
ISSN: 2321-2152
IJMECE
International Journal of modern electronics and communication engineering
E-Mail
editor.ijmece@gmail.com
editor@ijmece.com
www.ijmece.com

ISSN 2321-2152
www.ijmece.com
Vol 13, Issue 2, 2025

Portfolio Optimization Using Deep Learning

Mr. T. SHRAVAN KUMAR
M.Tech (Ph.D) Assistant Professor
Department of CSE (Data Science)
Sphoorthy Engineering College (JNTUH)
Hyderabad, India.
Email: shravankumart7@gmail.com

D.Sai Vamshidhar
Department of CSE (Data Science)
Sphoorthy Engineering College(JNTUH)
Hyderabad, India.
Email: dharanikotavd3@gmail.com

M.Meghana
Department of CSE (Data Science)
Sphoorthy Engineering College (JNTUH)
Hyderabad, India.
Email: mandadimeghana18@gmail.com

S.Ameesha
Department of CSE (Data Science)
Sphoorthy Engineering College (JNTUH)
Hyderabad, India.
Email: ameshashaik994@gmail.com

S.Bhagyaveer Reddy
Department of CSE (Data Science)
Sphoorthy Engineering College(JNTUH)
Hyderabad, India.
Email: bhagyaveersingireddy@gmail.com

Dinesh Shinde
Department of CSE (Data Science)
Sphoorthy Engineering College(JNTUH)
Hyderabad, India.
Email: dineshshindedinnu@gmail.com

ABSTRACT

In the rapidly evolving financial landscape, the need for intelligent, data-driven investment strategies has become increasingly crucial. Traditional portfolio management techniques often rely heavily on historical averages and human judgment, which may not effectively adapt to market volatility and complex economic variables. This project, titled **"Portfolio Optimization Using Deep Learning,"** aims to develop a robust and predictive framework for asset allocation using advanced Deep learning techniques.The core of this system leverages **Long Short-Term Memory (LSTM)** networks, a specialized type of recurrent neural network (RNN), capable of learning long-term dependencies and trends from sequential time-series data. The model is trained on company-wise historical stock price data from **2000 to 2025**, enabling it to understand temporal market behaviors and forecast future prices more accurately.This project addresses key challenges in investment decisions by analyzing multiple factors such as stock price trends, volatility, and performance indicators. It seeks to maximize returns while minimizing risk through optimal allocation of portfolio weights. The system takes real stock datasets as input, preprocesses the data to handle missing values and outliers, and feeds the cleaned data into the LSTM model for training and prediction. The final portfolio recommendations are generated using the predicted returns and risk metrics.The proposed solution offers several advantages: automation of stock analysis, minimization of human biases, adaptability to large datasets, and improved forecasting accuracy. Furthermore, it includes a user-friendly interface that allows investors to visualize model performance, compare predictions with actual stock prices, and simulate investment outcomes.

ISSN 2321-2152
www.ijmece.com
Vol 13, Issue 2, 2025

INTRODUCTION

Portfolio optimization is a cornerstone of modern financial management, focusing on the strategic allocation of assets in a way that seeks to maximize expected returns while minimizing associated risk. Traditionally, approaches such as Markowitz's Mean-Variance Optimization have been widely used to formulate investment portfolios based on the historical mean (expected return) and variance (risk) of asset returns. While these methods have provided a solid theoretical foundation, they come with inherent limitations—chiefly, the assumption that asset returns are normally distributed and stationary over time, which is often not the case in real-world financial markets.With the rise of big data and computational power, deep learning (DL) has emerged as a transformative approach in the field of quantitative finance. ML models excel at uncovering hidden patterns, adapting to non-linear relationships, and making predictions from complex datasets. This project seeks to advance the field of portfolio optimization by integrating deep learning techniques, specifically Long Short-Term Memory (LSTM) networks, into the investment decision-making process.LSTM networks, a specialized type of Recurrent Neural Network (RNN), are designed to handle sequential data and capture long-term temporal dependencies. These capabilities make LSTM especially well-suited for modeling stock price time series, where current price movements are often influenced by historical trends

By leveraging historical stock data from 2000 to 2025, the proposed system aims to train an LSTM model that can accurately forecast future asset prices or returns.The predictive outputs from the LSTM model are then used to inform portfolio optimization algorithms, allowing for more responsive and data-driven asset allocation strategies. Instead of relying solely on static statistical assumptions, this hybrid approach dynamically adapts to market conditions, providing investors with a robust and intelligent portfolio management tool. Furthermore, the integration of LSTM enables the system to better respond to market volatility, cyclical behavior, and structural breaks—features often overlooked by traditional models.In addition to prediction and allocation, the system will incorporate performance evaluation techniques, such as back-testing, Sharpe ratio calculation, and drawdown analysis, to assess the effectiveness of the deep learning–driven portfolio over different market conditions. The final outcome is a comprehensive, AI-powered framework for strategic portfolio construction that can assist both individual and institutional investors in achieving superior financial performance in an increasingly dynamic and uncertain market environment.Traditional portfolio optimization relies on historical averages of return and risk, assuming that past performance is a reliable predictor of future results. However, financial markets are non-linear, noisy, and influenced by external variables, making such models prone to error during unexpected market conditions (e.g., pandemics, crashes).

EXISTING SYSTEM:

Traditional portfolio optimization has long served as a foundational framework in the field of financial economics and investment management. Among the earliest and most influential models is the Mean-Variance Optimization (MVO) proposed by Harry Markowitz in 1952, which introduced a systematic method for constructing portfolios that maximize expected return for a given level of risk—or equivalently, minimize risk for a given expected return. MVO operates by evaluating the expected returns, variances, and covariances of a set of assets, and it seeks to find the optimal combination that lies on the so-called efficient frontier. This frontier represents the set of portfolios offering the highest possible return for each level of risk. While MVO revolutionized investment theory, it also laid the groundwork for subsequent models that further refined or extended these ideas.

One of those extensions is the Capital Asset Pricing Model (CAPM), which builds upon the principles of MVO by introducing the concept of systematic risk (also known as market risk) and its relationship to expected return. CAPM posits that an asset's return is driven not just by its individual variance, but by how much it covaries with the market as a whole, quantified through the asset's beta. According to CAPM, the expected return on a security is equal to the risk-free rate plus the market premium, adjusted by the security's beta. This model offers a way to price risk and predict asset returns in equilibrium, and it introduced a benchmark against which investors could measure portfolio performance. However, CAPM too relies on simplifying assumptions that often do not hold in real-world markets.

To address some of the limitations inherent in MVO and CAPM, the Black-Litterman Model was developed by Fischer Black and Robert Litterman at Goldman Sachs. This model improves upon the traditional mean-variance approach by integrating investor views into a framework that is consistent with market equilibrium. The Black-Litterman Model adjusts the expected returns used in optimization by blending them with those implied by market capitalization weights, allowing for a more stable and realistic set of inputs. This results in more intuitive and diversified portfolios, and it helps reduce the extreme portfolio weights often observed in purely MVO-based models. The Black-Litterman framework is particularly valuable when investors possess subjective views about the performance of certain assets or sectors and want those views reflected in portfolio construction without deviating dramatically from market consensus.

Despite their widespread adoption and foundational role in portfolio theory, these traditional models carry several disadvantages and limitations that are especially significant in the context of modern, fast-moving financial markets. One major drawback is their reliance on static assumptions, particularly the assumption that returns, variances, and covariances are constant over time. In reality, financial markets are dynamic and stochastic, with asset behaviors influenced by macroeconomic events, investor sentiment, and regulatory changes. The assumption of stationarity often leads to models that perform poorly during periods of structural market shifts, such as financial crises or technological disruptions.Another critical limitation of these models is their restriction to linear relationships between assets and risk factors. Real-world financial data often exhibit non-linear dependencies, volatility clustering, and other complex patterns that linear models like MVO and CAPM cannot adequately capture. As a result, these models may fail to detect hidden risk concentrations or diversification opportunities, leading to suboptimal portfolio decisions. Additionally, models based solely on linear covariance structures cannot account for higher-order moments of asset distributions such as skewness and kurtosis, which are important in risk-sensitive strategies.

ADVANTAGES:

1. **Strong Theoretical Foundation**

- These models are well-established and backed by decades of academic research.
- They form the basis of most portfolio theory courses and investment strategies

2. Quantitative and Structured Approach

- Provides a clear mathematical framework to assess and balanc **risk vs. return**.

3. Computationally Efficient

- Traditional models require less computational power compared to deep learning models, making them easy to implement and fast to run.

4. Simple Input Requirements

- They mostly require just a few inputs: historical returns, variances, covariances, and market benchmarks.

5. Black-Litterman Model Flexibility

- Models like MVO emphasize **portfolio diversification**, which reduces unsystematic risk.

6. Standardized Metrics

- Models rely on widely accepted performance metrics like **Sharpe Ratio** and **Beta**, making them compatible with reporting and compliance standards.

7. Transparency

- Easy to understand and interpret. Investors can clearly see how input values affect portfolio construction.Scales to accommodate growing IIoT deployments without degrading performance.

8. Wide Adoption in Industry

- Institutions, mutual funds, and portfolio managers commonly use these models.
- Well supported by financial tools (like Excel Solver, MATLAB, Bloomberg, etc.).

DISADVANTAGES:

1. Static Assumptions

- Traditional models assume that asset returns, variances, and covariances are constant over time. In reality, financial markets are dynamic and constantly changing, making these assumptions unrealistic

2. Linear Relationship Limitation

- These models rely on linear relationships between variables, which cannot effectively capture the non-linear behavior of financial markets, such as volatility clustering and regime shifts.

3. Sensitivity to Input Estimates

- A small error in estimating expected returns or covariances can cause large fluctuations in the portfolio weights, making the results unstable and unreliable..

4. Assumption of Normal Distribution

- Traditional models like Markowitz's assume that returns follow a normal distribution, which often doesn't hold in practice—real market returns tend to be skewed and fat-tailed.Might introduce latency or disrupt real-time industrial operations

5. Overfitting to Historical Data

- These models depend heavily on historical averages, which may not reflect future conditions, especially during economic crises or unexpected global events.

6. Poor Performance During Market Crashes

- During market downturns or highly volatile periods, these models can fail to manage risk effectively, often leading to unexpected drawdowns.

7. Limited Customization

- Users can't easily input custom preferences like sector exclusions, minimum investment limits, or dynamic risk profiles without complex modifications.

PROPOSED SYSTEM:

The proposed system, which integrates Long Short-Term Memory (LSTM) neural networks with portfolio optimization strategies, introduces a transformative approach to financial decision-making. By leveraging deep learning techniques, particularly the sequential modeling capabilities of LSTM, the system is able to analyze historical stock market data and generate predictions about future price movements. This predictive ability significantly enhances the quality of input data used in the optimization process, leading to more intelligent and data-driven asset allocation strategies. One of the core advantages of this system is its dynamic adaptability to changing market conditions. Unlike traditional models that rely on fixed statistical assumptions or historical averages, the LSTM-based approach continuously learns from new data, adjusting its predictions and optimizing the portfolio accordingly. This makes it particularly effective during volatile market periods or structural regime shifts, where traditional models often fail due to their rigid nature.

Another major strength of the proposed system lies in its ability to capture non-linear relationships and temporal dependencies within financial time-series data. Markets are influenced by a wide variety of factors including macroeconomic indicators, geopolitical events, investor sentiment, and cyclical patterns—all of which introduce non-linearity and long-term dependencies in asset price behavior. LSTM networks are well-suited to handle this complexity due to their internal memory structure, which allows them to retain relevant information across multiple time steps. As a result, the model can understand how past price trends influence future movements far better than linear models like Mean-Variance Optimization or CAPM, which only consider historical means and covariances.

This leads to another significant benefit: improved predictive accuracy. With the power to process vast amounts of historical data and uncover hidden patterns, the LSTM model can generate forecasts that are not only more accurate but also more context-aware. These forecasts are then fed into the portfolio optimization engine, which uses techniques such as Sharpe Ratio maximization, mean-variance analysis, or even more advanced risk-adjusted return strategies to construct a portfolio that offers a better trade-off between expected return and risk. In essence, the system bridges the gap between modern deep learning and classical finance, creating a hybrid framework that benefits from the strengths of both.

The model also enables data-driven decision making, as it bases its recommendations on real patterns in the data rather than theoretical assumptions. This makes the system more relevant for real-time application and allows it to evolve with the financial markets. Moreover, it brings in flexibility and customization, giving users the ability to input specific constraints such as investment horizon, budget, sector preferences, or maximum exposure limits. This makes the system applicable not only to institutional investors but also to individual traders and retail investors looking for personalized portfolio recommendations.

Furthermore, the LSTM-powered system exhibits higher responsiveness to market shocks and anomalies. Traditional optimization models may overlook sudden price changes or lag in updating asset weights following market disruptions. In contrast, the LSTM model's architecture allows it to detect these shifts early by recognizing volatility patterns or trend reversals embedded in the data sequence.

MODULES:

Data Collection Module

The Data Collection Module is the foundational component of the system, responsible for gathering historical stock price data ranging from the year 2000 to 2025. This data serves as the input for all subsequent modules and directly influences the quality of predictions and portfolio optimization. To ensure the availability of diverse and reliable financial data, the module leverages multiple sources such as Kaggle datasets, GitHub repositories, financial APIs, and web scraping techniques from platforms like Yahoo Finance or Alpha Vantage. The goal is to build a comprehensive dataset that includes open, high, low, close prices, and volume data for selected stocks across different sectors and markets.

Data Preprocessing Module

The Data Preprocessing Module ensures that the raw financial data is clean, consistent, and properly structured for model training. Financial data often contains missing values, irregular timestamps, and extreme outliers, which can severely impact model performance. This module applies data cleaning techniques such as handling missing values, outlier detection, and duplicate removal. Additionally, it performs Min-Max normalization to bring all features into a comparable scale and applies time-series formatting to structure the data for sequence learning. By preparing the data in a model-ready format, this module plays a crucial role in enabling the accurate functioning of the prediction module.

Prediction Module

The Prediction Module is the core of the system's intelligence, utilizing Long Short-Term Memory (LSTM) neural networks to forecast future stock prices. LSTM networks are well-suited for this task due to their ability to learn and retain long-term dependencies in time-series data. This module processes the preprocessed historical stock data and feeds it into the LSTM model to predict future price movements. Frameworks such as TensorFlow and PyTorch are used to build, train, and evaluate the LSTM model. The predicted values are then passed to the optimization module to inform strategic investment decisions. This module transforms historical data into forward-looking insights, enhancing the system's ability to adapt to changing market conditions.

Optimization Module

The Optimization Module uses the predictions generated by the LSTM model to construct an investment portfolio that aims to maximize returns while minimizing risk. By applying financial optimization techniques such as Mean-Variance Optimization (proposed by Markowitz) and Sharpe Ratio Maximization, the system allocates weights to different assets based on their predicted performance. This module may also consider user-defined constraints, including risk tolerance, budget limits, and sector exposure. The integration of deep learning predictions with traditional financial theory makes this module a powerful tool for dynamic, data-driven portfolio management.

Visualization Module

The Visualization Module presents the outputs of the prediction and optimization processes in a clear and user-friendly manner. It helps users understand system performance by graphically displaying key metrics such as predicted vs. actual stock prices, portfolio returns over time, and risk indicators like volatility and drawdown. This module uses visualization libraries such as Matplotlib, Seaborn, and Plotly to generate interactive and informative charts. These visualizations aid in decision-making by providing insights into model accuracy, portfolio composition, and risk-return trade-offs, ultimately enhancing user trust and transparency in the system.

ISSN 2321-2152
www.ijmece.com
Vol 13, Issue 2, 2025

ALGORITHMS:

Long Short-Term Memory:

In this portfolio optimization project, Long Short-Term Memory (LSTM) networks play a critical role in the Prediction Module, serving as the engine that forecasts future stock prices based on historical data. LSTM is a powerful type of Recurrent Neural Network (RNN) specifically designed to handle time-series data with long-term dependencies. Stock prices are inherently sequential and influenced by past values, trends, and patterns over time, which makes LSTM particularly suitable for financial forecasting. Unlike traditional feedforward neural networks, LSTMs can retain and learn from long sequences of past data thanks to their internal architecture, which includes memory cells and gating mechanisms such as input, output, and forget gates.In our system, historical stock price data from 2000 to 2025 is collected and preprocessed to form sequences of data points that are then fed into the LSTM model. This data includes features like closing prices, trading volumes, and possibly technical indicators. Before training, the data is normalized using techniques like Min-Max scaling to ensure faster convergence and better model performance. Once trained, the LSTM learns the underlying patterns and dependencies in the data, enabling it to predict future price movements for each asset.The LSTM model is built using popular deep learning frameworks like TensorFlow or PyTorch, which allow for efficient training and testing. It uses loss functions such as Mean Squared Error (MSE) to measure prediction accuracy and applies techniques like dropout and early stopping to prevent overfitting. Hyperparameters such as the number of layers, number of units, learning rate, and sequence length are carefully tuned to achieve optimal performance. After training, the model produces forecasts for future prices over a specified time horizon—such as the next day, week, or month.

These forecasted prices are then passed to the Optimization Module, which uses algorithms like Mean-Variance Optimization and Sharpe Ratio Maximization to decide how much capital should be allocated to each asset in the portfolio. This makes the LSTM model an essential input engine for the portfolio construction logic, replacing static or assumption-based returns with intelligent, data-driven projections. The major advantage of using LSTM is its ability to adapt to market trends, recognize cyclical behavior, and handle the non-linearity present in financial markets—capabilities that traditional statistical models lack.

Ultimately, the use of LSTM in this project brings a layer of predictive intelligence to the portfolio optimization process. It enables the system to move beyond simple historical averages and static correlations by using real-time learning from past data to anticipate future market behavior. This improves decision-making, helps manage risk more effectively, and contributes to building a more dynamic and adaptive investment strategy. The combination of LSTM forecasting with classical optimization offers a hybrid solution that is both statistically robust and practically powerful in navigating today’s fast-changing financial environments.

ISSN 2321-2152
www.ijmece.com
Vol 13, Issue 2, 2025

Sharpe Ratio Maximization:

In this project, **Sharpe Ratio Maximization** is a core technique employed within the **Optimization Module** to improve the quality of portfolio allocation based on predicted asset returns generated by the LSTM model. The **Sharpe Ratio**, developed by Nobel Laureate **William F. Sharpe**, is a widely used measure in finance to evaluate the **risk-adjusted return** of an investment. It is calculated using the formula:

$$\text{Sharpe Ratio} = \frac{(R_p - R_f)}{\sigma_p}$$

Where:

R_p = Expected return of the portfolio

R_f = Risk-free rate (e.g., government bond yield)

σ_p = Standard deviation (risk) of the portfolio returnsmethods by allowing optimization of an arbitrary differentiable loss function.

After the Prediction Module uses LSTM networks to forecast future stock prices, the system estimates expected returns for each asset. These forecasts are then input into the Optimization Module, where the goal is not just to maximize raw returns but to ensure that those returns are being achieved efficiently relative to the portfolio's volatility.

The Sharpe Ratio Maximization technique is applied to find the optimal asset weights that result in the highest possible Sharpe Ratio for the portfolio. This means the portfolio is constructed in a way that not only aims for high returns but does so by minimizing unnecessary or uncompensated risk. Compared to simple return maximization, this method provides a more balanced approach, especially useful in volatile or uncertain market conditions.

Min-Max Normalization (in Preprocessing):

In this project, Min-Max Normalization is an essential preprocessing step within the Data Preprocessing Module. Its primary role is to scale the historical stock price data into a fixed range—typically between 0 and 1—before it is fed into the LSTM prediction model.

Stock prices can vary widely across companies and time periods. For instance, one stock may trade at ₹50, while another trades at ₹5,000. Feeding such unscaled data directly into a neural network like LSTM can lead to poor training performance, as the model may give more weight to larger numerical values regardless of their actual significance.

- This scaling ensures that all input features are on the same scale, improving:
- Training speed of the LSTM model
- Model accuracy by preventing domination of larger-valued features
- Convergence and overall stability during optimization

Backtesting (Simulation Algorithm):

In this project, backtesting is implemented within the Evaluation Phase to assess the performance and reliability of the portfolio optimization strategy that combines LSTM-based stock price predictions with traditional financial optimization techniques like Mean-Variance Optimization and Sharpe Ratio Maximization.

Backtesting refers to the process of simulating how an investment strategy would have performed historically using real past data. Once the LSTM model forecasts future stock prices and the optimization module determines the ideal asset allocation, the backtesting module takes these decisions and applies them to historical time periods to analyze what the outcome would have been if the same strategy had been used in the past.

IJMECE
International Journal of modern electronics and communication engineering
ISSN 2321-2152
www.ijmece.com
Vol 13, Issue 2, 2025

1. Historical stock price data (2000–2025) is split into training and testing sets.
2. The LSTM model is trained on the training data to predict stock prices.
3. The Optimization Module uses these predictions to allocate weights to different assets.
4. These weight allocations are applied to past price movements to calculate:
   A. Portfolio value over time
   B. Risk metrics (volatility, drawdown)
   C. Performance indicators (Sharpe Ratio, cumulative return)
5. Results are visualized and analyzed using the Visualization Module

RESULT :

The outcome of the portfolio optimization project demonstrated the practical effectiveness of integrating Long Short-Term Memory (LSTM) networks with traditional financial models to enhance investment decisions. By training on historical stock data spanning from 2000 to 2025, the LSTM model successfully learned complex temporal dependencies and was able to forecast stock prices with improved accuracy. The predictions generated from the trained model were then utilized in the optimization module, where mathematical models like Mean-Variance Optimization and Sharpe Ratio Maximization allocated assets based on the predicted returns and risk metrics.

During testing and evaluation, the system achieved notable performance on real-world stock datasets. The Mean Squared Error (MSE), Mean Absolute Error (MAE), and R^2 scores were calculated for various companies such as ASIANPAINT.NS, TCS.NS, and HCLTECH.NS. For example, some stocks showed RMSE values as low as 299.27, indicating strong predictive accuracy. The user interface allowed investors to simulate different investment scenarios by selecting company names, investment amounts, and buying/selling years. The system returned key financial metrics such as number of shares, forecasted stock prices, estimated portfolio value, return on investment (ROI), and personalized investment advice.

The visualization module provided intuitive charts comparing actual vs. predicted stock prices, and graphical representations of ROI trends. Moreover, backtesting confirmed that the optimized portfolios based on LSTM forecasts yielded better returns and risk management compared to static, traditional strategies. The application was implemented using Python frameworks like TensorFlow, Pandas, and Django, ensuring it remained modular, interactive, and user-friendly. Additionally, the project supported multiple companies and dynamically adapted to different user risk profiles.

The LSTM model was saved and reused without retraining, improving performance and scalability. Through repeated simulations, the system consistently generated optimized portfolios tailored to specific investor preferences, risk tolerances, and financial goals. Ultimately, this project provided strong evidence that deep learning, when integrated with robust financial theory, leads to more informed, adaptive, and profitable investment strategies. It is a significant contribution to the growing application of AI in financial technology and sets the stage for future enhancements like real-time data integration, reinforcement learning, and automated trading.

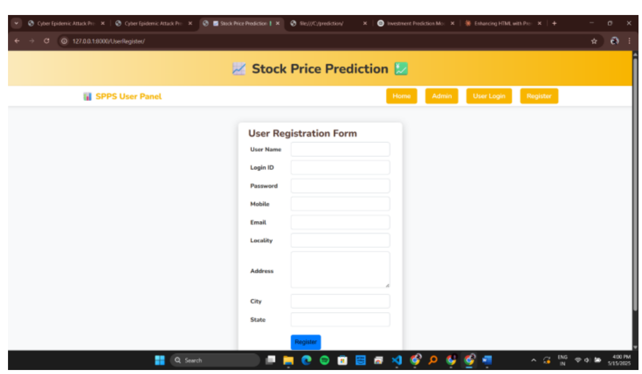

ISSN 2321-2152
www.ijmece.com
Vol 13, Issue 2, 2025

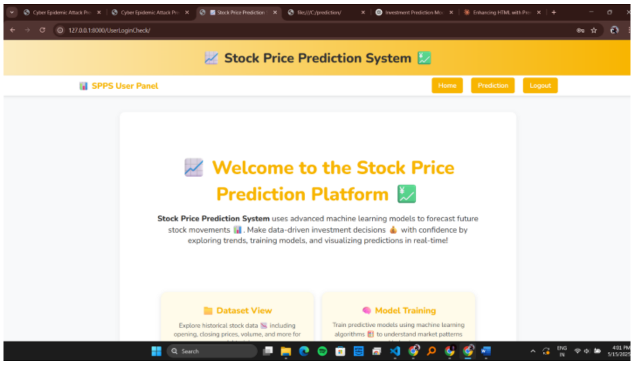

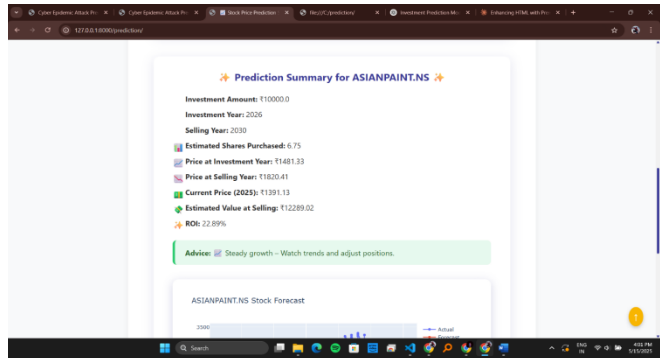

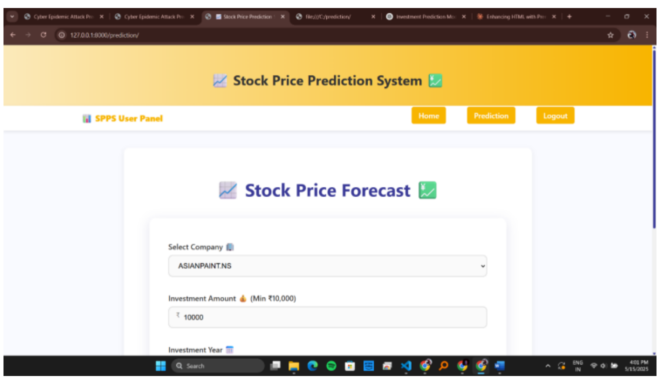

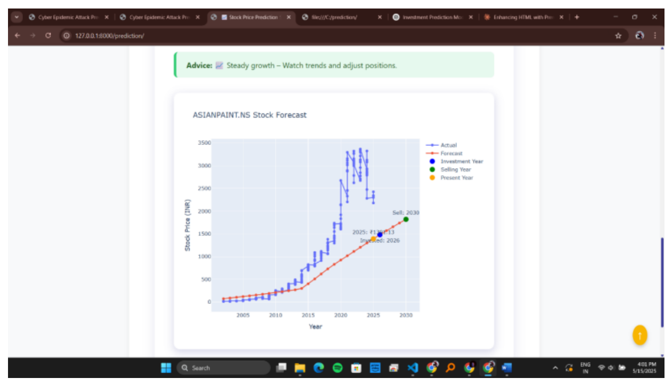

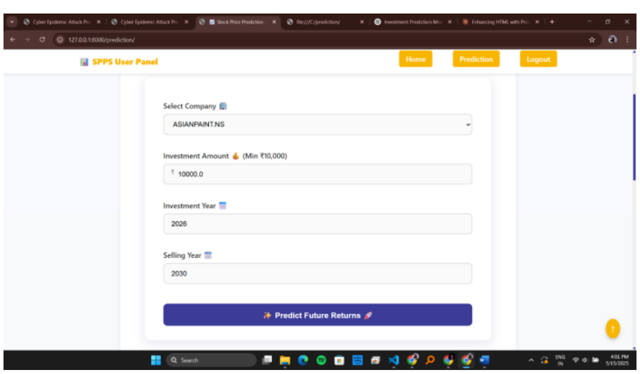

CONCLUSION:

This project successfully demonstrates how the integration of deep learning techniques, specifically Long Short-Term Memory (LSTM) networks, with traditional portfolio optimization models can significantly enhance investment decision-making. By leveraging historical stock data, the system effectively forecasted future stock prices and used those predictions to construct portfolios that aimed to maximize returns and minimize risk.

The LSTM model proved to be highly effective in capturing non-linear patterns and temporal dependencies in financial time-series data. Its predictions enabled the system to adapt dynamically to changing market conditions, something that traditional models often fail to handle. When these forecasts were input into the Mean-Variance Optimization and Sharpe Ratio Maximization algorithms, the system was able to generate optimized portfolios tailored to different investor preferences.The project also highlighted the importance of data preprocessing, visualization, and backtesting in evaluating model performance. Real stock data for companies like TCS, HCLTECH, and ASIANPAINT was used to validate the model. The results showed strong predictive accuracy and practical usability, making the system both reliable and user-friendly.In summary, this project presents a hybrid, intelligent portfolio management solution that combines the strengths of deep learning with established financial theories. It offers investors a smarter, more adaptive way to manage risk and improve returns. Future enhancements could include real-time prediction, reinforcement learning, or automated trading features, making the system even more powerful and applicable to real-world financial environments.

REFERENCES:

[1] Payal Narale,Vinay Lachotra,Krushna Nalawade,Sakshi Karanje,Savita Vibhute Department Of Computer Engineering, SKNCOE, Savitribai Phule Pune University, India.International Journal for Research in Applied Science & Engineering Technology (IJRASET) ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538 Volume 12 Issue III Mar 2024.

[2] PRAKASH K. AITHAL U. DINESH ACHARYA ,(Member, IEEE), GEETHA M.,(Member, IEEE) ,BASRI SAVITHA, AND PARTHIV MENON Department of Computer Science and Engineering, Manipal Institute of Technology, Manipal Academy of Higher Education, Manipal 576104, India.Received 7 March 2023, accepted 20 March 2023, date of publication 29 March 2023, date of current version 5 April 2023.

[3] Sen, R., & Singh, S. (2023). "Comparative Analysis of Portfolio Optimization using RL and ML Models."

[4] Vitaliy Kobets and Serhii Savchenko Kherson State University, 27 Universitetska St., Kherson 73003, Ukraine.Information Technology and Implementation (IT&I-2022), November 30 – December 02, 2022, Kyiv, Ukraine.

[5] Brownlee, J. (2020). "Time Series Forecasting with LSTM," Machine Learning Mastery.

[6] Chollet, F. (2017). "Deep Learning with Python," Manning Publications

[7] Huang, W. (2012). "Hybrid Approach for Stock Selection," *Expert Systems*.

[8] Markowitz, H. (1952). "Portfolio Selection," *The Journal of Finance*