Evaluating the Effectiveness of Machine Learning Algorithms in Stock Price Prediction Across Different Time Frames

Isabella J E Lim[a], **Amy L Y Low**[a], **Kenneth Y T Lim**[b]

[a] *independent scholar, Singapore*

[b] *National Institute of Education, Singapore* kenneth.lim@nie.edu.sg

1. Introduction

This project seeks to test various machine learning algorithms for stock price prediction in Singapore's banking sector over different time frames to optimize predictive accuracy, model performance, and their applicability to short-term and long-term investment strategies. Predicting stock price trends is not only vital for maximizing investment returns but also crucial for risk management. In addition, this project will offer an evaluation of different algorithms' performance across varying time periods, providing investors with better insights into making informed decisions when conducting trades in the stock market.

2. Method

The algorithms used in the study include Random Forest (RF), Support Vector Regression (SVR), K-Nearest Neighbor (KNN), Long Short-Term Memory (LSTM), and Artificial Neural Network (ANN). Each algorithm will be trained on historical stock data from three prominent Singaporean banks - DBS, OCBC, and UOB. Each machine learning algorithm functions differently, utilising unique computational methods to identify patterns and trends in the data. To evaluate the performance of the above algorithms for predicting stock closing prices, we utilized Python's scikit-learn library to implement, train, and test our models. The workflow involved data preprocessing, model training, and performance evaluation using selected metrics, ensuring a robust and systematic approach. We used the yahoofinance library to populate our dataset with historical stock prices of DBS UOB and OCBC, which were cleaned and prepared for analysis. Missing values were handled by manually inputting the values from other historical stock prices databases where necessary. Each model was initialized with default hyperparameters. Once trained, the predictions generated by each model were compared against the actual stock prices using the following evaluation metrics:

Based on the predicted closing price and the absolute closing price, we then calculated these metrics, Percentage Error, R-squared (R^2), Mean Absolute Error (MAE), and Mean Squared Error (MSE). These metrics were chosen to provide insights into the accuracy, reliability, and explanatory power of the models.

Percentage error measures the absolute deviation of the predicted price from the actual price as a percentage. It helps translate absolute errors into relative terms, making it easier to interpret the accuracy across different scales of stock prices. A lower percentage error signifies better prediction accuracy.

R^2 quantifies the proportion of variance in the actual stock price explained by the model. R2 values closer to 1 suggest a stronger explanatory power, indicating the model's ability to account for variations in the stock prices.

MAE measures the average absolute differences between predicted and actual values. Smaller MAE values indicate a higher level of prediction accuracy by minimizing overall errors.

MSE evaluates the average squared differences between predicted and actual values. Unlike MAE, MSE squares the residuals, giving disproportionately higher weights to larger errors. This makes it more sensitive to outliers and places a higher penalty on significant deviations from the actual values. Lower MSE values reflect better prediction accuracy, with an emphasis on penalizing larger errors due to the squaring of residuals.

R^2 complements MAE and MSE by evaluating the proportion of variance in the actual target values that is explained by the model's predictions, offering insights into the model's overall explanatory power. While MAE and MSE primarily focus on quantifying the magnitude of errors, R^2 provides a broader perspective by assessing how effectively the model captures the variability within the dataset, thereby indicating its ability to generalize to unseen data. The percentage error metric further enhances analysis by contextualizing the significance of the errors identified by MAE and MSE, especially in datasets with diverse price ranges. Together, these metrics provide a holistic set of metrics for evaluation, enabling a comprehensive assessment of model performance.cited.

3. Results and Discussion

RF consistently shows high R^2 values and low percentage errors across all three stocks in short-term predictions, emphasizing its strength in capturing immediate trends. As the timeframe increases, its percentage error generally increases while R^2 value remains relatively constant. Despite maintaining relatively constant R^2 values in mid-term and long-term predictions, its percentage errors increase significantly. This suggests that RF may overfit to training data, making it less adaptive to the volatility and unpredictability of long-term financial data. Hence, RF is highly effective for short-term forecasting but less reliable for long-term predictions due to its deterministic nature and sensitivity to overfitting.

SVR shows a consistent trend of having the lowest percentage errors amongst all other algorithms, with a general increasing trend with increasing time frame. It also maintains high R^2 values across all stocks, but tends to have a slightly lower R^2 value for short-term predictions as compared to mid-term and long-term predictions. This may suggest that it prioritizes accuracy over generalizability in this timeframe, restricting its adaptability in highly

Evaluating the Effectiveness of Machine Learning Algorithms in Stock Price Prediction Across Different Time Frames

Isabella J E Lim[a], **Amy L Y Low**[a], **Kenneth Y T Lim**[b]

[a] *independent scholar, Singapore*

[b] *National Institute of Education, Singapore* kenneth.lim@nie.edu.sg

volatile markets.

KNN generally maintains low percentage errors for short-term predictions, but has a significant increase in error for long-term predictions. Additionally, its R^2 value for short-term predictions is significantly lower than that for mid-term and long-term predictions, indicating that it struggles to capture the overall variability in stock price trends in the short term. However, for longer time frames, the R^2 values improve significantly, indicating its ability to generalize better with larger datasets and longer trends, but comes at the cost of increasing percentage errors, particularly for long-term predictions. This suggests that its reliance on proximity-based relationships becomes less effective in highly dynamic and volatile environments.

LSTM was unable to capture trends for shorter time frames, performing poorly when given data for less than one month. For longer time frames, it consistently showed high percentage errors and low R^2 values. This suggests its limited capability to handle the volatility and non-linear relationships inherent in stock price data over extended periods.

ANN had a general increasing trend of percentage errors with increasing time frames, with a few notable anomalies. Its R^2 value was the lowest amongst all the algorithms for short-term predictions, but significantly improved for mid-term and long-term predictions. This shows that it struggles to capture meaningful variability in small datasets, but is able to generalize trends over time, albeit at the cost of accuracy. ANN is adaptable for mid-term and long-term predictions but requires extensive optimization of hyperparameters and datasets to enhance its accuracy and minimize errors.

Random Forest and SVR consistently outperform other models, demonstrating low error rates and stable results across short- to long-term predictions. For example, Random Forest achieves an MAE of 0.1464 and MSE of 0.0216 for DBS stock over 2 days and maintains strong performance even over 3 years. SVR shows similarly effective performance, only being slightly worse than Random Forest in longer time horizons.

In contrast, ANN and LSTM are less reliable, with error rates increasing significantly in long-term predictions. LSTM, while theoretically strong for time-series data, struggles to outperform other models consistently. It performs poorly in both medium and long-term predictions, with high error rates. These results suggest that LSTM's requirement for large datasets and careful tuning limits its practical applicability in this context. KNN performs inconsistently, with its error rates rising noticeably in the long term. Overall, Random Forest and SVR are the most effective, while ANN and LSTM require improvements in tuning and data handling for better long-term accuracy.

4. Conclusion

The study's findings are highly dependent on the dataset used, focusing on Singapore's banking sector (DBS, OCBC, UOB stocks), which may not generalize to other industries or markets. Additionally, algorithms like ANN and LSTM are highly sensitive to hyperparameter tuning, leading to variable performance across different scenarios. Moreover, the analysis is purely based on statistical data, primarily using closing prices, which might not fully capture market dynamics. Supplementing additional features like trading volume, economic indicators, and volatility analysis could improve model performance and robustness. It also does not account for external market factors, such as geopolitical events or sudden economic shifts, which can significantly impact stock prices. The addition of sentiment analysis of external sources to the models' dataset could greatly improve their predictive capabilities.

In conclusion, this project highlights the performance of various machine learning algorithms in predicting stock price trends within Singapore's banking sector over different time frames. Among the models analysed, Random Forest (RF) and Support Vector Regression (SVR) consistently performed remarkably across both short and long-term horizons, with RF excelling in short-term predictions and SVR maintaining high accuracy and adaptability over longer timeframes. While K-Nearest Neighbors (KNN) shows potential in mid to long-term predictions, its reliability diminishes in volatile market conditions. On the other hand, Artificial Neural Networks (ANN) and Long Short-Term Memory (LSTM) models, although theoretically well-suited for time-series data, struggle with higher error rates and lower R^2 values, particularly for long-term predictions. These results underline the importance of balancing algorithm choice with the specific demands of prediction timeframes and data variability.

In light of the growing shift toward algorithmic trading, a transformative trend in financial markets where decisions are increasingly made by AI-driven systems, it has become even more paramount to gain a better understanding of how different algorithms work and their possible applications to the stock market. Algorithmic trading relies on automated, rule-based strategies to execute trades at high speeds, often leveraging machine learning models to forecast trends and identify profitable opportunities.

References

[1] Najem, R., M. F. Amr, A. Bahnasse, and M. Talea, "Advancements in Artificial Intelligence and Machine Learning for Stock Market Prediction: A Comprehensive Analysis of Techniques and Case

Evaluating the Effectiveness of Machine Learning Algorithms in Stock Price Prediction Across Different Time Frames

Isabella J E Lim[a], **Amy L Y Low**[a], **Kenneth Y T Lim**[b]

[a] *independent scholar, Singapore*

[b] *National Institute of Education, Singapore* kenneth.lim@nie.edu.sg

Studies," Procedia Computer Science, vol. 231, pp. 198–204, Jan. 2024, doi: https://doi.org/10.1016/j.procs.2023.12.193.
[2] "A Comparative Study of Supervised Machine Learning Algorithms for Stock Market Trend Prediction - IEEE Conference Publication," Ieee.org, 2018. https://ieeexplore.ieee.org/document/8473214
[3] Bao, W., J. Yue, and Y. Rao, "A deep learning framework for financial time series using stacked autoencoders and long-short term memory," PLOS ONE, vol. 12, no. 7, p. e0180944, Jul. 2017, doi: https://doi.org/10.1371/journal.pone.0180944.
[4] Chang, C.-C., and C.-J. Lin, "LIBSVM: A Library for Support Vector Machines," 2001. Available: https://www.csie.ntu.edu.tw/~cjlin/papers/libsvm.pdf
[5] Platt, J. C., "Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods," vol. 10, no. 3, Jun. 2000, Available: https://www.researchgate.net/publication/2594015_Probabilistic_Outputs_for_Support_Vector_Machines_and_Comparisons_to_Regularized_Likelihood_Methods
[6] Van Houdt, G., C. Mosquera, and G. Nápoles, "A review on the long short-term memory model," Artificial Intelligence Review, vol. 53, no. 8, May 2020, doi: https://doi.org/10.1007/s10462-020-09838-1.