Analysis of Intraday Financial Market Using ML and Neural Networks for GBP/USD Currency Pair Price Forecasting

Melis Zhalalov and Vitaliy Milke[a]

Computing and Information Science, Anglia Ruskin University, East Road, Cambridge, U.K.

Keywords: Artificial Neural Networks, Machine Learning, K-Nearest Neighbors, Logistic Regression, Decision Trees, Random Forest, Support Vector Machines (SVM), MLP, LSTM, Intraday Trading.

Abstract: This study employs a range of machine learning and artificial neural network techniques for financial market price prediction. The approach involves data preprocessing, feature engineering, and model evaluation using daily and 5-minute interval records. Leveraging methods like K-Nearest Neighbors, Logistic Regression, Decision Trees, Random Forest, Support Vector Machines, Multi-Layer Perceptron and Long Short-Term Memory networks, the models exhibit distinct strengths and limitations. Notably, the LSTM model achieved an accuracy of 63%, while Random Forest demonstrated 60% accuracy, indicating promising results for intraday trading. It is essential to acknowledge that due to the exclusion of night hours, the approach is tailored specifically for intraday trading. This study offers a valuable approach to exchange rate prediction, providing an additional practical resource for practitioners and researchers in the field of financial market forecasting.

1 INTRODUCTION

Forecasting financial market trends, particularly in foreign exchange and stock markets, presents a longstanding challenge due to their inherently volatile and unpredictable nature. Reasonable predictions are paramount in investment decision-making, risk management, and portfolio optimisation. The advent of advanced technology and the availability of extensive historical data have paved the way for data-centric approaches, including machine learning (ML) and artificial neural networks (ANN), to address this complicated task.

Conventional forecasting methods often fail to capture the nuanced patterns steering market movements. These movements predominantly focus on a single method, as outlined in a study conducted by Altman (1968), who used univariate approaches to predict corporate bankruptcy. This study addresses limitations by employing machine learning algorithms to unveil concealed trends. Notably, the focus is exclusively on short-term intraday market movements, ranging from minutes to several hours. This domain encompasses rapid fluctuations and complex interactions influencing currency prices within a single trading day. It is imperative to distinguish this focus from high-frequency trading (HFT), which operates at rapid speeds measured in milliseconds to seconds and requires expensive infrastructure investments, as illustrated in the study conducted by MacKenzie (2019). By narrowing the scope to short-term intraday movements, the aim is to recognise the underlying forces driving these market shifts and construct predictive models effectively.

The primary objective of this research is to develop models to forecast the direction of the following price movement of the GBP/USD currency pair and the level of the future closing price of the selected time frame relative to the opening price (higher or lower). Market data spanning from 2022 to mid-2023 is utilised using a data-centric approach. The emphasis lies in employing various ML classification techniques, encompassing Support Vector Machines, Decision Trees, Random Forests, K-nearest Neighbors, Logistic Regression, and Long-Short-Term Memory (LSTM) to discern relationships within the data and hidden patterns for making informed financial decisions.

Furthermore, this study seeks to address another research gap by identifying early indicators of significant price movements within flat swings—periods characterised by relatively stable prices.

[a] https://orcid.org/0000-0001-7283-2867

Zhalalov, M. and Milke, V.
Analysis of Intraday Financial Market Using ML and Neural Networks for GBP/USD Currency Pair Price Forecasting.
DOI: 10.5220/0012386700003636

In *Proceedings of the 16th International Conference on Agents and Artificial Intelligence (ICAART 2024) - Volume 3*, pages 734-741
ISBN: 978-989-758-680-4; ISSN: 2184-433X

Recognising these cues can significantly enhance short-term price prediction and the trading system's potential profit. By excluding external factors, a targeted method for prediction is provided, with a pronounced emphasis on historical data to isolate and analyse patterns associated with flat periods.

The distinctiveness of this approach lies in its emphasis on optimising accuracy using exclusively digital historical datasets, but simultaneously for bid and ask prices in the order book. This is pivotal, as accessibility to other big data, such as news or text comments from market participants, may vary, and processing it can be resource-intensive. By relying on numerical historical data, the need for extensive preprocessing is circumvented, thereby simplifying computation and processing time in various formats, and does not require expert weighing of the importance of such data. Additionally, this approach directly addresses the challenge of extracting insights from currency exchange market data in real-time decision-making. In contrast to research integrating factors such as news sentiment analysis and extensive Twitter data (Maqsood et al., 2020), the concentration is solely on historical data, streamlining the analysis process. It is assumed that all news sentiments are already included in the market prices of classical financial instruments such as stocks, currencies, bonds, ETFs (exchange-traded funds), but the market is still inefficient.. By zeroing in on the most probable flat movements independent of external influences, this research contributes to a more sustainable and efficient prediction model.

The successful application of these models carries wide-ranging advantages. Individual investors may potentially gain valuable insights for informed decision-making and optimized investment strategies. Financial institutions bolster their credibility by offering expertise in data-driven investment decisions. Effective risk management is facilitated through improved predictions, aiding in identifying and mitigating risks.

The research is based on data from the Repository of (Dukascopy Bank Sa, 2023), a recognized source for accurate currency rate datasets. The research follows a two-step approach involving data engineering techniques for preprocessing and organizing the data, followed by applying Machine learning (ML) methods and Neural Network (NN) architectures to develop predictive models.

The rest of this paper is organized as follows: Section 2 investigates the literature and state-of-the-art studies on the topic; Section 3 describes the proposed method; Section 4 outlines the results; finally, the conclusions are drawn in Section 5.

2 RELATED WORK

The foundational work of Schierholt and Dagli (1996) pioneers the application of AI techniques in stock market prediction, explicitly focusing on data preprocessing for the Standard & Poor's 500 Index. Their utilisation of neural networks aligns closely with this study's goal of predicting currency exchange rate movements. While their study concentrates on stock market movements using neural network structures for prediction, it resonates with the intention to forecast currency exchange rates.

Building on this foundation, Zhanggui, Yau, and Fu (1999) introduce an innovative approach centred on data preprocessing for pattern analysis and relaxation classification for stock price prediction. This method effectively captures changes in stock prices over time, demonstrating its potential for refining investment decisions. This research aligns with the current research focusing on employing advanced techniques for preprocessing financial data as inputs for neural networks, predicting currency exchange rates.

The research conducted by Islam and Hossain (2021) proposes a model for predicting future closing prices of FOREX currencies. The model combines Gated Recurrent Unit (GRU) and Long Short-Term Memory (LSTM) neural networks. The study focuses on four major currency pairs: EUR/USD, GBP/USD, USD/CAD, and USD/CHF. The experiment was conducted for 10-minute and 30-minute timeframes, and the performance of the model was evaluated using regression metrics. The proposed hybrid GRU-LSTM model demonstrated high accuracy in predicting currency prices for both short and medium-timeframes. It outperformed standalone GRU and LSTM models and a simple moving average (SMA) model.

The research paper by Abedin (2021) proposes an ensemble deep learning approach that combines Bagging Ridge (BR) regression with Bi-directional Long Short-Term Memory (Bi-LSTM) neural networks. This integrated model, referred to as Bi-LSTM BR, is used to predict the exchange rates of 21 currencies against the USD, including GBP, during both pre-COVID-19 and COVID-19 periods. The study compares the proposed Bi-LSTM BR approach with traditional machine learning algorithms (such as Regression Tree, SVM and Random Forest regression) as well as deep learning-based algorithms like LSTM and Bi-LSTM regarding prediction error. However, the performance of the model varies significantly across different currencies and during non-COVID-19 and COVID-19 periods, highlighting the importance of prediction models in highly volatile foreign currency markets.

Incorporating feature selection and a composite classifier, Vignesh's (2018) comparative analysis between SVM and LSTM for stock price prediction highlights LSTM's effectiveness in temporal processing. This aligns with this study's approach of utilising LSTM alongside other AI methodologies for predicting currency exchange rates.

In contrast to the papers listed above, one distinctive feature of our approach is excluding night hours from trading activities. This decision is rooted in the understanding that market activity is reduced during these hours, and the spread between ask and bid prices is often too high, rendering trading non-advantageous. This strategic exclusion optimises trading performance by focusing on periods of heightened market activity and reduced bid-ask spread differentials.

Additionally, this research addresses the categorical nature of the task, which aligns seamlessly with intraday trading. By categorising potential directions of short-term price movements (of which there are only two: up or down), the model can make more precise predictions, increasing the likelihood of profitable trades due to the greater statistical probability of achieving take profit on short-term flat movements. Furthermore, this approach is designed for real-time intraday trading without HFT technologies and news analysis, eliminating the need for big data or additional non-numerical data sources. This streamlined process allows the model to remain adaptable and efficient for practical intraday trading scenarios without using expensive computer hardware.

3 PROPOSED APPROACHES

In contrast to other state-of-the-art papers that emphasize regression-based approaches for forecasting the Close market price, the models selected in this paper for predicting short-term price movements encompass a variety of techniques renowned for their effectiveness in classification tasks and time series analysis. This approach proves more advantageous for intraday trading, where understanding the price movement direction holds greater significance than precise numerical price values. The models chosen are K-Nearest Neighbors (KNN) (Aha, Kibler, and Albert, 1991), Support Vector Machines (SVM) (Keerthi, 2001), Logistic Regression (Cessie and Houwelingen, 1992), Decision Trees, Random Forest (Breiman, 1996), Multilayer Perceptron (MLP) (Pedregosa et al., 2011) and Long Short-Term Memory (LSTM) networks (Goodfellow, Bengio, and Courville, 2016).

The development process entails several essential stages:

- Feature Engineering: This involves enhancing the data to extract meaningful insights, such as incorporating lagged features and technical indicators (Long, Lu, and Cui, 2019).
- Training the models with Hyperparameter Tuning: Fine-tuning parameters used in this research are essential for optimal model performance, particularly in the case of deep neural networks (Hoque and Aljamaan, 2021).
- The evaluation and validation techniques used in this research comply with the accepted rigorous standards described (Nauta et al., 2023) and are employed to ensure models generalise well to new data, simulating real-world conditions.
- Interpretability and Visualisation: Given the complexity of some models, like LSTMs, it is crucial to employ interpretability techniques to bridge the gap between algorithmic insights and human understanding (Samek et al., 2019).
- Addressing Market Volatility: The approach accounts for market volatility by incorporating measures like the Volatility Index (VIX), allowing for more accurate predictions in rapidly changing market conditions (Engle, Ghysels, and Sohn, 2013).

The research leverages Python 3.7.1 as the primary programming language, executed in the Jupyter Notebook environment. It relies on key libraries like Scikit-Learn (Pedregosa et al., 2011), TensorFlow (Abadi et al., 2016) with Keras (Chollet et al., 2015), Pandas for efficient data manipulation, and Matplotlib and Seaborn for visualisation.

3.1 Data Used

Detailed, clean, trustworthy data is critical to applying machine learning in finance. This research uses data obtained from reliable sources (Dukascopy Bank Sa, 2023). This dataset encompassed massive historical market price and volume information, providing a comprehensive view of past market dynamics. This dataset included an array of vital features, including 5-minute and daily open, high, low and close prices for ask and bid orders in the order book, trading volume executed separately at bid and ask prices and additional relevant financial metrics. These features were carefully curated to serve as the foundation for these predictive models.

It is worth noting that the dataset incorporates data from January 2022 to June 2023 (inclusive), comprising approximately 220,000 rows, ensuring a robust foundation for the machine learning analysis.

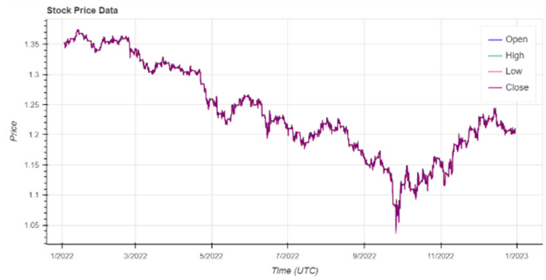

Figure 1: Time Series Plot of GDP/USD for the Whole Year (2022).

Figure 1 offers a macroscopic view of GBP/USD price movements throughout the entire year of 2022. This visualisation provides a high-level perspective, allowing for discerning long-term trends, seasonal patterns, and significant events that have impacted the financial markets during the year.

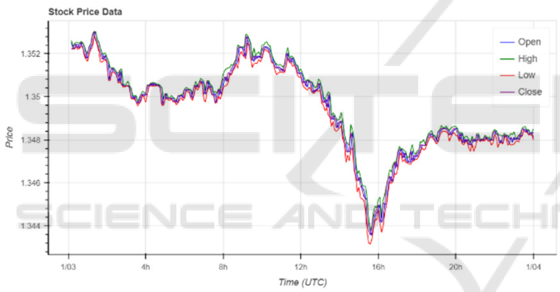

Figure 2: Time Series Plot of GDP/USD for One Day (03/01/2022).

In contrast, Figure 2 zooms in on a single day, focusing specifically on March 1, 2022. This fine-grained time series plot enables detailed exploration of intraday flat fluctuations, pinpointing volatility patterns and exploring the nuanced dynamics of market prices on a micro-timescale.

3.2 Data Preprocessing

Data preprocessing was pivotal in ensuring this dataset's cleanliness, consistency and suitability for model training.

3.2.1 Outliers Removal

In this phase, the outliers were identified, which were data points significantly deviating from the norm. These outliers were primarily associated with specific hours, roughly from 9 p.m. to 1 a.m., when trading activity was notably low (Table 1). Market prices exhibited erratic and unpredictable movements during these periods, posing challenges for accurate modelling, as shown in Figure 3.

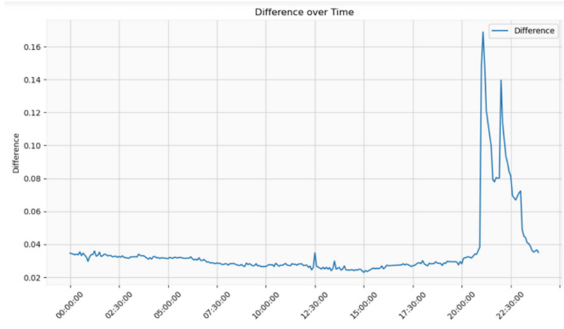

Figure 3: Hourly difference between Ask and Bed for 2022.

To address this issue, the noisy segment from the dataset was removed. This decision was motivated by the substantial spread between ask and bid prices during these hours, making trading economically impractical due to the significant potential loss.

These outliers were predominantly linked to noisy hours characterized by less predictable fluctuations due to reduced trading activity, so techniques like clipping or transformation were implemented. These outliers clipping involved capping extreme values to a predefined range, effectively limiting the influence of outliers on the analyses. Transformation techniques allowed for the adjustment of the scale or distribution of the data, rendering it more suitable for modelling. By identifying and handling outliers, especially those linked to noisy hours, the aim was to construct a dataset that reflected stable and typical trading conditions. This approach significantly bolstered the robustness and dependability of the predictive models. It ensured their effectiveness in capturing meaningful patterns while mitigating the impact of irregular fluctuations during specific hours.

In practical trading, taking profit from all potential target price fluctuations is not so important, but it is crucial to keep the deposit by avoiding incorrect entries into the market when predictions are not obvious and/or erroneous. Consequently, a strategic decision was made to exclude these hours data (with unpredictable volatility and large spreads) from the initial dataset and refrain from trading during this period. This novation further contributed to the better accuracy of the modelling efforts.

3.2.2 Feature Engineering

To enhance the model's predictive capability, additional feature engineering was performed, allowing, in addition to the short-term patterns, to add medium-term trends averaged over time. This involved creating new features based on the historical data, including but not limited to Moving averages of different time frames (Figure 4), Relative Strength Index (RSI) (Gumparthi, 2017) and Moving Average Convergence Divergence (MACD) (Aguirre et at., 2020).

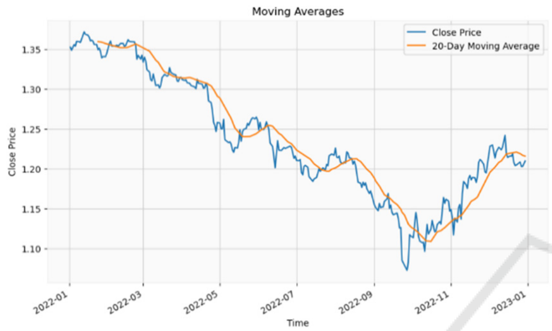

Figure 4: 20-Day Moving Average.

Table 1: The biggest difference between ask and bid throughout 2022.

Time (UTC)	Open_df_ask	Open_df_bid	Difference	time
2022-06-26 21:05:00	1.22880	1.22500	0.00380	21:05:00
2022-10-02 21:00:00	1.11850	1.11504	0.00346	21:00:00
2022-05-29 21:05:00	1.26437	1.26093	0.00344	21:05:00
2022-05-29 21:10:00	1.26437	1.26093	0.00344	21:10:00
2022-07-24 21:05:00	1.20187	1.19863	0.00324	21:05:00

This section describes the process of label generation, a crucial step in preparing the dataset for predictive modelling. The label, representing the target variable in the supervised learning framework, signifies the anticipated movement of currency rates for the subsequent trading period.

When calculating labels for intraday trading, it is essential to consider that the spread in the Forex market is usually quite large; therefore, to correctly train ML models and neural networks, it is necessary to take into account not only the currency rate behaviour of one price parameter (for example, only Bid or vs only Ask), as is often used in most other studies. This research meticulously assesses the absolute differences between the 'High Bid', 'Low Ask', and 'Open Ask' values as crucial indicators of currency rate behaviour. It is necessary to consider that market orders will make trades with a loss of spread (Figure 5). Thus, comparing the differences mentioned above, it is possible to evaluate whether it is profitable to enter into positions and whether the currency rate will move up or down in the selected timeframe is sufficient, taking into account the losses on the spread.

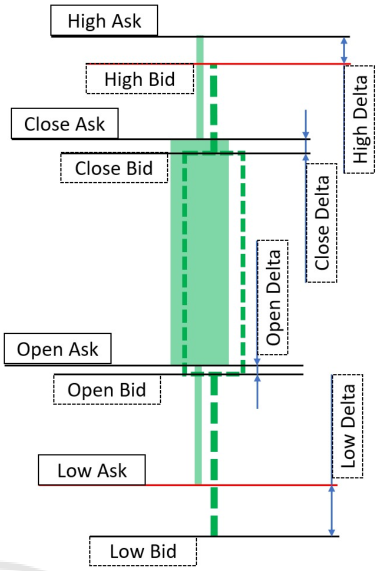

Figure 5: Calculation deltas, spreads and differences based on a random five-minute Japanese Candlestick.

The novation described above looks simple at first glance, but it allows authors to clearly mark datasets for further supervised learning. Subsequently, the 'Result' column is integrated into the dataset, which represents the currency rate movement for each specific period, denoted by 'Up' for an anticipated increase and 'Down' for an expected decrease. A reasonable shift of the differences and the 'Result' column by one position is made to ensure alignment with the relevant data points. This adjustment guarantees the 'Label' column encapsulates the movement prediction for the ensuing trading period. This preprocessing culminates in creating 4-hour blocks as parts of the initial dataset for using cross-validation over time.

Besides, within these blocks, there are no transitions between different trading days since this study analyses only intraday trading, which involves closing all positions overnight.

3.3 Model Implementation

The model implementation phase involves the following:

- converting conceptual frameworks into operational code.

Then, there is a cyclical process of the following crucial stages such as:

- data preprocessing within the framework of cross-validation over time,
- model development and hyperparameter improvements, and
- evaluation of each stage.

This process leverages the Python programming language, in conjunction with ML libraries such as Scikit-learn and TensorFlow-2 with Keras, to ensure the precise implementation of each constituent element. With the preprocessed dataset, various classification models are trained. In this research, five classical machine learning methods and two neural network architectures, such as K-Nearest Neighbors (K-NN), Logistic Regression, Support Vector Machines (SVM), Decision Trees, Random Forest, Multi-Layer Perceptron (MLP) and Long Short-Term Memory (LSTM) networks.

Achieving optimal model performance hinges on meticulous hyperparameter tuning. This entails a methodical approach involving extensive experimentation to pinpoint the ideal configuration that balances model complexity and accuracy. Through this iterative process, various aspects were fine-tuned, including setting the hyperparameters for the Random Forest model with 40 estimators, configuring the LSTM and MLP with one hidden layer, and optimising KNN with seven neighbours. This strategic refinement resulted in a significant enhancement of predictive accuracy.

In terms of training and validation, recognition is given to the sequential and time-dependent nature of the data. Traditional cross-validation techniques are deemed unsuitable, as they might disrupt the chronological order. Instead, this research uses a cross-validation approach that is often adopted for large time series, known as the 'Blocking Time Series Split'. It allows for evaluating models on unknown data from a learning point of view while ensuring that the information flow aligns with real-world scenarios.

For binary classification tasks, the focus is placed on metrics tailored to specific objectives. These include Accuracy, Precision, Recall, and F1-Score. Each metric is crucial in gauging the models' abilities to make accurate predictions while considering trade-offs between true positives and false positives.

Throughout this process, careful consideration is given to factors like model suitability, complexity, interpretability, resource constraints, and the trade-off between model complexity and performance. It is presumed that the models capture essential patterns in the data by iteratively adjusting model architectures, experimenting with different hyperparameters, and fine-tuning algorithms.

This comprehensive approach to model implementation, training, and validation, along with consideration of performance metrics, establishes a foundation for subsequent stages of analysis that are carried out when new trading data becomes available.

4 RESULTS

The performance of various machine learning and artificial neural network models was evaluated for the binary classification task of predicting currency rate movements (up or down). The results obtained in this research are presented in Table 2.

Table 2 also demonstrates a comparison of the results obtained in this research with a state-of-the-art paper with similar research objectives, such as (Pande et al., 2021), which found that the KNN algorithm outperformed the Naïve Bayes algorithm in terms of recall, precision, accuracy, and f-score. The Naïve Bayes algorithm yielded an accuracy of 50%, precision of 43%, recall of 55%, and f-score of 49%, while the KNN algorithm achieved an accuracy of 53%, precision of 54%, recall of 56%, and f-score of 55%. This comparison emphasises the superiority of some neural network architectures, such as LSTM, over classical ML methods in the context of the movement direction prediction of forex price for GBP/USD for intraday trading.

The accuracy, although seemingly modest at first glance, is commendable within the real-time prediction of the direction of short-term movements of exchange prices. Attaining an accuracy exceeding 51% is considered a favourite in this context. This contrasts with tasks like image recognition, where an accuracy below 90% might be considered suboptimal.

It is important to note that there is much more research on predicting specific price values (regression task) of financial instruments (mainly for medium-term time frames) than predicting the direction of future movements (classification task) for short-term intraday trading, which is potentially more profitable (Milke et al., 2020). Predicting short-term movements' directions of the financial market is inherently more challenging due to the dynamic and complex nature of the domain. The significance of this threshold (51%) lies in the fact that the predictions yield profits more frequently than not. In intraday trading, where transactions occur rapidly, this establishes a statistically advantageous position, significantly outweighing the 1% due to using the compound interest formula with multiple small increments of the deposit. It underscores the effectiveness of these models in navigating the inherently dynamic nature of financial markets, even given their intraday volatility, taking into account that intraday trading does not take the risk of a possible overnight gap.

Table 2: Results of AI models.

Algorithms\ Metrics	Accuracy	Precision	Recall	F1
Machine learning (ML) classification algorithms				
KNN	54%	56%	46%	51%
Logistic Regression	52%	66%	25%	36%
SVM	52%	52%	97%	68%
Decision Trees	55%	59%	40%	48%
Random Forest	60%	77%	31%	44%
Artificial Neural Network (ANN) models				
MLP	52%	52%	96%	68%
LSTM	63%	72%	48%	57%
(Pande et al., 2021) results				
Naïve Bayes	50%	43%	55%	49%
KNN	53%	54%	56%	55%

In summary, the variance in performance among these algorithms can be attributed to their inherent characteristics. Models like Logistic Regression and Random Forest excel in precision but encounter challenges with recall, whereas others like LSTM strike a balance between these metrics and accuracy, which reached 63%. The intricacies of predicting financial markets are rooted in their complex, dynamic, constantly changing and noisy nature, where past movements may not always anticipate future trends. These results offer valuable insights into the strengths and weaknesses of each algorithm in the context of the currency market prediction.

5 CONCLUSIONS

This research, focusing on foreign exchange rate prediction of GBP/USD currency rate, uses several machine learning and artificial neural network techniques; data was systematically pre-processed, and models were implemented to gain insights into the intricate dynamics of financial markets.

The technical outcomes reveal nuanced performance across various models. Each model, ranging from K-Nearest Neighbors to Long Short-Term Memory Networks, demonstrated distinctive strengths and limitations. While some exhibited commendable accuracy in market price prediction, others necessitated further refinement and parameter tuning. The selection of models was thoughtfully guided by their suitability for the temporal granularity of the data.

The developed framework fulfils the original technical research requirements: by utilising both traditional machine learning models and neural networks, a balanced approach was achieved, harnessing the strengths of each methodology. This endeavour has deepened technical proficiency and conferred a better understanding of the intricacies of the financial market.

As a further exploration, integrating additional external factors like macroeconomic indicators and news sentiment analysis and adding data from the order book could bolster predictive accuracy. Exploring advanced deep learning architectures and ensemble techniques holds potential for additional insights. Furthermore, real-time data integration and creating a user-friendly interface could significantly enhance the practical utility of the framework.

The authors hope that, with a focus on future refinements, this research will continue to serve as a valuable resource for both researchers and practitioners in the field of financial forecasting.

REFERENCES

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M. and Kudlur, M. (2016). {TensorFlow}: a system for {Large-Scale} machine learning. *In 12th USENIX symposium on operating systems design and implementation (OSDI 16)* (pp. 265-283).

Abedin, M.Z., Moon, M.H., Hassan, M.K. and Hajek, P. (2021). Deep learning-based exchange rate prediction during the COVID-19 pandemic. *Annals of Operations Research*, pp.1-52.

Aguirre, A.A.A., Medina, R.A.R. and Méndez, N.D.D. (2020). Machine learning applied in the stock market through the Moving Average Convergence Divergence (MACD) indicator. *Investment Management and Financial Innovations,* 17(4), p.44.

Aha, D.W., Kibler, D. and Albert, M.K. (1991). Instance-based learning algorithms. *Machine learning, 6,* pp.37-66.

Breiman, L. (1996). Bagging predictors. *Machine learning, 24*, pp.123-140.

Cessie, S.L. and Houwelingen, J.V. (1992). Ridge estimators in logistic regression. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 41(1), pp.191-201.

Chollet, F. and others. (2015). Keras. Avable at: https://github.com/fchollet/keras. / (Accessed: 25 October 2023).

Engle, R.F., Ghysels, E. and Sohn, B. (2013). Stock market volatility and macroeconomic fundamentals. *Review of Economics and Statistics, 95(3),* pp.776-797.

Gumparthi, S. (2017. Relative strength index for developing effective trading strategies in constructing optimal portfolio. *International Journal of Applied Engineering Research, 12(19),* pp.8926-8936.

Goodfellow, I., Bengio, Y. and Courville, A. (2016). *Deep learning.* MIT press.

Dukascopy Bank Sa. (2023). Historical data feed, Swiss Forex Bank-ECN broker. *Swiss FX Trading Platform Dukascopy Bank SA.* Available at: https://www.dukascopy.com/swiss/english/marketwatch/historical/ (Accessed: 25 October 2023).

Hoque, K.E. and Aljamaan, H. (2021). Impact of hyperparameter tuning on machine learning models in stock price forecasting. *IEEE Access, 9*, pp.163815-163830.

Keerthi, S.S., Shevade, S.K., Bhattacharyya, C. and Murthy, K.R.K. (2001). Improvements to Platt's SMO algorithm for SVM classifier design. *Neural computation, 13(3)*, pp.637-649.

Islam, M.S. and Hossain, E. (2021). Foreign exchange currency rate prediction using a GRU-LSTM hybrid network. *Soft Computing Letters, 3*, p.100009.

Long, W., Lu, Z. and Cui, L. (2019). Deep learning-based feature engineering for stock price movement prediction. *Knowledge-Based Systems, 164*, pp.163-173.

MacKenzie, D. (2019). How Fragile Is Competition in High-Frequency Trading. *Tabbforum, March*, 26.

Maqsood, H., Mehmood, I., Maqsood, M., Yasir, M., Afzal, S., Aadil, F., Selim, M.M. and Muhammad, K. (2020). A local and global event sentiment based efficient stock exchange forecasting using deep learning. *International Journal of Information Management, 50,* pp.432-451.

Milke, V., Luca, C., Wilson, G. and Fatima, A. (2020). Using Convolutional Neural Networks and Raw Data to Model Intraday Trading Market Behaviour. *In ICAART* (2) (pp. 224-231).

Nauta, M., Trienes, J., Pathak, S., Nguyen, E., Peters, M., Schmitt, Y., Schlötterer, J., van Keulen, M. and Seifert, C. (2023). From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable ai. *ACM Computing Surveys, 55(13s)*, pp.1-42.

Pande, K.S.Y., Divayana, D.G.H. and Indrawan, G. (2021). March. Comparative analysis of naïve bayes and knn on prediction of forex price movements for gbp/usd currency at time frame daily. *In Journal of Physics: Conference Series* (Vol. 1810, No. 1, p. 012012). IOP Publishing.

Pedregosa, F. et al. (2011). Scikit-learn: Machine learning in Python*.* *Journal of machine learning research, 12(Oct),* pp.2825–2830.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V. and Vanderplas, J. (2011). Scikit-learn: Machine learning in Python. *The Journal of machine Learning research, 12*, pp.2825-2830.

Samek, W., Montavon, G., Vedaldi, A., Hansen, L.K. and Müller, K.R. eds. (2019). *Explainable AI: interpreting, explaining and visualizing deep learning* (Vol. 11700). Springer Nature.

Schierholt, K. and Dagli, C.H. (1996). March. Stock market prediction using different neural network classification architectures. *In IEEE/IAFE 1996 Conference on Computational Intelligence for Financial Engineering (CIFEr)* (pp. 72-78). IEEE.

Vignesh, C.K. (2018). Applying machine learning models in stock market prediction. *IRJET, Journal*.

Waskom, M. et al. (2017). Mwaskom/seaborn: v0.8.1 (September 2017), Zenodo. Available at: https://doi.org/10.5281/zenodo.883859.

Zhanggui, Z., Yau, H. and Fu, A.M. (1999). July. A new stock price prediction method based on pattern classification. In IJCNN'99. International Joint Conference on Neural Networks. Proceedings (Cat. No. 99CH36339) (Vol. 6, pp. 3866-3870). IEEE.