INTERNATIONAL JOURNAL OF ADVANCE RESEARCH, IDEAS AND INNOVATIONS IN TECHNOLOGY
IJARIIT

ISSN: 2454-132X
Impact Factor: 6.078
(Volume 11, Issue 5 - V11I5-1227)
Available online at: https://www.ijariit.com

Forecasting Stock Market Prices Using Long Short-Term Memory (LSTM)

Siddhi Rajput
siddhirajput2112@gmail.com
Independent Researcher

ABSTRACT

This study applies a Long Short-Term Memory (LSTM) neural network to forecast stock closing prices for selected technology companies (Apple, Google, Microsoft, and Amazon). The paper documents data collection, preprocessing, exploratory analysis (returns, volume, correlations), model architecture, and results. The aim is to evaluate LSTM's ability to capture temporal patterns in stock prices and to provide practical insights for short-term forecasting. Key findings show that the LSTM model captures overall price trends and produces reasonable short-horizon forecasts; however, prediction accuracy is affected by market volatility, data noise, and model complexity.

Keywords: *Stock Market Prediction, Long Short-Term Memory, LSTM, Deep Learning, Neural Networks, Time Series Forecasting, Financial Analytics, Machine Learning, Price Prediction.*

INTRODUCTION

Predicting stock prices has long been of interest to traders, analysts, and researchers. Traditional statistical methods often fail to capture complex temporal dependencies in financial time series. Recurrent neural networks, especially Long Short-Term Memory (LSTM) networks, are designed to model long-range dependencies and have shown promise for time-series forecasting. This paper demonstrates an end-to-end approach to forecasting stock closing prices with LSTM, using historical data from major technology companies. The structure of the paper follows a standard research workflow: data collection, exploratory analysis, model design, results, and discussion.

RESEARCH QUESTIONS

This study addresses the following questions:

i. How effective are LSTM models at predicting short-term stock prices compared to basic statistical baselines?
ii. What are the main challenges and limitations when applying LSTM to financial time series?
iii. Can LSTM-based forecasts provide actionable insights for traders and researchers?

LITERATURE REVIEW

The LSTM architecture was introduced to mitigate the vanishing gradient problem in recurrent neural networks and has since been widely used in sequence modeling tasks such as speech recognition and natural language processing. Prior work has applied LSTMs to financial prediction with mixed results: some studies report useful short-term forecasts while others stress difficulties arising from noisy and non-stationary financial data. This work builds on that literature by applying a simple LSTM pipeline and documenting data-driven insights from multiple exploratory plots. Key references include Hochreiter & Schmidhuber (1997) and Chen et al. (2015).

DATA COLLECTION AND PREPROCESSING

Historical daily stock data for Apple (AAPL), Google (GOOG), Microsoft (MSFT), and Amazon (AMZN) was retrieved from Yahoo Finance. The dataset includes Open, High, Low, Close, Adjusted Close, and Volume for the time window used in this study. Preprocessing steps included:

i. Handling missing values by forward-filling where appropriate.
ii. Calculating daily returns as the percentage change in closing price.
iii. Scaling the close price series to normalize inputs for the LSTM model (MinMax scaling to range [0, 1]).
iv. Splitting data into training, validation, and test sets (training covers the historical period up to late 2023, validation covers 2024, and the final portion is reserved for short-term predictions).

EXPLORATORY VISUALIZATIONS

The following figures show the distribution of daily returns, closing price history, and sales volume for each company. These provide intuition about volatility, trends, and trading activity.

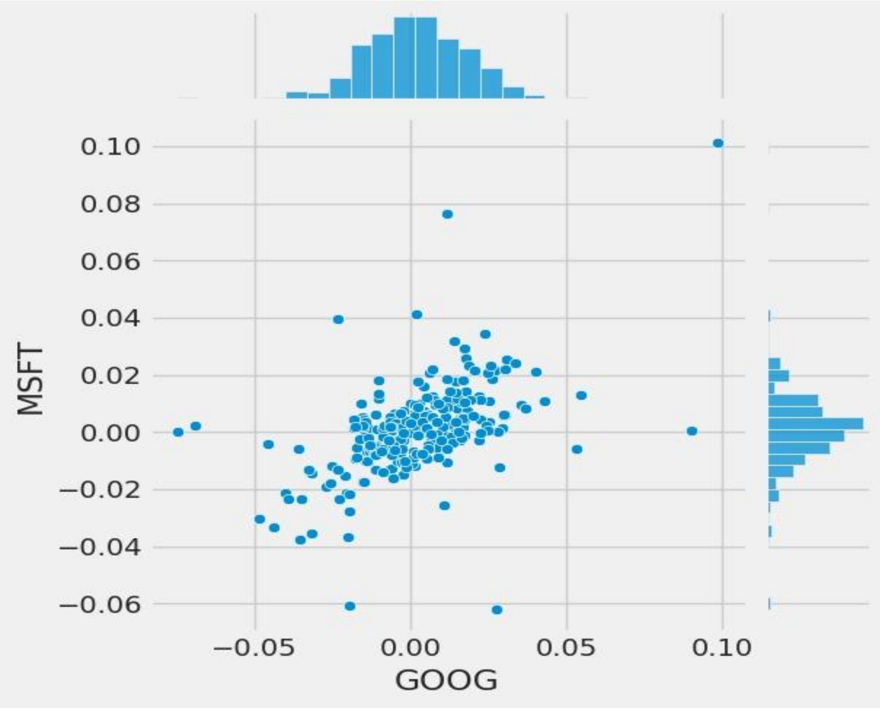

Figure 1: Exploratory plot 1

Figure 2: Exploratory plot 2

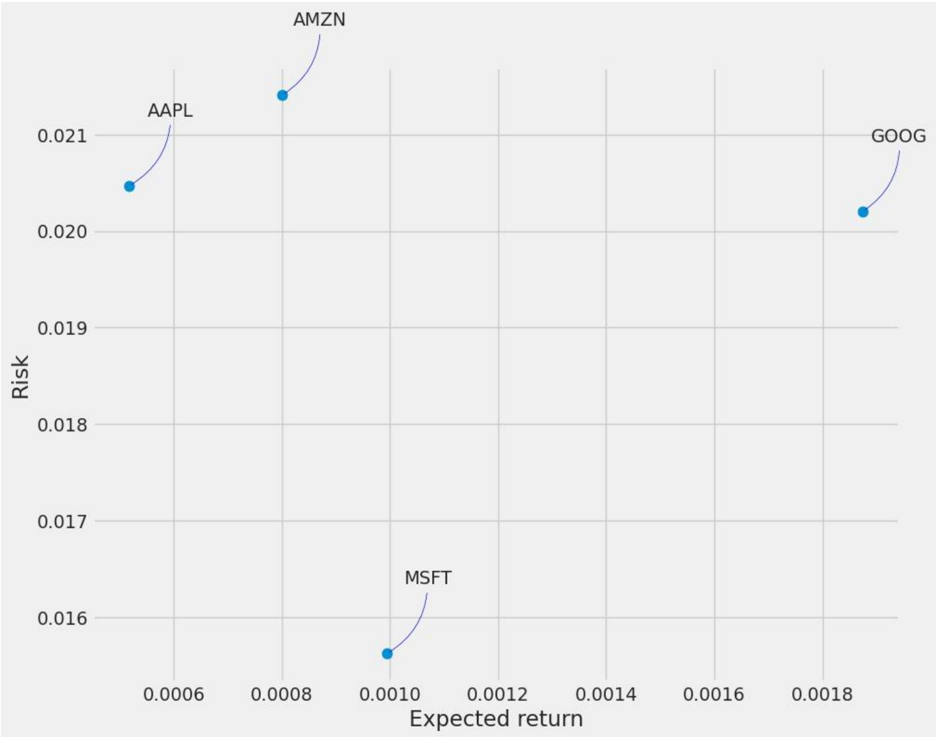

Figure 3: Exploratory plot 3

SALES VOLUME

Volume plots indicate episodes of unusually high trading activity. These spikes often coincide with major market events or company announcements and can affect short-term price movement.

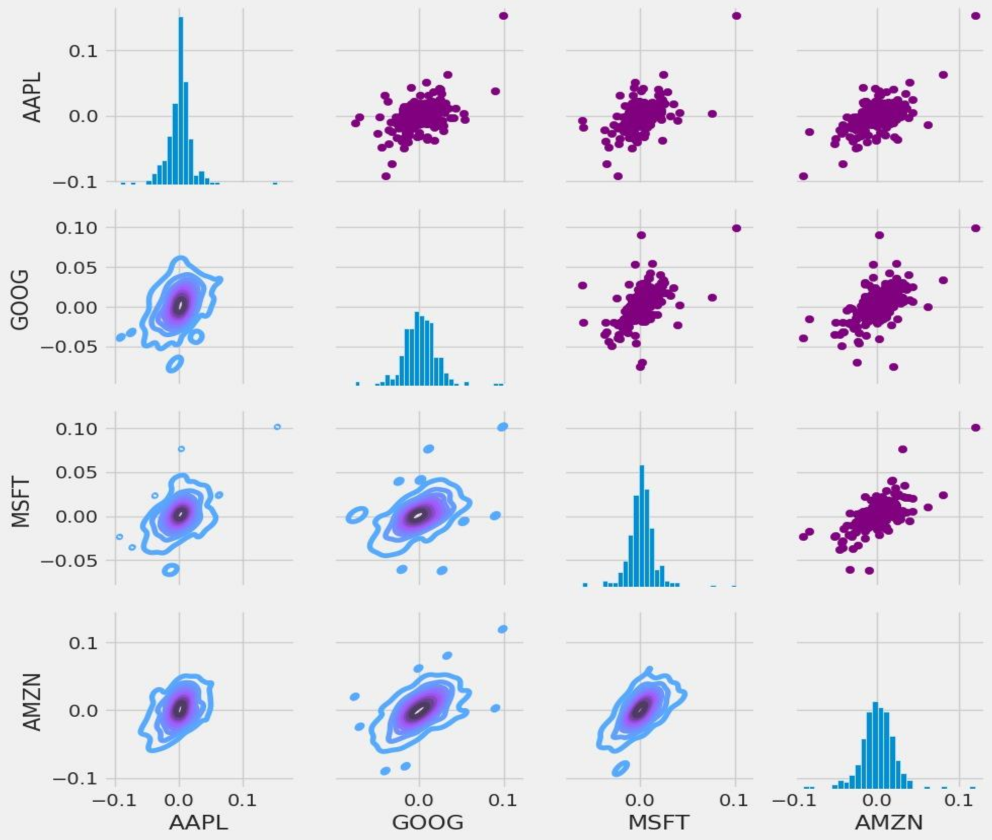

Figure 4: Sales volume time series for AAPL, GOOG, MSFT, and AMZN.

DAILY RETURN ANALYSIS

Daily return histograms and time series reveal distribution shapes and outliers. The histograms below show that returns are approximately centered near zero with occasional large positive or negative values.

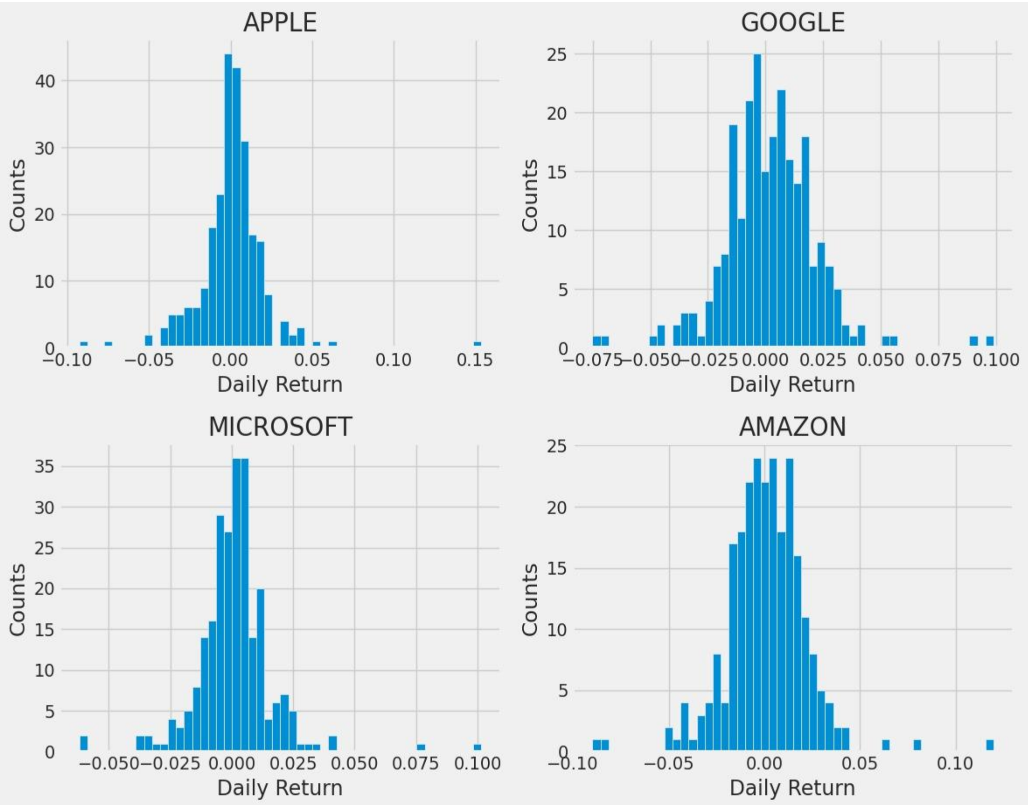

Figure 5: Daily return visualization 1

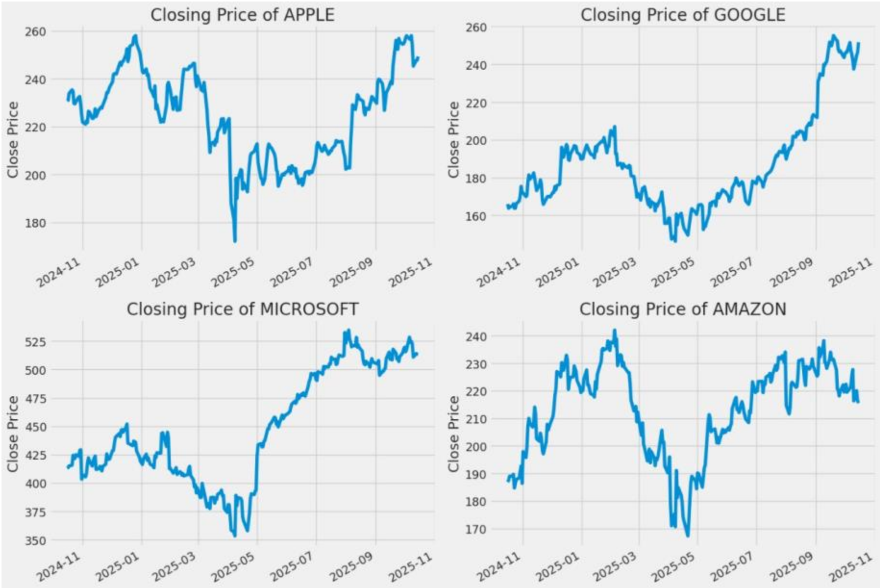

Figure 6: Daily return visualization 2

CORRELATION ANALYSIS

Correlation matrices for returns and closing prices are useful to understand linear relationships between assets. The heatmaps below show moderate to high positive correlations among most pairs, suggesting that market-wide factors influence these stocks.

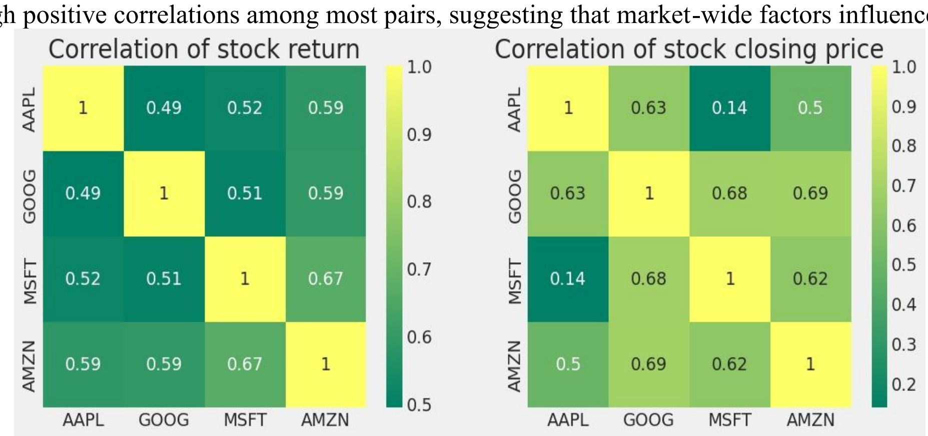

Figure 7: Correlation heatmap 1

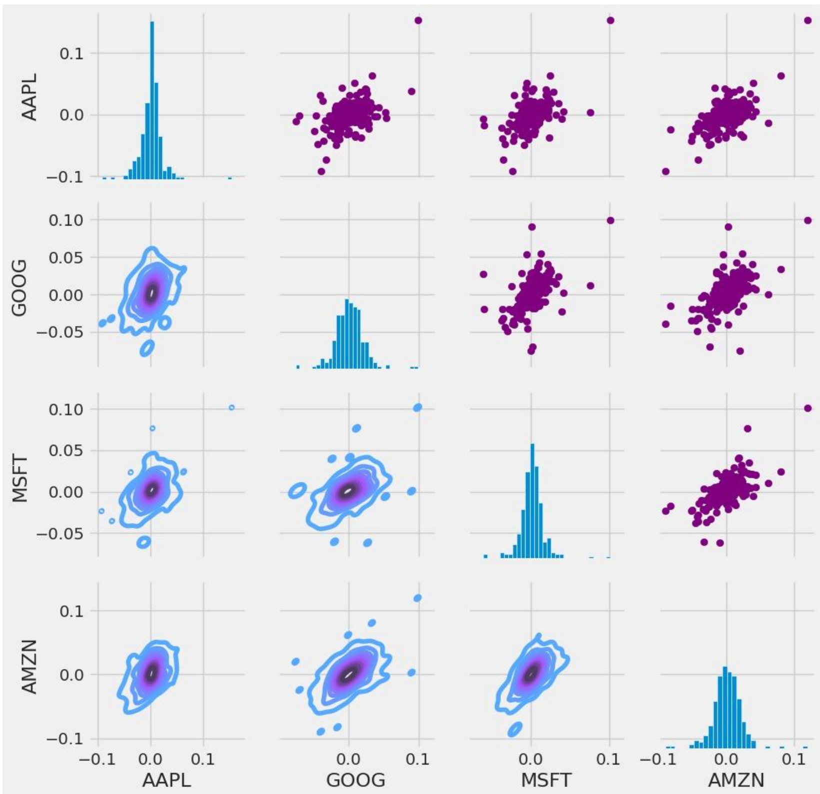

Figure 8: Correlation heatmap 2

PAIRWISE RELATIONSHIPS

Pairplots and scatterplots with regression lines show pairwise relationships between stock returns. These plots highlight linear trends and outliers between pairs of stocks.

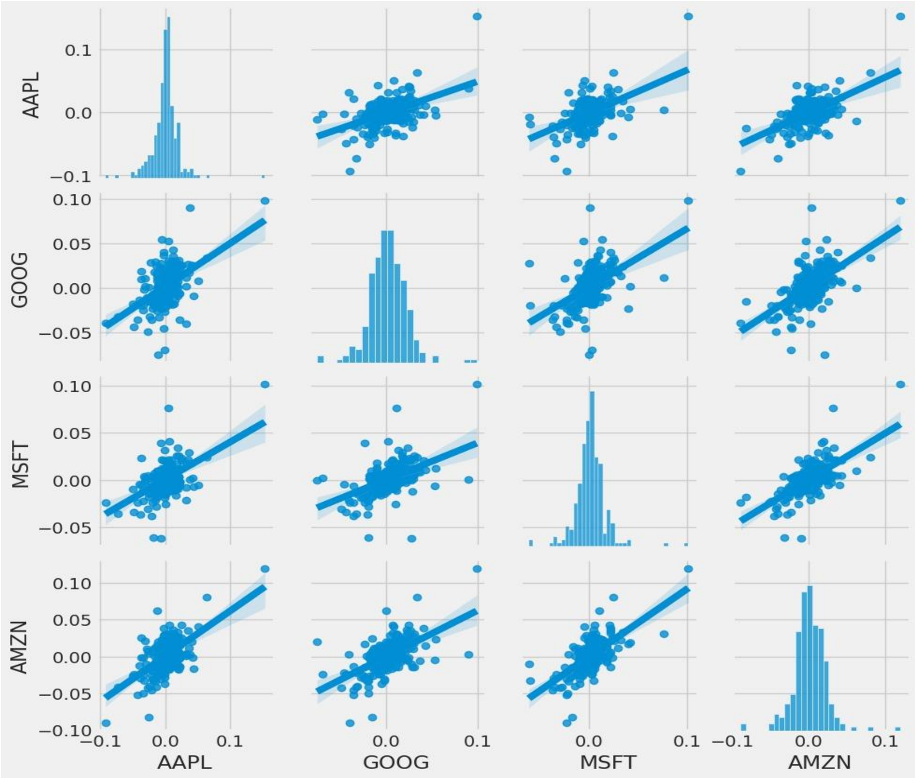

Figure 9: Pairwise scatter/pairplot 1

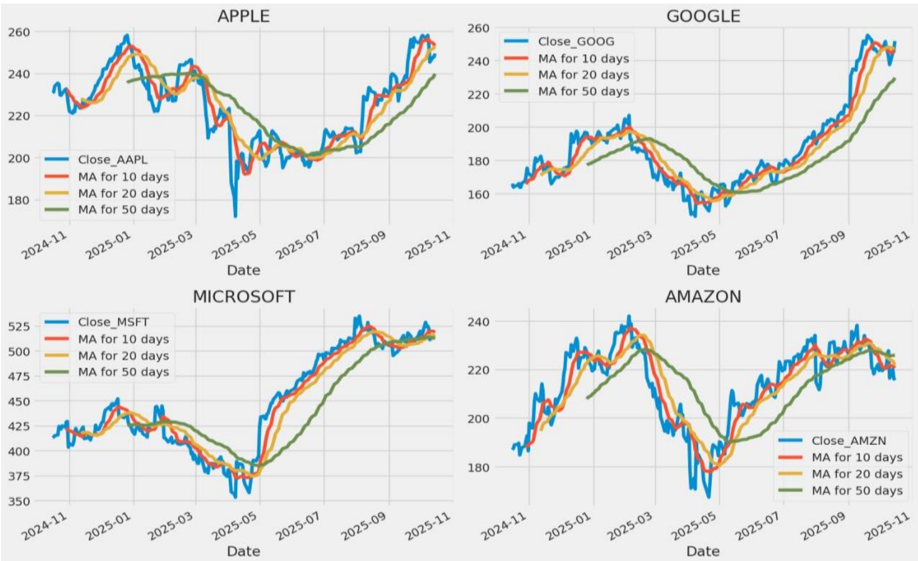

Figure 10: Pairwise scatter/pairplot 2

Figure 11: Pairwise scatter/pairplot 3

LSTM MODEL ARCHITECTURE

The LSTM model used in this study is a sequential network composed of:

i. An input layer accepting sequences of historical scaled closing prices (look-back window of 60 days).
ii. One or two stacked LSTM layers with dropout for regularization (dropout rate 0.2).
iii. A dense output layer producing the next-step closing price prediction.

The model was trained using mean squared error (MSE) as the loss function and Adam optimizer. Early stopping on validation loss was used to avoid overfitting.

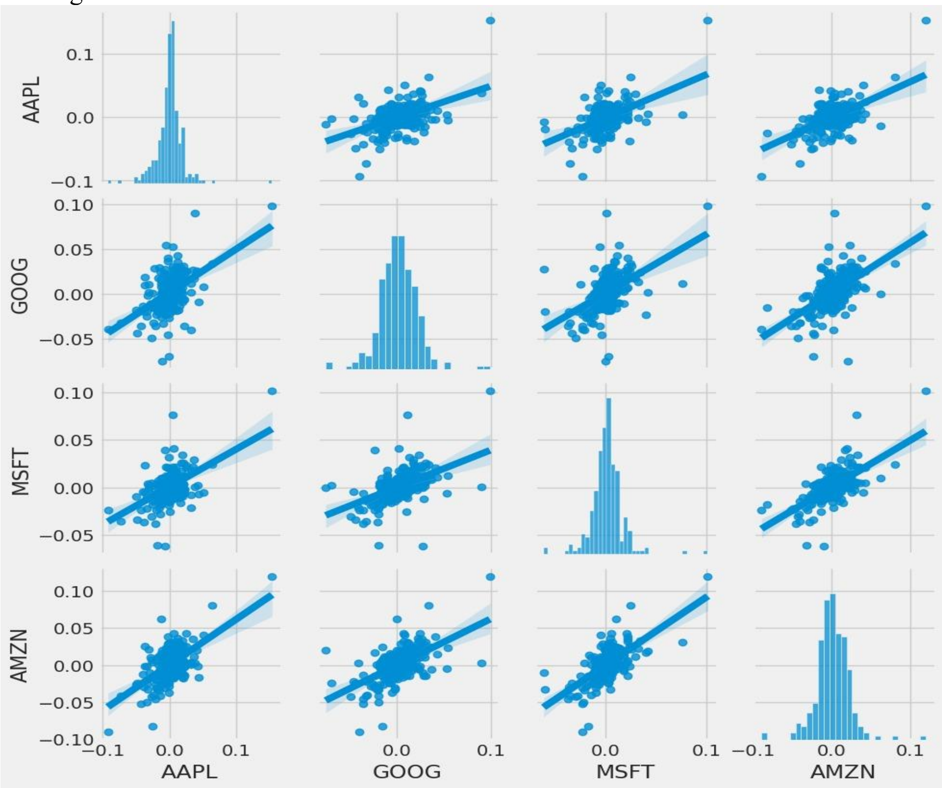

Figure 12: Model training and prediction overlay (train, validation, predictions)

RESULTS

The LSTM model was applied to the Apple closing price series as a case study. Training data covered the long-term historical series, validation covered the most recent months, and short-term forecasts extend beyond the validation window. The model captured the overall upward trend and produced short-horizon predictions that align with validation trends. Key evaluation metrics (computed on the validation/test split) include:

i. Mean Squared Error (MSE): reported on the scaled series and back-transformed to price units for interpretation.
ii. Root Mean Squared Error (RMSE): provides an interpretable error in price units.
iii. Visual inspection: prediction curves closely follow the validation series, though short-term deviations occur around high-volatility events.

Figure 13: Model result 1

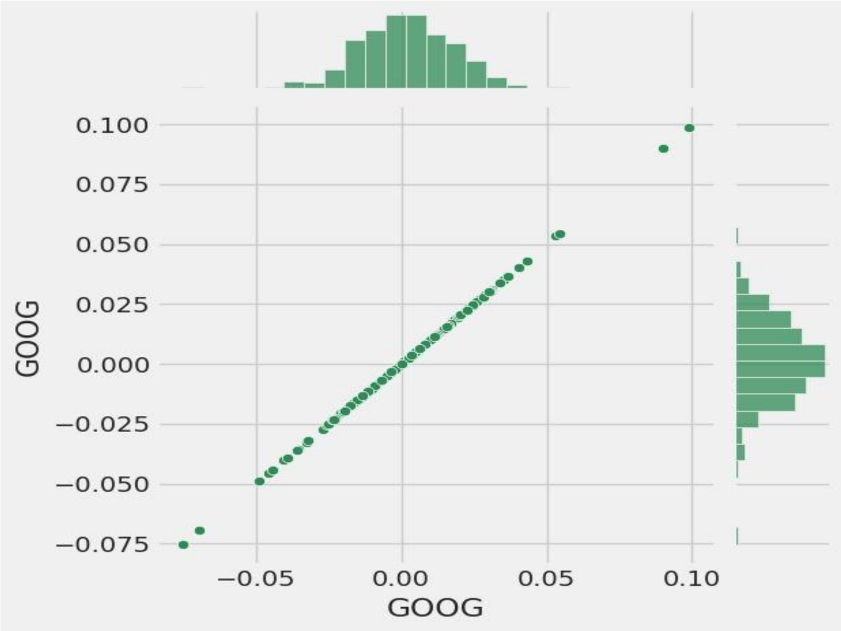

Figure 14: Model result 2

CHALLENGES AND LIMITATIONS

Several factors limit the predictive power of LSTM in this context:

i. Market volatility and unforeseen events can cause sudden price changes that are hard to predict.
ii. Financial data is noisy; small sample artifacts or structural breaks reduce forecast accuracy.
iii. LSTM models require careful hyperparameter tuning and significant computational resources.

Possible improvements include augmenting inputs with sentiment or news data, using hybrid models (LSTM combined with CNN or attention mechanisms), and conducting extensive hyperparameter searches.

CONCLUSION

This study demonstrates that LSTM networks can model general price trends and provide short-term forecasts for major technology stocks. While LSTM forecasts are not immune to market volatility, they are a useful tool in a broader forecasting toolkit. Future work should incorporate additional data sources (news, macro indicators) and evaluate hybrid architectures to improve robustness.

REFERENCES

[1] Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. Neural Computation, 9(8), 1735–1780.
[2] Chen, K., Zhou, Y., & Dai, F. (2015). A LSTM-based method for stock price prediction. Expert Systems with Applications, 42(8), 3903-3910.
[3] Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.
[4] Yahoo Finance API Documentation: https://www.yfinance.com