LSTM-based Portfolio Optimization Strategy

Xiaotong Li *

Glorious Sun School of Business and Management, Donghua University, Shanghai, China

* Corresponding Author Email: 220613105@mail.dhu.edu.cn

Abstract. In the highly competitive and constantly evolving landscape of the investment realm, the construction of an optimal investment portfolio has emerged as a cornerstone for investors aiming to adeptly navigate the complex interplay between risks and returns. With the rapid advancement of technology, the performance and trends of stocks have become increasingly intertwined with technological developments. This study analyzes five tech-dependent stocks and applies the Long Short-Term Memory (LSTM) model for price prediction. The combination of LSTM and Mean-Variance Optimization (MVO) model outputs daily portfolio weights using two strategies: minimizing variance and maximizing Sharpe ratio. The minimal variance model produces a cumulative return of 193.1164% during a 30-day testing period from day 71 to day 100; the maximum Sharpe ratio model produces a return of 61.2246%, both exceeding the return of the CSI 300 Index at just 14.6193%. The outcomes validate the effectiveness of the suggested approach and provide insightful analysis for portfolio managers choosing their course of action.

Keywords: Long short-term memory, Portfolio optimization, Mean-variance, Investment Strategy.

1. Introduction

Over the years, considerable study has gone toward portfolio building. A pillar of financial strategy, building an ideal investment portfolio helps investors to properly control risks and maximize profits [1]. A well-considered portfolio not only reduces possible losses in market downturns but also uses development possibilities to reach long-term financial goals.

Significant study throughout the years has advanced portfolio building techniques. Introduced in 1952 [2], Markowitz's revolutionary Mean-variance Optimization (MVO) model created the foundation for contemporary portfolio theory by balancing risk and return via asset weight optimization. Recent developments build on this basis by using advanced methods including clever approaches to improve prediction precision. For stock price index predictions, Kim and Han, for example, used evolutionary algorithms housed within artificial neural networks [3]. Chen and Lin paired feature selection techniques [4] with Support Vector Machines (SVM). Ballings et al. evaluated a number of classifiers for direction of stock price prediction [5]. Based on random forest [6], Deng, Li, and Wen propose an original method for stock price prediction. Using machine learning techniques, Bhardwaj and Tiwari investigated stock price prediction [7].

However, with regard to the application of LSTM, few studies aim to integrate LSTM into the portfolio optimization process and the viability of this approach. The majority of the prior studies are put emphasis on the applicability and superiority of LSTM on stock prediction. For example, Roondiwala et al. addressed accuracy [8] and used LSTM to forecast market indices. Although they deployed LSTM networks on portfolios composed of SP500 members, Fischer & Krauss largely concentrated on the differences with other memoryless machine learning models as logistic regression and random forest [9]. Some academics regarded classical forecasting as a classification problem and broke creatively from conventional wisdom. For example, some studies have focused on applying machine learning techniques to predict stock market value trajectory [10], including neural networks and tree-like classifiers [11], rather than a regression problem of predicting closing price itself.

This paper connects LSTM with portfolio management and makes the following contributions. Firstly, the daily closing price data of five stocks within a specific time period will be collected. Following outlier corrections and integrity checks on these data, they will be split correspondingly into test and training sets. Following that, the LSTM model will be applied to record, on the training set for prediction, the long-term dependency of stock prices. The mean-variance optimization model will then be fed the projected data coupled with historical return data, risk-free rate, and estimated covariance matrix. The returns and hazards of the investment portfolio will be computed step by step using a sliding window approach; subsequently, the asset weight allocation will be found. Lastly, by means of a comparison between the daily and cumulative returns of the investment portfolio under several models with the CSI 300 Index, the performance and efficiency of this investment portfolio optimization strategy will be fully assessed, so aiming to offer a more valuable approach for investors to build and control their portfolios in the face of market uncertainty.

2. Data

This study concentrates on five stocks, namely Heertai (002402), Hundsun Electronics (600570), Great Wall of China (000066), Eastmoney (300059), and Shanghai Electric (601727). Their daily closing prices were collected from Investing.com (https://www.investing.com) for the period from June 26, 2024, to November 21, 2024, encompassing 100 trading days.

The selection of these stocks was guided by multiple factors. They are all in industries influenced by technology, with Hundsun Electronics and Eastmoney being prominent in the fintech space, Heertai in intelligent control technology, and the others also leveraging technological advancements in their respective operations. Additionally, they are relatively large players in their fields, with significant market capitalization and the potential for global reach or influence in their industries. Their efficient business models contribute to industry progress.

After gathering the data, it was carefully inspected. Fortunately, there were no missing values. However, a routine check for outliers was conducted, and a small number of minor outliers were adjusted to ensure data integrity. After that, the data was separated into a test set including the thirty days remained and a training set involving the first seventy days.

Table 1 shows the descriptive data of the closing prices therefore offering a more complete picture of the features of the stocks. The mean closing prices vary, indicating different value levels. Hundsun Electronics and Eastmoney with negative kurtosis possess relatively stable price trends and lower probabilities of extreme prices, suitable for risk-averse investors but with limited potential for high returns in volatile markets. Heertai has larger skewness and kurtosis, implying a higher likelihood of upward extreme values and significant price fluctuations. The standard deviation differences also reflect the varying degrees of price volatility among the stocks.

Table 1. Descriptive statistics of five stocks

	002402	600570	000066	300059	601727
Mean	11.0342	20.6620	10.5157	15.2994	5.0232
Std	2.7910	5.6608	4.1156	6.6807	2.3731
Skew	2.6589	0.9297	1.7489	0.8475	1.6375
Kurt	6.9399	-0.6410	1.6814	-0.9800	1.1655

Figure 1 illustrates the cumulative returns of five stocks over a 100-trading-day period. These returns are calculated as the cumulative percentage change in closing prices relative to the initial trading day.

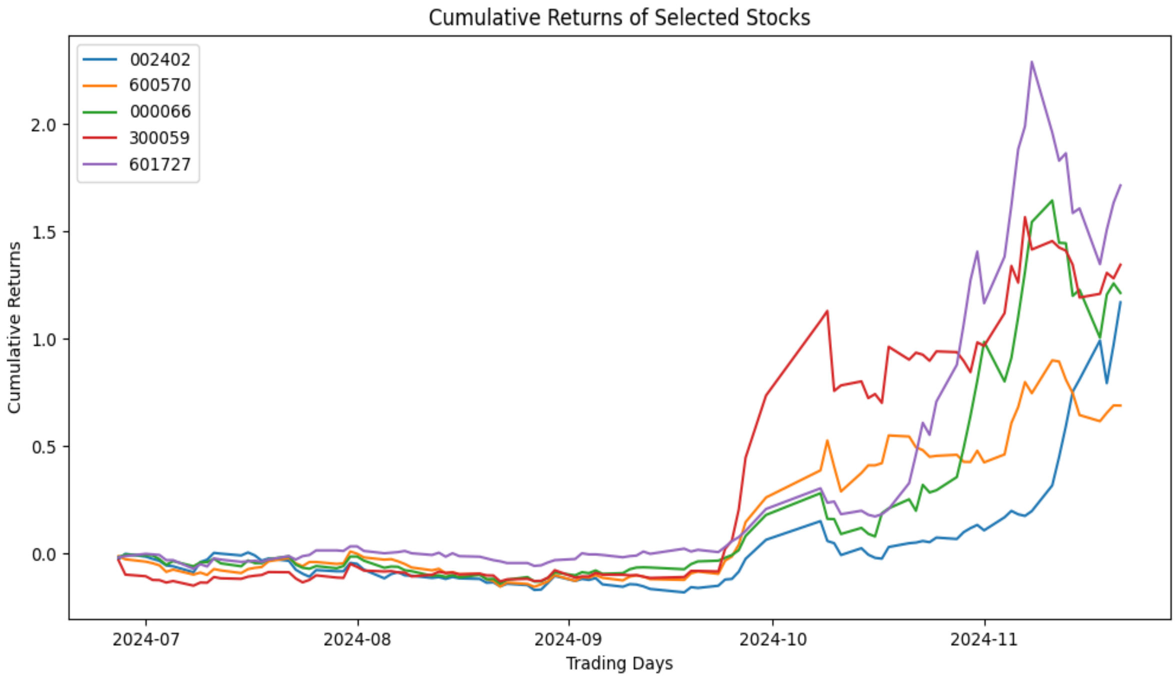

Figure 1. Cumulative returns of selected stocks

3. Method

3.1. Long Short-Term Memory (LSTM) Model

Processing sequential data like stock price time series is a strength of the specialized recurrent neural network (RNN) LSTM. Correctly identifying long-term dependencies in the data depends on overcoming the vanishing gradient problem affecting conventional RNNs. Here in an outline is its stock price prediction system.

3.1.1. Data Input and Initialization

As input, the LSTM model handles a sequence of past stock price data. Every time step starts with the internal state and hidden state of the LSTM cell initializing themselves. Usually designated as h_{t-1}, the hidden state and the cell state, c_{t-1}, are set to starting values, often zeros or tiny random numbers.

3.1.2. Forget Gate

The forget gate of an LSTM cell decides which information from the previous cell state, c_{t-1}, should be kept or thrown away first in its functioning. Using x_t as the current input and h_{t-1} as the past hidden state, this gate sends both through a sigmoid activation function. The forget gate, f_t, generates a vector of values between 0 and 1 whereby a value near 0 denotes that the associated information from the previous cell state should be forgotten and a value near 1 suggests that it should be kept. The forget gate is computed numerically as:

$$f_t = \sigma\left(W_f \cdot [h_{t-1}, x_t] + b_f\right) \tag{1}$$

W_f is the weight matrix; b_f is the bias vector connected with the forget gate; σ is the sigmoid function.

3.1.3. Input Gate and Cell State Update

Two elements define the input gate: the gate itself and the candidate cell state.

i_t, the input gate, chooses which fresh information from the previous hidden state h_{t-1} and the present input x_t should be included into the cell state. It is calculated also with a sigmoid activation function:

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \quad (2)$$

Where W_i and b_i are the weight matrix and input gate bias vector respectively.

Using a hyperbolic tangent (tanh) activation function, one computes the candidate cell state, $\widetilde{c_t}$. It produces a vector of possible new cell state values depending on the concealed state and current input:

$$\widetilde{c_t} = tanh(W_c \cdot [h_{t-1}, x_t] + b_c) \quad (3)$$

Where the weight matrix W_c and bias vector b_c correspond for the candidate cell state.

Combining the forget gate output and the input gate and candidate cell state then updates the new cell state c_t. The candidate cell state is multiplied by the input gate output then joined together after the prior cell state is multiplied by the forget gate output:

$$c_t = f_t \odot c_{t-1} + i_t \odot \widetilde{c_t} \quad (4)$$

Where $\odot$ represents element-wise multiplication.

3.1.4. Output Gate and Hidden State Update

For the current time step, the output gate o_t decides which portions of the updated cell state $\mathrm{c_t}$ should be output as the hidden state h_t. It is computed applying a sigmoid activation function:

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \quad (5)$$

Where the weight matrix W_o and bias vector b_o apply to the output gate.

Multiplying the output gate output with the hyperbolic tangent of the updated cell state then determines the hidden state h_t:

$$h_t = o_t \odot tanh(c_t) \quad (6)$$

3.1.5. Sequence Processing and Prediction

For every time step in the input sequence, the LSTM cell carries out these tasks. Predicting uses the last hidden state or a combination of hidden states once the whole sequence is processed. In stock price forecasting, the output of the model could be a straightforward prediction like the closing price of the next day or more intricate outputs including trends or ranges of likely pricing. Minimizing a loss function, such as mean squared error, which gauges the variance between expected and actual stock prices in the training dataset—is what training the LSTM is all about. By means of optimization techniques such stochastic gradient descent or variations, this procedure modifies the weights and biases of the model.

LSTM gates' selective memory properties help the model to efficiently capture long-term dependencies and trends in stock price data. This makes it an excellent instrument for stock price analysis and prediction. Figure 2 shows an LSTM's architecture.

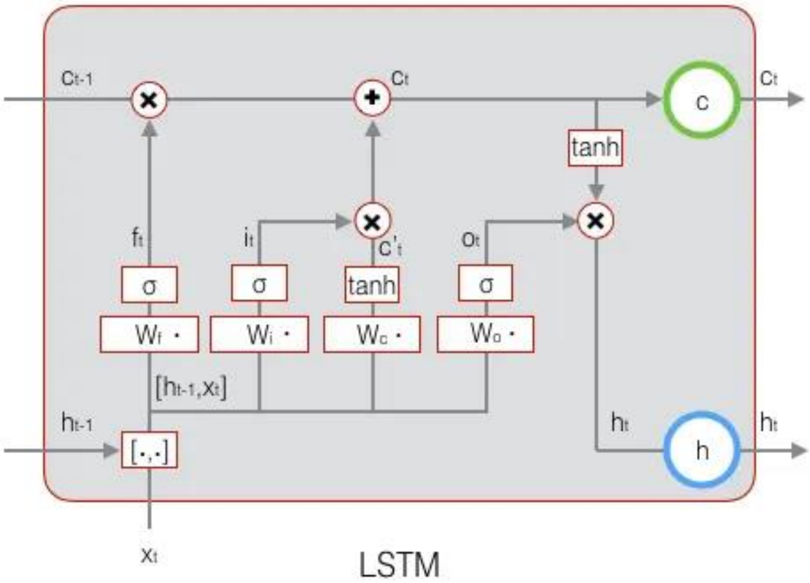

Figure 2. Structure of LSTM

3.2. Mean-Variance Optimization Model

Introduced by Harry Markowitz, the Mean- Variance Optimization (MVO) model is pillar of contemporary portfolio theory. It aims to balance the trade-off between projected returns and portfolio risk, expressed as variance, therefore producing an optimal portfolio.

3.2.1. Data Input and Parameter Estimation

- Historical Return Data: For each stock i in the portfolio, the return at time t, $r_{i,t}$, is calculated as

$$r_{i,t} = \frac{P_{i,t} - P_{i,t-1}}{P_{i,t-1}} \tag{7}$$

Where $P_{i,t}$ is the stock price of stock i at time t.

- Risk-Free Rate: Denoted as R_f, it is an important benchmark return.
- Covariance Matrix Estimation: The covariance between stock i and stock j is $\text{cov}(r_i, r_j)$. The covariance matrix is estimated from historical return data.

3.2.2. Portfolio Return and Variance Calculation

- Portfolio Return: If a portfolio has n stocks with weights w_i and the expected return of stock is $E(r_i)$, then the portfolio return $E(R_P)$ is given by

$$E(R_P) = \sum_i w_i E(r_i) \tag{8}$$

- Portfolio Variance: The variance of the portfolio σ_P^2 is calculated as

$$\sigma_P^2 = \sum_{ij} w_i w_j cov(r_i, r_j) \tag{9}$$

3.2.3. Optimization Objective and Constraints

- Optimization Objective: The Mean-variance Optimization model has two basic optimization goals. One is getting the Sharpe ratio as maximum,

$$S = \frac{E(R_P) - R_f}{\sigma_p} \tag{10}$$

Which balances risk and return, aiming for the highest return per unit risk and is favored by investors seeking such a trade-off. The other is minimizing the portfolio variance σ_P^2, focusing on reducing risk and appealing to highly risk-averse investors.By means of a thorough comparison, this study seeks to offer insightful analysis of the performance and features of portfolios built under every objective, therefore enabling more informed investment decision-making.

Constraints: The sum of weights must equal 1, i.e., $\sum_i w_i = 1$. Additionally, there may be constraints like $0 \leq w_i \leq 1$ for each i to limit short selling and over-concentration.

This model provides a quantitative framework for investors to balance return and risk in portfolio management.

4. Results

4.1. Sliding Window Weight Calculation and Portfolio Return Calculation

4.1.1. Sliding Window Operation

Leveraging the LSTM model trained with the data from days 1 - 70, we predicted the prices of each asset on day 71. By comparing the actual prices on day 70 with the predicted prices on day 71, the returns of each asset on day 71 were computed. The covariance matrix calculated from the data of days 1 - 70 was adopted as the covariance matrix for day 71. Within the framework of the minimum variance model (or the maximum Sharpe ratio model), the investment portfolio weights on day 71 were determined.

Subsequently, the data window was advanced. Each time the window was shifted, the data within the new window was employed to predict the asset prices for the ensuing day. New returns and covariance matrices were calculated, and consequently, the investment portfolio weights for the next day were obtained. Through continuous window sliding, the investment portfolio weights for each day from day 71 to day 100 were derived.

4.1.2. Portfolio Return Calculation

The true return of the portfolio under this investment strategy was computed for every day within the period from day 71 to day 100 by combining the investment portfolio weights of the day with the true returns of each asset. Thus, a sequence of the daily returns of the investment portfolio from day 71 to day 100 was generated.

4.2. Cumulative Return Calculation and Line Chart Drawing

The total returns of these portfolios from day 71 to day 100 were computed respectively depending on the sequences of daily returns of the portfolios built by the minimum variance model and the maximum Sharpe ratio model. The cumulative return of the CSI 300 Index from day 71 to day 100 was also calculated using its daily true returns. All three cumulative return series (minimum variance portfolio, maximum Sharpe ratio portfolio, and CSI 300 Index) were then plotted in Figure 3. This chart clearly shows the trends and comparisons of the cumulative returns over the 30 - day testing period.

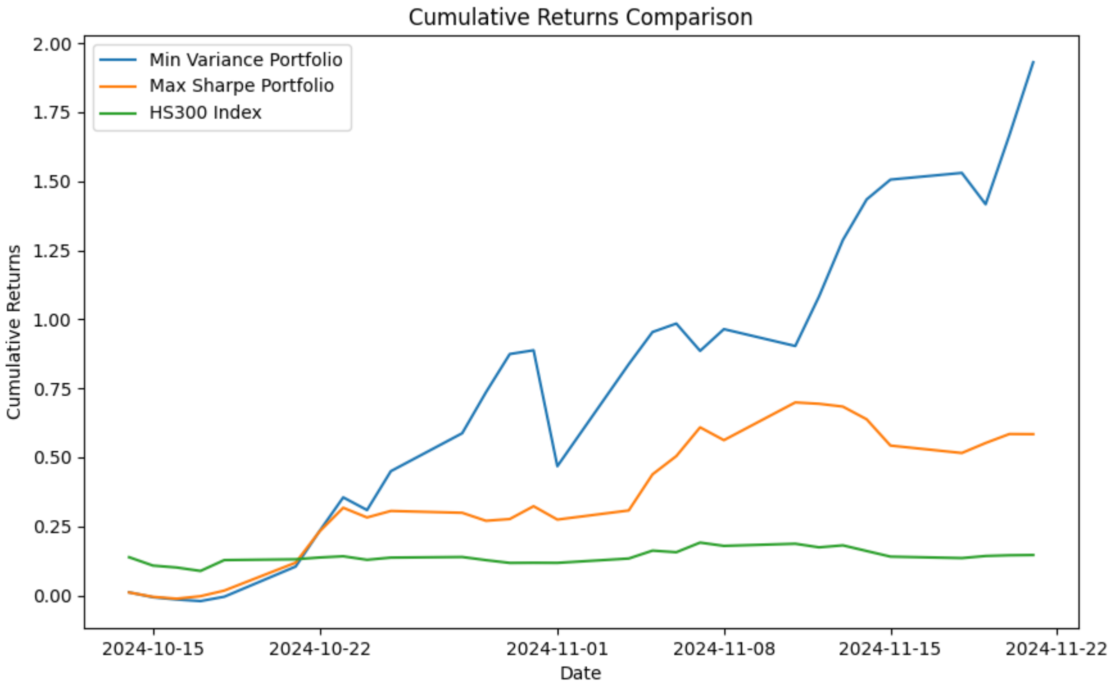

Figure 3. Cumulative Returns Comparison

4.3. Result Display Table and Analysis

4.3.1. Result Table

Table 2 is a brief table displaying the CSI 300 Index's cumulative returns on day 100 together with the investment portfolio's results.

Table 2. Comparison of Cumulative Returns of Models and CSI 300 Index (Day 100)

Model/Index	Cumulative Return (%)
Minimum Variance Model	193.1164
Maximum Sharpe Ratio Model	61.2246
CSI 300 Index	14.6193

4.3.2. Analysis

From the line chart, the trends of the cumulative returns of the three can be clearly observed.

Over the whole 30-day testing period, the minimum variance portfolio's cumulative return had a constant increasing trend. It began to climb gradually from the outset and kept on expanding free from major obstacles. By day 100 it had a total return of 193.1164%. The cumulative return of the greatest Sharpe ratio portfolio showed a pattern of first gradual development then more marked rising trend. There were some fluctuations in the middle period, but overall, it also showed an upward - moving trajectory, reaching a cumulative return of 61.2246% by day 100. The CSI 300 Index exhibited a relatively flat growth trend over the 100-day period, with minor fluctuations and a cumulative return of 14.6193%.

On cumulative returns over the testing period, however, both the minimal variance model and the highest Sharpe ratio model much exceeded the CSI 300 Index. Of the two, the minimum variance model shown better performance. This result shows that, given both models attained superior cumulative returns compared to the market index during the testing period, the suggested portfolio creation technique clearly optimized investment selections.

5. Summary

This work presents a mean-variance optimization (MVO) model combined with the Long Short-Term Memory (LSTM) model into an investment portfolio optimization technique. For examination were five technology-driven stocks with notable market impact. Daily closing price data for these stocks, spanning from June 26 to November 21, 2024, was collected and preprocessed. The LSTM model was trained using the first 70 days of data to predict future stock prices. Subsequently, within the Mean-Variance Optimization model framework, portfolio weights were calculated using the predicted prices and historical data via a sliding window method. The results show that throughout the 30-day test period, both the minimal variance and maximum Sharpe ratio models exceeded the CSI 300 Index in cumulative return; the lowest variance model fared better than the maximum Sharpe ratio model. This confirms how well the suggested approach will maximize investment portfolios.

Future studies could concentrate on extending the stock selection to increase portfolio diversity and test the resilience of the strategy; integrating other advanced models or algorithms, like deep reinforcement learning, for better performance; including more factors such macroeconomic and fundamental data for a more accurate decision basis; and doing longer-term simulations and backtests to enhance practical insights. This would contribute to the evolution of investment portfolio optimization in finance.

Acknowledgement

I truly appreciate the great direction and helpful criticism that Professor Zhang, my adviser, provided, which helped this project to be successfully completed.

I would wish to congratulate Donghua University's Glorious Sun School of Business and Management for fostering scholarly growth in an environment. Additionally, much valued are my coworkers and peers for their perceptive comments and unwavering support over this effort.

At last, I owe a great lot to my friends and family for their understanding and support. Their unflinching faith in me has been a continual inspiration.

References

[1] ASLAN, Y., & ÖZKAN, Ö. (Eds.). (2022). Innovative Approaches to Accounting, Finance and Auditing-4. Efe Akademi Yayınları.

[2] Markowitz, H. (1952). Portfolio Selection. The Journal of Finance, 7 (1), 77 - 91.

[3] Kim, K. J., & Han, I. (2000). Genetic algorithms approach to feature discretization in artificial neural networks for the prediction of stock price index. Expert Systems with Applications, 19 (2), 125 - 132.

[4] Chen, P. H., & Lin, C. J. (2006). Combining SVMs with various feature selection strategies. In Proceedings of the Fourth International Conference on Machine Learning and Cybernetics (pp. 3199-3204). IEEE.

[5] Ballings, M., Van den Poel, D., Hespeels, N., & Gryp, R. (2015). Evaluating multiple classifiers for stock price direction prediction. Expert Systems with Applications, 42 (20), 7046 - 7056.

[6] Deng, X., Li, Y., & Wen, J. R. (2017). A new method for stock price prediction based on random forest. In Proceedings of the 36th Chinese Control Conference (CCC) (pp. 10098-10102). IEEE.

[7] Bhardwaj, G., & Tiwari, A. K. (2019). Stock price prediction using machine learning techniques: A survey. Journal of King Saud University-Computer and Information Sciences, 31 (4), 480 - 495.

[8] [8] Roondiwala, M., Patel, H., & Varma, S. (2015). Predicting Stock Prices Using LSTM. 6 (4).

[9] Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, 270 (2), 654 – 669.

[10] Khaidem L., Saha S. & Roy Dey S. (2016). Predicting the direction of stock market prices using random forest. arXiv preprint arXiv: 1605.00003.

[11] Basak, S., Kar, S., Saha, S., Khaidem, L., & Dey, S. R. (2019). Predicting the direction of stock market prices using tree-based classifiers. The North American Journal of Economics and Finance, 47, 552 – 567.