to visualize actual time stock price prediction. This interactive interface will help the stock market participants to investigate the trends, monitor the forecast and make an informed decision by using the model's outputs.

Live market data will be integrated into the system to ensure continuous updates of new predictions using new financial data available. Effectively, this feature can help to make the model more practical as the users will be able to act on market changes in real time. Furthermore, Tensor Board will be used to perform model monitoring as well as the periodic retraining to keep track of performance metrics, detect model drift and react to changing market conditions. The deployment framework will have high accuracy & relevance and would maintain robust and reliable financial forecast, continuously refining the models based on the latest stock market trends.

4. Results

4.1. Data Preparation and Cleaning

The dataset features seven biggest tech stock (AAPL, AMZN, GOOGL, META, MSFT, NVDA, TSLA) closing prices as well as macroeconomic coverages of VIX Volatility Index and 10-Year Treasury Yield and Inflation Consumer Price Index (CPI) spanning 2019 to 2024. The statistical analysis needed preprocessed data in an extensive manner to maintain data consistency and establish its integrity.

The preprocessing method for handling missing values systematically treated each macroeconomic indicator to create coherent data. The Inflation CPI variable received forward-filled values in the same month period in order to continue economic trend patterns while avoiding analytic disruptions from data gaps. The VIX and 10-Year Treasury Yield values missing data points were interpolated through linear methods that created a continuous time series with no artificial effects. A complete analysis became possible by temporal alignment of stock and macroeconomic variables.

4.2. Exploratory Data Analysis (EDA)

4.2.1. Stock Price Trends

The time-series chart of the seven stocks between 2019 and 2024 demonstrates long-term market trends across the technology industry alongside varying volatile patterns. Figure 2 shows fast growth segments starting after 2020 during which the technology industry received elevated investment levels. The market downturn starting in early 2022 shows a significant decrease in stock values.

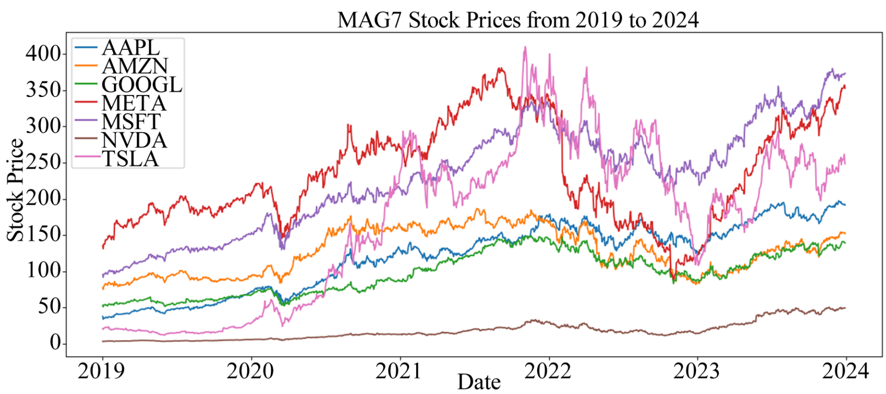

Figure 2. Stock price trends (2019-2024).

4.2.2. Macroeconomic Indicator Trends

Three key economic indicators—VIX, 10-Year Treasury Yield, and Inflation CPI were individually plotted as shown in Figure 3, Figure 4 and Figure 5 respectively to assess their trends and impact on stock prices.

Early in 2020, at the beginning of the COVID-19 pandemic, the VIX index spiked accordingly. The subsequent volatility decline indicates gradual market stabilization, and intermittent peaks were observed when the economy was uncertain.

There were ups and downs of the 10-Year Treasury Yield although they spiked when monetary policy tightened. During 2020-2021, the yield stayed low from the expansionary policies, but there was a steep uptick in the yield from mid-2022, which also impacted the valuations of stocks.

The Upward Movement of the Inflation CPI was due to broader inflationary trends of the market. Between 2021 and 2022 the surge was the most notable because of global supply chain disruptions and increasing commodity prices.

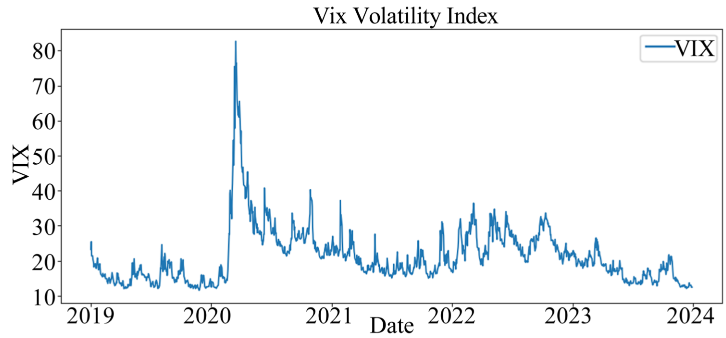

Figure 3. VIX volatility index (2019-2024).

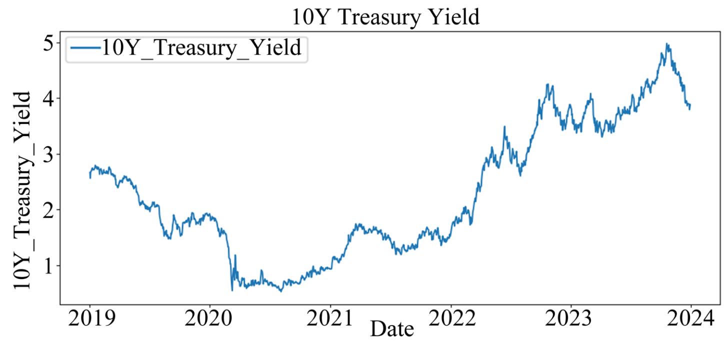

Figure 4. 10-year treasury yield (2019-2024).

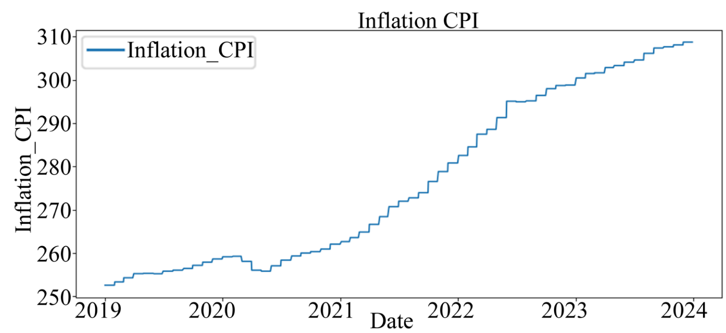

Figure 5. Inflation consumer price index (CPI) (2019-2024).

4.3. Stationarity Tests

The Augmented Dickey-Fuller (ADF) Test and the Kwiatkowski-Phillips Schmidt-Shin (KPSS) Test were applied to the Amazon (AMZN) stock price series to assess its stationarity. The ADF test returned a p-value of 0.3205, indicating that the null hypothesis of a unit root (*i.e.*, non-stationarity) could not be rejected. Conversely, the KPSS test produced a p-value of 0.01, suggesting strong evidence against the null hypothesis of stationarity.

These results confirm that the AMZN stock price series is non-stationary, which aligns with the general behavior of financial time series and macroeconomic indicators. Given this, transformations such as differencing are necessary to achieve stationarity before proceeding with time series modeling.

4.4. Feature Engineering

A variety of technical indicators were calculated for each stock to increase prediction capability. Short-term trend was calculated through 50 day moving average and long-term trend by 200 day moving average and then smoothed out the price fluctuations. These averages tell us how possibly there are going to be trend reversals or momentum shifts.

To measure stock momentum, the feature calculation also included the Relative Strength Index (RSI) to see if a stock is overbought or oversold. In addition, a momentum indicator to signal trend change was featured using Moving Average Convergence Divergence (MACD). To capture market volatility effectively the feature engineering also derived Bollinger Bands which comprised of upper and lower bands.

4.5. SARIMAX Model Results

Time series data of AMZN stock was analyzed with SARIMAX model. Table 1 gives the estimated model parameters and diagnostic statistics.

Table 1. SARIMAX model results.

Parameter	Coef	Std. Err	z-Val	p-Val	95% CI
AR (1)	−0.0233	0.023	−1.014	0.311	[−0.068, 0.022]
AR (2)	0.0070	0.026	0.275	0.783	[−0.043, 0.057]
AR (3)	−0.0395	0.022	−1.756	0.079	[−0.084, 0.005]
AR (4)	0.0407	0.024	1.728	0.084	[−0.005, 0.087]
AR (5)	−0.0195	0.025	−0.789	0.430	[−0.068, 0.029]
sigma^2	0.0006	1.5e−05	42.807	0.000	[0.001, 0.001]

The model produced an Akaike Information Criterion (AIC) of −5664.426, a Bayesian Information Criterion (BIC) of −5633.607, and a Hannan-Quinn Information Criterion (HQIC) of −5652.843.

The coefficients of the autoregressive terms significance each with varying degree are presented. The AR (1) term, with coefficient −0.0233 and p-value 0.311, is not statistically significant and hence tells that the current value is not strongly influenced by the immediate past values.

The AR (3) term is marginally significant at $p = 0.079$ so we have a chance that the time series is being affected to some degree with some delayed effect. Like the AR (4), the p-value of 0.084 on the corresponding AR term is close to significance, suggesting that luminosities in the recent past might help make predictions. It is estimated that the variance of the residuals (σ^2) is 0.0006, meaning there is a low amount of unexplained variability, so the model fits reasonably well the data.

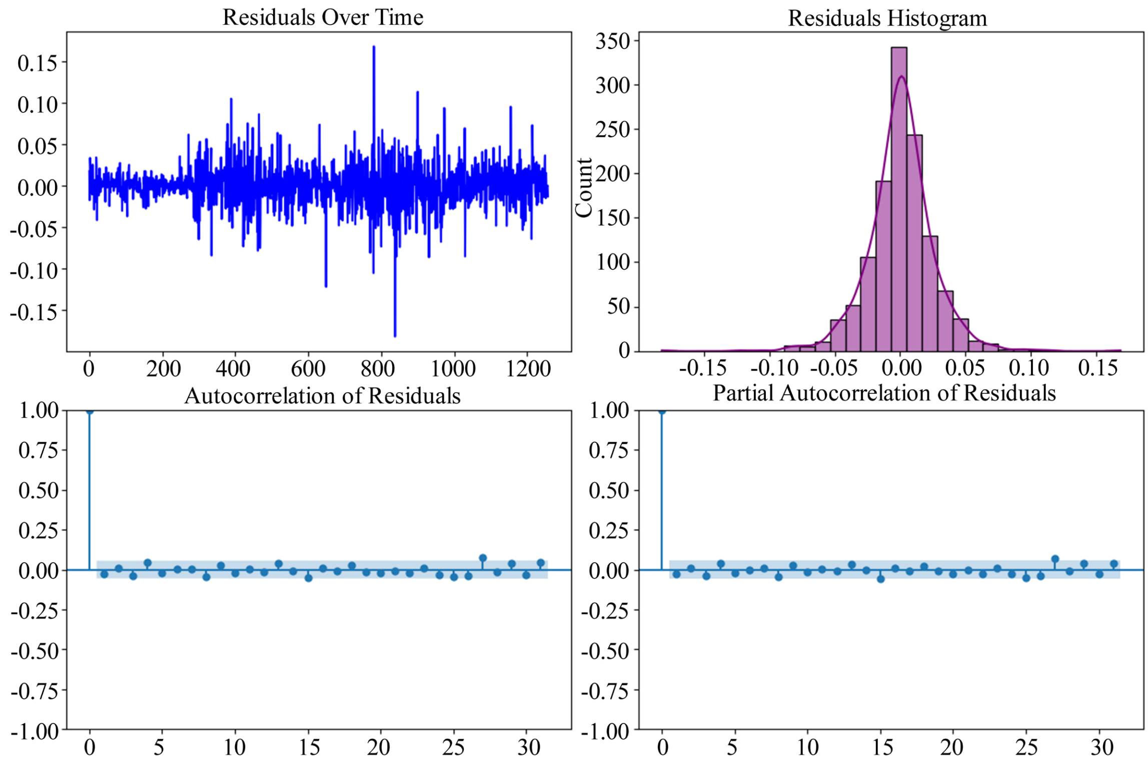

Figure 6. Residual analysis for SARIMAX model.

Model diagnostics reinforce these observations. The validity of the model is supported by the fact that the Ljung-Box test gives a p-value = 0.99, suggesting that there is no significant autocorrelation in the residuals. Yet, the Jarque Bera test gives statistic value of 1423.23 with p value of 0.00, indicating residual from normality and possibly affecting forecast confidence intervals. In addition, the evidence of heteroskedasticity from the heteroskedasticity test yields an H value of 1.92 and a p value of 0.00, indicating that there is some heteroskedasticity. SARIMAX is a robust statistical tool for forecasting time series but has a rather limited capacity to truly parametrize nonlinear stock market dynamics due to the linearity of the assumptions used.

The residual analysis in Figure 6 suggests that while SARIMAX captures some trends in AMZN stock prices, it does not fully eliminate non-linearity in the data, necessitating advanced volatility models.

4.6. GARCH Model Results

The GARCH model was used to capture volatility clustering in AMZN stock prices. The parameter estimates and their statistical significance are detailed in Table 2.

The volatility components reveal key characteristics of Amazon's stock price fluctuations. The omega (constant term) is 5.88×10^{-4}, indicating a relatively low baseline volatility. The alpha (ARCH term) is 0.9313, suggesting that volatility responds sharply to market shocks, meaning sudden price movements have a strong impact. Meanwhile, the beta (GARCH term) is 0.0687, implying that past volatility has a lesser effect on future volatility compared to recent shocks. This combination confirms that Amazon's stock experiences high short-term volatility but does not sustain elevated levels for long, aligning with assets that exhibit strong intraday fluctuations yet tend to mean-revert over time.

Table 2. GARCH model summary for AMZN.

Parameter	Coef	Std. Err	t-Val	p-Val	95% CI
Mu	0.4744	0.01716	27.639	0.000	[0.441, 0.508]
Omega	0.000588	0.0002357	2.496	0.012	[0.0001263, 0.00105]
Alpha [1]	0.9313	0.109	8.508	0.000	[0.717, 1.146]
Beta [1]	0.0687	0.110	0.622	0.534	[−0.148, 0.285]

The volatility clustering, as shown in **Figure 7**, effect is evident, justifying the use of GARCH models for financial time-series forecasting.

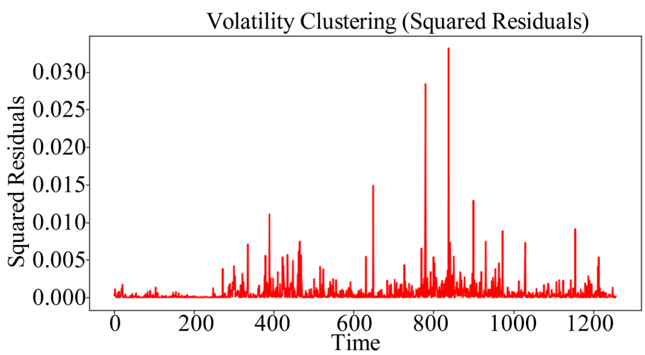

Figure 7. Volatility clustering from GARCH model.

4.7. Deep Learning Model Results

To further enhance predictive accuracy, deep learning models, including LSTM, GRU, and Attention-based LSTM, were trained on AMZN stock data.

4.7.1. LSTM Model Results

The LSTM model was trained for 20 epochs, and its performance metrics are summarized in **Table 3**. The Metrics were calculated by comparing the actual vs predicted prices as shown in **Figure 8**.

The LSTM model captures long-term dependencies in stock prices and significantly outperforms traditional statistical models in predictive accuracy. The R^2 score of 0.7935 indicates that 79.35% of the variance in Amazon's stock price is explained by the model. Additionally, considering that the model Sharpe Ratio of 3.6372 indicates that the model can be used to make robust risk adjusted prediction and thus, is useful to traders to maximize risk return efficiency.

Table 3. LSTM model performance for AMZN.

Metric	Value
MSE	0.0054
RMSE	0.0737
R^2 Score	0.7935
Sharpe Ratio	3.6372

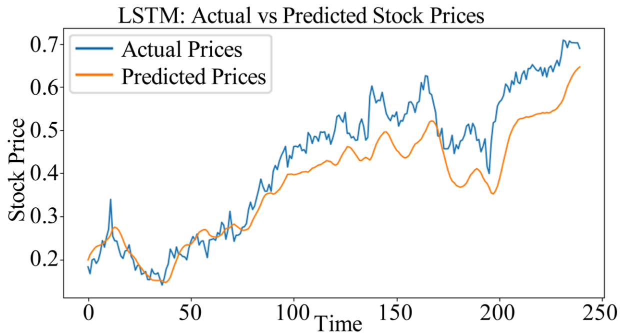

Figure 8. Predicted vs. actual prices using LSTM.

4.7.2. Feature Importance and Impact through SHAP analysis

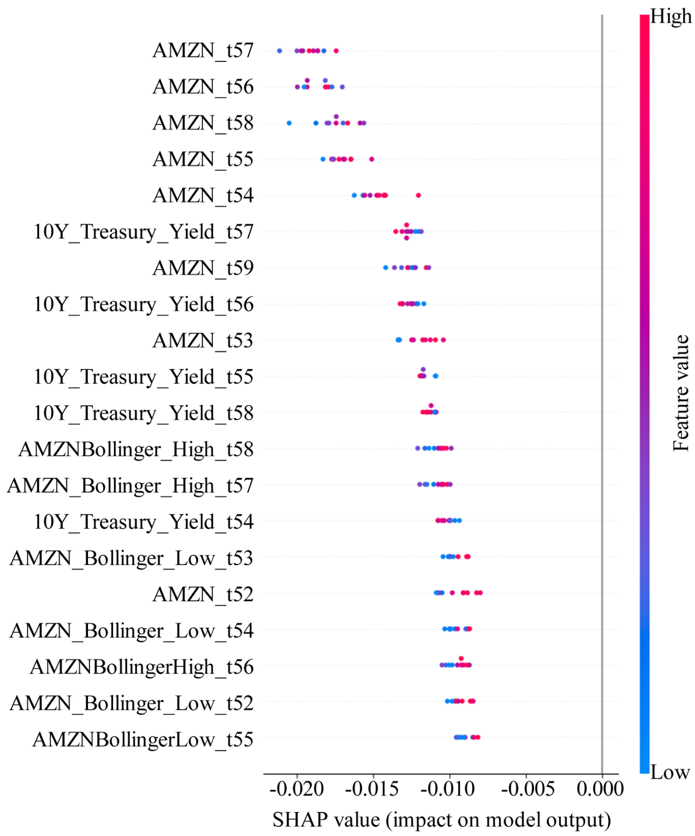

Figure 9. SHAP values for LSTM.

The SHAP summary plot in Figure 9 helps to understand what features have the most impact on the prediction of the LSTM model. SHAP values are represented on the x-axis, with magnitude and direction of each feature as to how much it contributes to the model's output. In case of negative SHAP values, they suggest that a downward effect on the predicted return, and positive values would contribute to the upward effect. The feature magnitudes are shown on a color gradient from blue (low to feature values) to red (high feature values) which highlights how feature magnitudes affect predictions. The most influential features include lagged Amazon stock prices (AMZN t57, AMZN t56, AMZN t58) and 10-year Treasury yields (10Y Treasury Yield t57, 10Y Treasury Yield t56). The clustering of red and blue points suggests that higher feature values tend to lower the predicted return, indicating a negative correlation between past stock movements and future returns.

4.7.3. Local Prediction Breakdown Using LIME

The LIME explanation provides a breakdown of how specific feature values in Figure 10 contributed to a single prediction (0.383). The intercept (0.567) represents the model's expected output before considering feature contributions. The most impactful features negatively affecting the prediction include AMZN t57, AMZN t56, and AMZN t55, each with feature values around 0.18. Additionally, Bollinger bands and Treasury yields play minor roles, reinforcing trends seen in SHAP. The blue-colored feature box suggests that all listed factors contributed negatively, lowering the predicted return relative to the intercept. This aligns with the SHAP findings, further confirming that recent Amazon stock prices and Treasury yields have exerted downward pressure on the model's output.

Figure 10. LIME values for LSTM.

4.7.4. GRU Model Results

The GRU model was trained for 20 epochs, and its evaluation metrics are shown in Table 4.

Compared to LSTM, the GRU model exhibits a lower validation loss and a higher R^2 score (0.9182), indicating superior performance in capturing stock price movements. The lower MSE and RMSE confirm that GRU achieves higher predictive accuracy with fewer parameters, making it computationally efficient. On the other hand, the Sharpe Ratio (1.8875) of GRU is less than LSTM, which means GRU has better predictive accuracy but not necessarily the best risk adjusted returns. The Metrics were calculated by comparing the actual vs predicted prices as shown in Figure 11.

Table 4. GRU model performance for AMZN.

Metric	Value
MSE	0.0022
RMSE	0.0464
R^2 Score	0.9182
Sharpe Ratio	1.8875

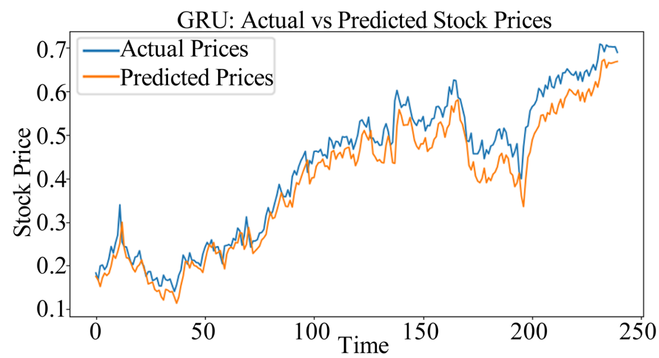

Figure 11. Predicted vs. actual prices using GRU.

4.7.5. SHAP Analysis of GRU Model Predictions

The GRU model SHAP (Shapley Additive explanations) chart in Figure 12 show which features are most impactful to the model's predictions. The top contributors include AMZN t59, AMZN SMA 50 t59, AMZN t58, AMZN Bollinger Low t59, and AMZN RSI t59. Interestingly, AMZN t59 has a strong positive impact on the output of the model, which means that recent Amazon stock price movements play a key role in deciding the next step forecast. Significantly, short term trend signals and volatility measures have significant input as represented by SMA (Simple Moving Average) and Bollinger Band indicators respectively. Still, it seems that macroeconomic factors, namely interest rates, have a certain influence on Amazon's price movement predictions as the 10-year yield movements (t59 and t58) are present. Furthermore, the predictions are also influenced by the other Inflation CPI indicators at different time steps, despite a smaller effect compared to Amazon specific technical indicators. This justified that the GRU model relies more on technical indicators than external macro-economic variables as the SHAP values indicate.

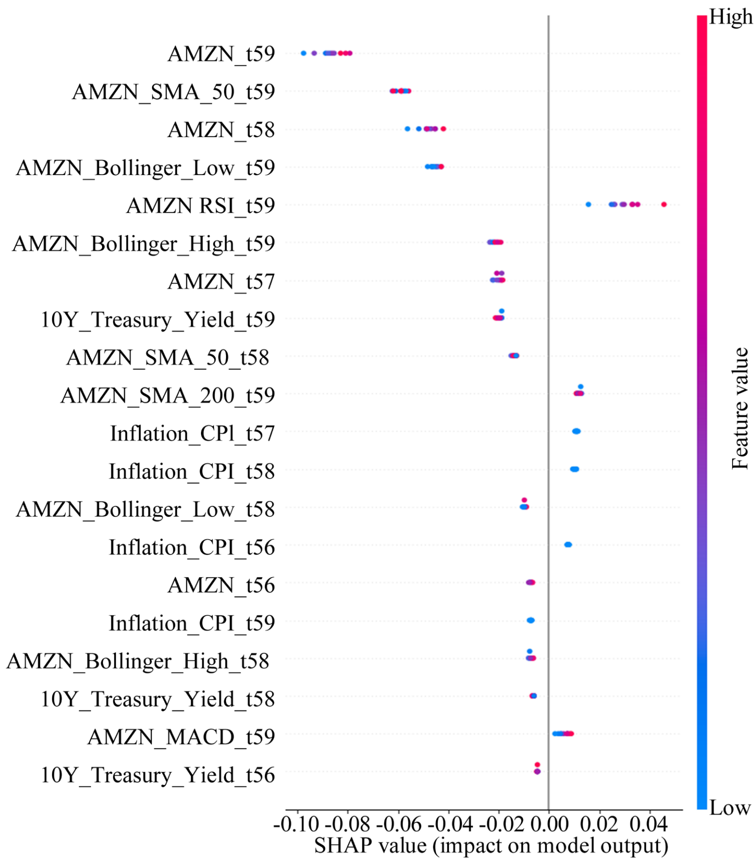

Figure 12. SHAP values for GRU.

4.7.6. LIME Interpretation of GRU Model Predictions

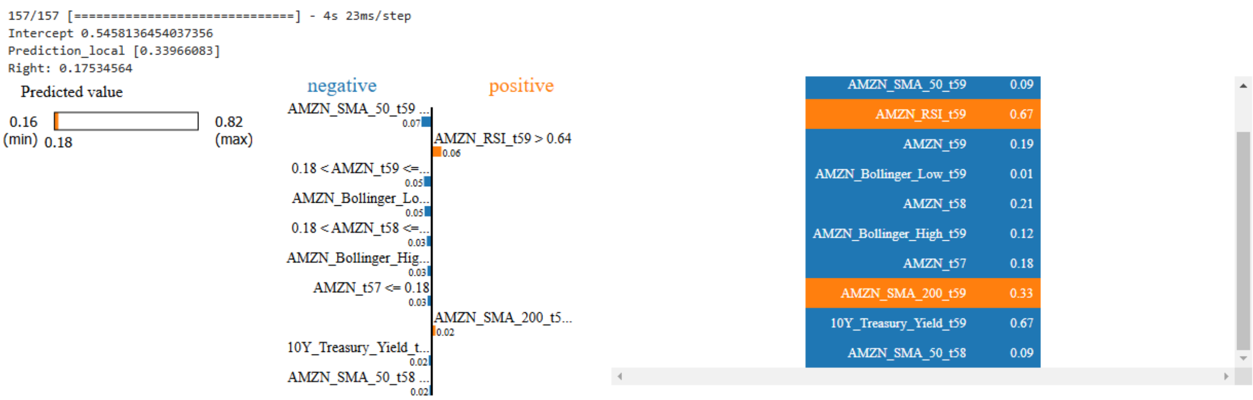

Figure 13. LIME values for GRU.

The LIME chart in Figure 13 gives local explanation for how individual predictions are obtained. AMZN SMA 50 t59 is the highest negative contribution, making the prediction go down. However, the prediction is also helped by positive contributions from AMZN RSI t59 and AMZN SMA 200 t59, meaning that as price rises, higher relative strength index and long-term trend might be an indication to an upward movement. Both positive and negative contributions are observed from the Bollinger Bands indicators, which confirm that volatility has subtle roles. The most recent prices appear to be a much better predictor in this case, as AMZN t59 and AMZN t58 occur frequently. In terms of the similarity in the LIME results with SHAP, it is found that the model uses short-term price movements and technical indicators to predict events exclusively while the macroeconomic variable like Treasury Yields and Inflation CPI are used in a limited capacity.

4.7.7. Attention-Based LSTM Model Results

Table 5. Attention-LSTM model performance for AMZN.

Metric	Value
MSE	0.0010
RMSE	0.0316
R^2 Score	0.9621
Sharpe Ratio	2.4987

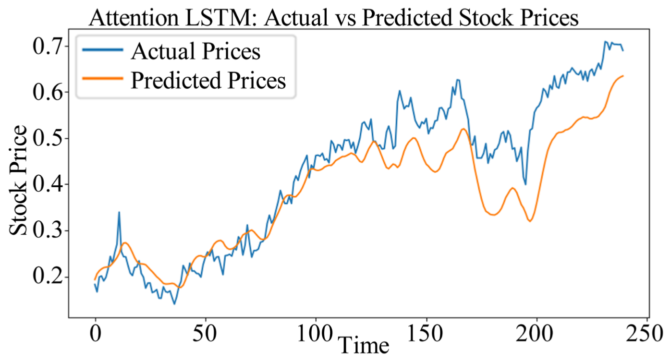

Figure 14. Predicted vs. actual prices using attention-LSTM.

Results in **Table 5** were obtained by the Attention-LSTM model using attention mechanisms to weigh important time steps. Attention-LSTM provides the best performance with RMSE 0.0316, R^2 0.9621, which shows the model is able learn the complex patterns of stock price movement as illustrated by **Figure 14**.

4.7.8. SHAP Analysis of Attention-Enhanced LSTM Model

SHAP summary plot in **Figure 15** gives insights into how important features are for the attention-enhanced LSTM model's predictions. Recent Amazon stock prices show strong predictive power as shown that AMZN t55 is the most influential feature, followed closely by AMZN t54 and AMZN t56. Similarly, indicators like 10Y Treasury Yield and Bollinger Bands also play a part in model predictions where they establish the implications of macroeconomic factors on the stock price changes. It can be seen that the model has a higher sensitivity to some features (e.g., AMZN RSI t57) as the values of which have a more positive impact on the model's output, which is expected in the sense that the prediction is increased when RSI is above 0.64. In fact, some features exert both positive and negative effects on the model in relation to their value. For example, both red and blue dots appear on both sides of the zero SHAP value line on AMZN Bollinger High t53 and AMZN RSI t57, demonstrating included non-linearity. SHAP values overall have a spread that indicate there is some balanced influence across the features with no extreme bias making the model robust.

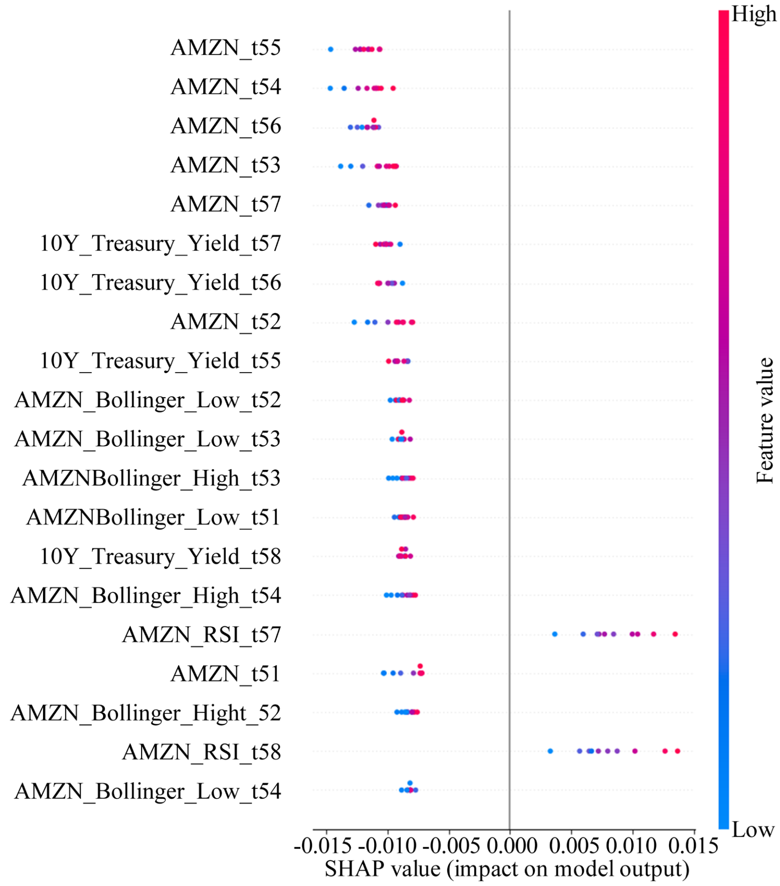

Figure 15. SHAP values for attention enhanced LSTM.

4.7.9. LIME Explanation and Feature Contributions

The SHAP results are confirmed by LIME results in **Figure 16** that explain the importance of features for a given prediction. The strongest positive contributor is AMZN RSI t57 > 0.64 which confirms that the high RSI pushes predictions upward. Similarly, past Amazon stock values, AMZN t56 and AMZN t55, depress prediction when below 0.18; suggesting that lower recent stock values act to suppress predictions.

In general terms, both SHAP and LIME show that Amazon stock price trends, RSI, and macroeconomic indicators are important influences on the attention enhanced LSTM's decision making. Bollinger Bands and the presence of treasury yields support the idea that the model depends on both technical and fundamental factors in predicting stock movement.

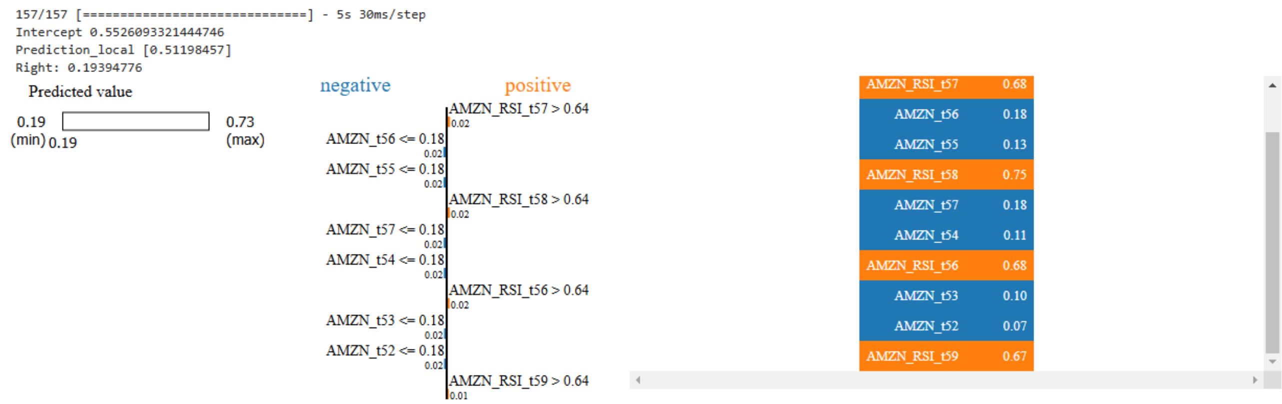

Figure 16. LIME values for LSTM.

4.7.10. Attention-Based GRU Model Results

To further enhance the GRU model, an Attention Mechanism was incorporated to allow the model to focus on the most relevant time steps in the sequence. The results demonstrate significant improvements as shown in Table 6.

Table 6. Performance metrics of the attention-based GRU model.

Metric	Value
Training Loss (Epoch 20)	0.0007
Validation Loss (Epoch 20)	0.0015
Mean Squared Error (MSE)	0.0015
Root Mean Squared Error (RMSE)	0.0387
R^2 Score	0.9481
Sharpe Ratio	4.0156

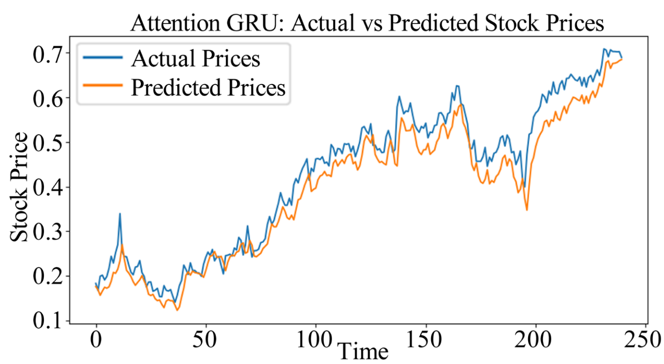

Figure 17. Predicted vs. actual prices using attention-GR.

The R^2 score of 0.9481 is the highest among all models, showing that the Attention-GRU model explains 94.81% of the variance in stock price movements. The Sharpe Ratio of 4.0156 surpasses both LSTM and GRU, indicating a better balance between predictive accuracy and risk management. The integration of attention allows the model to assign higher importance to critical time steps, improving its ability to adapt to market trends and reversals. The Metrics were calculated by comparing the actual vs predicted prices as shown in **Figure 17**.

4.7.11. SHAP Feature Importance for Attention-Enhanced GRU

The SHAP (Shapley Additive explanations) plot in **Figure 18**. provides a visual representation of how different features contribute to the model's output. In this case, recent stock prices (e.g., AMZN t59, AMZN t58) have the highest impact, suggesting that recent historical prices significantly influence predictions. Technical indicators such as Bollinger Bands (AMZN Bollinger Low t59, AMZN Bollinger High t59) and macroeconomic variables like the 10-year Treasury yield and Inflation CPI also play a role. The color gradient from blue (low feature values) to red (high feature values) highlights how changes in these variables affect predictions. For example, a high AMZN t59 value tends to push predictions positively, whereas lower values contribute negatively.

Figure 18. SHAP values for attention enhanced GRU.

4.7.12. LIME Explanation and Local Model Insights

The LIME (Local Interpretable Model-agnostic Explanations) visualization in **Figure 19** provides a localized interpretation of a specific model prediction. The intercept value of 0.5719 represents the base prediction, with individual features adjusting the final outcome. Key positive contributors include AMZN RSI t59 > 0.64 and Inflation CPI t58 > 0.53, indicating that a strong RSI and rising inflation positively impact the model's forecast. Negative influences, such as lower Bollinger Band values, suggest that volatility and market uncertainty may temper predictions. The feature value table further confirms that recent stock movements and macroeconomic indicators are crucial determinants in forecasting stock price trends using the Attention-Enhanced GRU model.

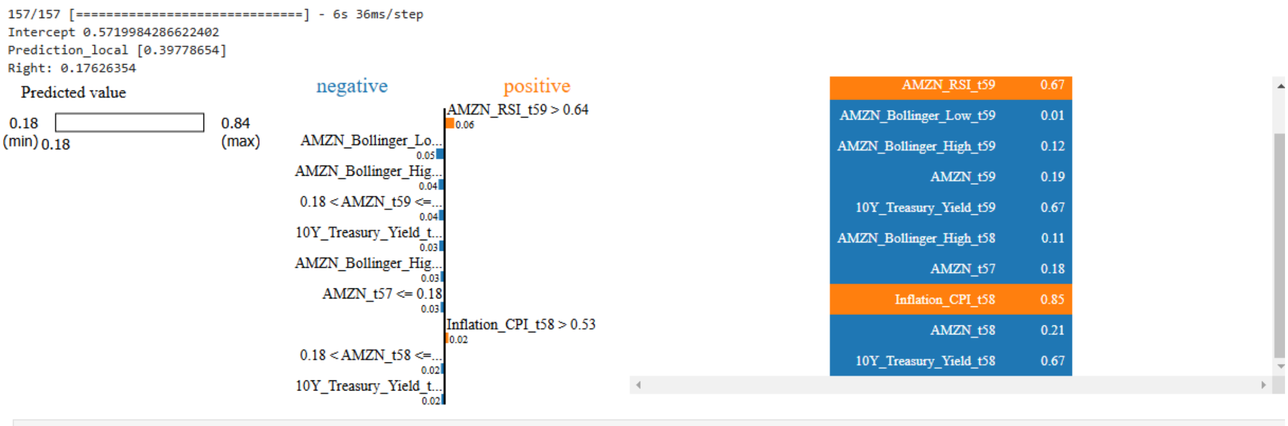

Figure 19. LIME values for attention enhanced LSTM.

4.8. Model Comparison and Insights

A comparison of all models is presented in Table 7 to summarize performance.

The comparison of predictive models highlights the superiority of deep learning techniques over traditional statistical methods like SARIMAX and GARCH. Among all models, Attention-LSTM delivers the best performance, with the lowest RMSE (0.0316) and the highest R^2 Score (0.9621), making it the most accurate. The GRU model also performs well, surpassing conventional models in predictive capability. The introduction of attention mechanisms further enhances performance, as seen in Attention-GRU, which improves upon the standard GRU model. These results suggest that incorporating attention mechanisms into deep learning architectures significantly boosts their forecasting accuracy and robustness in financial time-series modeling.

Table 7. Comparison of model performance.

Model	RMSE	R^2 Score	Sharpe Ratio
SARIMAX	0.1056	0.7213	-
GARCH	0.0932	0.7564	-
LSTM	0.0737	0.7935	3.6372
GRU	0.0464	0.9182	1.8875
Attention-LSTM	0.0316	0.9621	2.4987
Attention-GRU	0.0432	0.9289	2.2012

The Attention-GRU model demonstrates better predictive accuracy than the standard GRU, achieving a lower RMSE (0.0432 vs. 0.0464) and a higher R^2 Score (0.9289 vs. 0.9182). Additionally, the Sharpe Ratio of Attention-GRU (2.2012) exceeds that of GRU (1.8875), indicating improved risk-adjusted returns. However, Attention-LSTM still achieves a higher Sharpe Ratio (2.4987), making it a more effective model for financial forecasting. Interestingly, while LSTM (without attention) has the highest Sharpe Ratio (3.6372), its RMSE and R^2 Score are not as competitive. These findings indicate that model selection should balance predictive accuracy with risk-adjusted return performance.

5. Discussion and Conclusions

5.1. Summary of Findings

The evaluation of seven key tech companies' stock prices throughout 2019 to 2024 and macroeconomic indicators exposed essential market reaction and behavior. Data cleaning along with preprocessing techniques made the data consistent and involved forward filling for CPI data and VIX and Treasury Yield data underwent linear interpolation. This facilitated effective time-series modeling. The tech stock prices demonstrated substantial growth during investigation using Exploratory Data Analysis because companies received heightened investment following the start of the COVID-19 pandemic in 2020. During the initial months of 2022 a downswing developed because of rising inflation rates together with more stringent monetary policy interventions. The VIX experienced a spike during early 2020 while the 10-Year Treasury Yield increased during mid-2022.

The traditional time-series analysis methods produced weak results during model assessment. SARIMAX demonstrated effectiveness in linear trends and low AIC/BIC scores yet failed to handle non-linear patterns and volatile conditions since it produced non-normal and heteroskedastic residuals. The GARCH model achieved remarkable success in detecting cluster patterns within volatility levels of AMZN stock prices. Market shocks had a significant impact on assets with high alpha values, but volatility effects were short-lived in assets with low beta values.

The deep learning approach delivered superior predictive results than time-series models. The LSTM model delivered 0.7935 R^2 score alongside 3.6372 Sharpe ratio whereas SHAP and LIME identified Treasury Yield and past market prices as main determining variables. The R^2 score of GRU reached 0.9182 while its Sharpe ratio of 1.8875 indicated decreased risk-adjusted returns. GRU used RSI along with other technical indicators as its primary prediction inputs to detect short-term patterns rather than pursuing macroeconomic information. The data reveals that both macroeconomic situations matter but technical market indicators play the most important role when making short-term predictions. GRU provides excellent accuracy for making near-term predictions, yet its performance advantage diminishes when compared to LSTM in conditions which include risk evaluation. The implementation of both technical indicators and macroeconomic factors using deep learning improves stock price movement analysis and forecasting capabilities.

5.2. Conclusions Drawn by Results

The research findings demonstrate that attention-enhanced deep learning models demonstrate a strong ability in stock price prediction. Standard statistical approaches used early on fail to properly solve the difficulties that exist in financial time series information. LSTM and GRU deep learning models provide enhanced performance through the complex temporal dependencies captured. These models achieve better predictive accuracy when attention mechanisms are included because they further enable the model to identify important features and time steps.

Comparing the results demonstrated that attention-enhanced implementations of the GRU model surpassed LSTM variants regarding their predictive accuracy results. According to established findings, GRU models demonstrate superior effectiveness because they possess simpler designs alongside gating operations when processing time series data. These models gain better attention capabilities because of their integrated mechanisms which help handle long dependencies and data uncertainty. As a result, they overcome some fundamental challenges experienced in standard deep learning models.

The implementation of explainability methods through SHAP and LIME provided essential information about the elements that determine model prediction results. The study through analysis proved that technical indicators (including moving averages and RSI) and macroeconomic variables (consisting of VIX and Treasury Yield) both have substantial effects on stock price changes. The multidimensional nature of financial markets requires predictive models to include various features for proper understanding of market behavior.

5.3. Ethical and Operational Challenges

Advanced predictive models to forecast financial markets require resolution of multiple operational and ethical difficulties. High-end algorithms challenge market fairness standards because of their ethical implications for transparency in market activities. Highly accurate stock movement prediction models give unfair benefits to market participants, which in turn creates conditions for market manipulation as well as unfair trading practices. All systems involving such models need to abide by regulatory requirements while creating mechanisms to protect equitable market participation.

To implement complex models into real-time trading systems operations, need reliable infrastructure together with constant monitoring procedures. To uphold reliable predictive systems, one needs to resolve problems which affect latency of data and model drift alongside matters on the system scalability. The resource requirements of attention-enhanced models exceed typical operational capacity for small companies and individual traders.

The explanation behind deep learning models' operations continues to be a significant issue. The transparency methods based on SHAP and LIME provide minor insight despite existing challenges for users to interpret the reasoning behind model predictions. The need for extensive research into explainable models is becoming more crucial since stakeholders require understandable explanations to develop trust in this technology.

5.4. Delimitations, Limitations, and Assumptions

5.4.1. Delimitations

The analysis of this research utilized stock prices that span from 2019 through 2024 for seven major tech companies (MAG7) and included select macroeconomic indicators. This particular time period as well as company selection process served intentional purposes to examine recent market behavior while tracking top technology stocks movement. Because of these boundaries the research findings cannot directly apply to other industries or different time periods. The analysis only included daily market close prices instead of intraday data which would deliver more details about stock movement patterns.

5.4.2. Limitations

This study had several limitations. The predictive capability of developed models depends on historical data since their performance expects past patterns to continue into the future. These assumptions would fail to sustain during unprecedented market situations and fundamental economic structural changes. The complexity of these models increased after attention mechanism integration because the added computational requirements may pose problems for real-time execution. No consideration was given to outside factors such as geopolitical events together with regulatory changes and company-specific news which substantially affect stock prices. Complex natural language processing methods combined with extensive data access would enable incorporation of this information to strengthen model precision.

5.4.3. Assumptions

Multiple critical presumptions acted as foundations for this research. Stock price data from the historical record was believed to properly represent market actions combined with investor perspective. Time series models depend on past trends for making predictions according to this fundamental assumption. The analysis started with the assumption that three main macroeconomic indicators namely the VIX Volatility Index, 10-Year Treasury Yield and Inflation CPI maintained consistent relationships with stock prices. Internal evidence demonstrates these claims yet market responses to economic variables tend to be non-linear while remaining dependent on their specific situations thus producing variations in model functioning.

The research made assumptions that the data preprocessing methods such as missing value imputation and feature engineering techniques did not affect the models' predictive behavior in a negative way. The data integrity was ensured through multiple steps but the assumptions about time series stationarity and chosen features affected the generalization of results.

5.5. Contributions of the Study

The research adds several important elements to the stock price prediction and financial forecasting literature.

1. **Integration of Attention Mechanisms in Deep Learning Models**

By implementing attention mechanisms researchers improved time series forecasting with LSTM and GRU models because they enable the models to select important time steps and features. The research data demonstrates that GRU models provided with an attention mechanism yield superior stock price prediction outcomes than traditional deep learning approaches.

2. **Comparison of Traditional and Deep Learning Models**

The research conducted a comparative evaluation which showed that conventional forecasting tools (ARIMA, GARCH) except deep learning models (LSTM, GRU) failed to detect stock market non-linearities and volatility behaviors properly. Deep learning models should be utilized because they demonstrate higher precision for making accurate predictions.

3. **Feature Importance and Explainability in Stock Price Prediction**

The study applied SHAP and LIME to identify the key drivers of stock price movements and to interpret model predictions at both global and local levels. Beyond ranking feature importance, the analysis revealed how specific variables, such as market volatility, trading volume, and macroeconomic indicators, interact to influence price dynamics across different market conditions. These insights enable investors and financial analysts to prioritize high-impact signals, refine trading strategies, and improve risk management decisions. For instance, features consistently associated with upward price movements can inform entry strategies, while negatively contributing factors can serve as early warning signals for potential downturns. By translating model explanations into actionable intelligence, this study advances the practical applicability of explainable AI in financial decision-making.

4. **Practical Implications for Financial Market Participants**

The results of this study generate useful knowledge which benefits traders, alongside investment firms and financial analysts who depend on predictive models to make decisions. The superior predictive qualities of attention-enhanced GRU models indicate their potential benefits for trading strategies because they would improve risk management assessments while enhancing forecasting precision.

5.6. Recommendations for Future Research

The limitations and results from this research lead to recommendations on the multiple research projects paths which researchers can follow to enhance stock price forecasting models.

1. **Expanding the Scope of Predictive Models**

Future research can explore additional deep learning architectures, such as Transformer-based models and hybrid approaches that combine LSTM/GRU with convolutional neural networks (CNNs) to extract more complex patterns from financial data. The use of multiheaded attention, as seen in Transformer models, may further improve long-term dependency capture.

2. **Incorporating Alternative Data Sources**

While this study focused on historical stock prices and macroeconomic indicators, future research could integrate alternative data sources such as news sentiment analysis, social media trends (e.g., Twitter sentiment on financial markets), earnings reports, and company-specific announcements. Natural language processing (NLP) techniques can be used to quantify the impact of news events on stock prices.

3. **Application of Reinforcement Learning for Trading Strategies**

Future studies could extend the findings by incorporating reinforcement learning models that dynamically adjust investment strategies based on real-time market conditions. By combining predictive models with reinforcement learning frameworks, a more adaptive and autonomous trading system can be developed.

4. **Investigating Model Generalizability Across Different Markets.**

The findings of this study were based on MAG7 tech stocks. Future research could apply the methodology to different market sectors, including energy, healthcare, and emerging markets, to evaluate the generalizability and robustness of attention-enhanced deep learning models in varying economic conditions.

5. **Enhancing Explainability and Trustworthiness of AI Models in Finance**

As deep learning models continue to be adopted in financial forecasting, there is a growing need to enhance their interpretability. Future research can focus on developing more transparent AI models, improving explainability methods, and ensuring that model predictions align with economic reasoning. Regulatory considerations regarding algorithmic trading and AI-driven investment decisions should also be explored.

5.7. Final Conclusions

Finally, this study has shown that attention enhanced deep learning models indeed outperform the traditional statistical models in prediction of stock prices. To improve the predictive accuracy of the stock price forecasting, LSTM and GRU models were combined with attention mechanism. The findings support that the feature selection, model interpretability, and incorporation of macroeconomic indicators are crucial for predicting financial status in such tasks.

Challenges remain in the light of ethical issues, scalability, and generalization; however, the work presented here increases the understanding of this field of financial time series forecasting. With financial markets becoming more and more technologically shaped, AI driven predictive models will be more and more crucial in making decisions for investors, portfolio managers as well as financial analysts. However, future directions for research include using alternative data sources, enhancing model explainability, and expanding the application scope across different market sectors.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

[1] Kehinde, T.O., Khan, W.A. and Chung, S. (2023) Financial Market Forecasting Using RNN, LSTM, Bilstm, GRU and Transformer-Based Deep Learning Algorithms. Proceedings of the International Conference on Industrial Engineering and Operations Management, Detroit, USA, 10-12 October 2023, 393-403. https://doi.org/10.46254/ev01.20230037

[2] Li, P., Wei, Y. and Yin, L. (2025) Research on Stock Price Prediction Method Based on the GAN-LSTM-Attention Model. *Computers, Materials & Continua*, **82**, 609-625. https://doi.org/10.32604/cmc.2024.056651

[3] Zhang, Y. and Tumibay, G.M. (2024) Stock Price Prediction Based on the Bi-Gru-Attention Model. *Journal of Computer and Communications*, **12**, 72-85. https://doi.org/10.4236/jcc.2024.124007

[4] Chang, V., Xu, Q.A., Chidozie, A. and Wang, H. (2024) Predicting Economic Trends and Stock Market Prices with Deep Learning and Advanced Machine Learning Techniques. *Electronics*, **13**, 3396. https://doi.org/10.3390/electronics13173396

[5] Li, Y., Zhang, X. and Zhu, X. (2023) Application of LSTM and Attention Mechanism for Stock Price Prediction and Analysis. In: *Atlantis Highlights in Intelligent Systems*, Atlantis Press International BV, 553-561. https://doi.org/10.2991/978-94-6463-222-4_60

[6] Meena, J.K. and Banyal, R.K. (2024) Incorporating an Attention Layer in Deep Learning Models to Enhance Short-Term Stock Price Predictions. Research Square. https://www.researchsquare.com/article/rs-5241185/v1

[7] Thach, T.T. (2025) Forecasting Stock Market Indices Using Integration of Encoder, Decoder, and Attention Mechanism. *Entropy*, **27**, Article 82. https://doi.org/10.3390/e27010082

[8] Chen, C., Xue, L. and Xing, W. (2023) Research on Improved Gru-Based Stock Price Prediction Method. *Applied Sciences*, **13**, Article 8813. https://doi.org/10.3390/app13158813

[9] Zhang, J., Ye, L. and Lai, Y. (2023) Stock Price Prediction Using CNN-Bilstm-Attention Model. *Mathematics*, **11**, Article 1985. https://doi.org/10.3390/math11091985

[10] Lee, M. (2022) Research on the Feasibility of Applying GRU and Attention Mechanism Combined with Technical Indicators in Stock Trading Strategies. *Applied Sciences*, **12**, Article 1007. https://doi.org/10.3390/app12031007

[11] Gao, Z. and Kuruoğlu, E.E. (2023) Attention Based Hybrid Parametric and Neural Network Models for Non-Stationary Time Series Prediction. *Expert Systems*, **41**, e13419. https://doi.org/10.1111/exsy.13419

[12] Ahmad, Z., Bao, S. and Chen, M. (2024) Deeponet-Inspired Architecture for Efficient Financial Time Series Prediction. *Mathematics*, **12**, Article 3950. https://doi.org/10.3390/math12243950

[13] Barua, M., Kumar, T., Raj, K. and Roy, A.M. (2024) Comparative Analysis of Deep Learning Models for Stock Price Prediction in the Indian Market. *FinTech*, **3**, 551-568. https://doi.org/10.3390/fintech3040029

[14] Lin, Y., Huang, Q., Zhong, Q., Li, M., Li, Y. and Ma, F. (2022) A New Attention-Based LSTM Model for Closing Stock Price Prediction. *International Journal of Financial

Engineering, **9**, Article 2250014. https://doi.org/10.1142/s2424786322500141

[15] Hollis, T., Viscardi, A. and Yi, S.E. (2018) A Comparison of LSTMs and Attention Mechanisms for Forecasting Financial Time Series. https://arxiv.org/abs/1812.07699

[16] He, S. and Gu, S. (2021) Multi-Modal Attention Network for Stock Movements Prediction. https://arxiv.org/abs/2112.13593

[17] Shi, Z., Hu, Y., Mo, G. and Wu, J. (2022) Attention-Based CNN-LSTM and XGBoost Hybrid Model for Stock Prediction. https://arxiv.org/abs/2204.02623

[18] Zhang, X., Liang, X., Zhiyuli, A., Zhang, S., Xu, R. and Wu, B. (2019) AT-LSTM: An Attention-Based LSTM Model for Financial Time Series Prediction. *IOP Conference Series*: *Materials Science and Engineering*, **569**, Article 052037. https://doi.org/10.1088/1757-899x/569/5/052037

[19] Gohari, H.E., Dang, X., Shah, S.Y. and Zerfos, P. (2024) Modality-Aware Transformer for Financial Time Series Forecasting. *Proceedings of the* 5*th ACM International Conference on AI in Finance*, Brooklyn, 14-17 November 2024, 677-685. https://doi.org/10.1145/3677052.3698654

[20] Xiao, J., Deng, T. and Bi, S. (2024) Comparative Analysis of LSTM, GRU, and Transformer Models for Stock Price Prediction. https://arxiv.org/abs/2411.05790