Forex Price Predictions using Hybrid TCN-LSTM and LSTM-TCN Models

[1*]Caroline, [2]Wulan Sri Lestari, [3]Mustika Ulina
[1]Department of Information System, Universitas Mikroskil, Indonesia
[2,3]Department of Information Technology, Universitas Mikroskil, Indonesia
Email: [1]caroline.chong@mikroskil.ac.id, [2]wulan.lestari@mikroskil.ac.id, [3]mustika.ulina@mikroskil.ac.id

Article Info
Article history:
Received Sep 21th, 2025
Revised Oct 24th, 2025
Accepted Nov 25th, 2025

Keyword:
Foreign Exchange
Hybrid
Long Short-Term Memory
Prediction
Temporal Convolutional

ABSTRACT
Forecasting financial market prices, particularly foreign exchange (forex) rates, remains a substantial difficulty due to the market's inherent unpredictability, intricacy, and turbulent characteristics. By combining the Temporal Convolutional Network (TCN) and Long Short-Term Memory (LSTM) models into a hybrid framework, this study overcomes this difficulty and improves prediction accuracy. The MinMaxScaler function was used to standardize the input data prior to training, thereby bringing all values within a range of 0 to 1. An 80% training segment and a 20% testing segment were then separated from the prepared dataset. We tested two different hybrid architectures, the LSTM-TCN and the TCN-LSTM, with the EUR/USD, AUD/USD, and GBP/USD value pairs. With uniform parameters applied to both models during training, the Root Mean Squared Error (RMSE) measure was used for all performance evaluations to ensure a fair comparison and determine which model performed better. The LSTM-TCN architecture proved to be the superior predictor on the testing set. It recorded a lower average RMSE of 0.003911. This result contrasts with the TCN-LSTM model's performance, which yielded a higher average RMSE of 0.004181.

Corresponding Author:
Caroline,
Department of Information System,
Universitas Mikroskil,
Jl. M.H.Thamrin No. 112, 124, 140 Medan, Indonesia, 20212.
Email: caroline.chong@mikroskil.ac.id

DOI: http://dx.doi.org/10.24014/ijaidm.v8i3.38354

1. INTRODUCTION

Today, the foreign exchange market, or Forex, is the biggest financial market in the world, distinguished by its enormous daily trading volume [1]-[4]. Its popularity continues to increase due to high liquidity, easy accessibility, and significant profit potential [5]. However, despite the potential for huge profits in Forex trading, the market is characterized by an extremely high level of volatility, which inherently poses a significant risk of loss [6]. The uncertainty, complexity, and chaotic nature of the data make accurate Forex price prediction exceptionally difficult [5], [7]. Therefore, accurately predicting price movements remains a central and significant challenge in the world of Forex trading.

The urgency of this research stems directly from the intrinsic characteristics of the Forex market. First, high price volatility, characterized by fast and extreme price movements, makes the time series data highly unstable and difficult to model accurately [6]. Furthermore, the Forex market exhibits an extreme level of complexity, with prices influenced by numerous factors, including macroeconomic indicators, geopolitical events, central bank interventions, and complex, nonlinear market sentiment [5], [7]. Most crucially, these markets often exhibit a strong chaotic and non-linear nature [5], [7]. This means that small changes to initial inputs can result in large, unpredictable differences in output, rendering the task of predicting future prices exceptionally difficult and risky if relying solely on traditional statistical models.

Addressing this challenge requires advanced deep learning models capable of capturing both long-term sequential dependencies and local feature hierarchies [5], [7].

For predicting Forex prices, a variety of deep learning models have been thoroughly investigated, including Convolutional Neural Networks (CNNs), Deep Neural Networks (DNNs), Long Short-Term Memory (LSTM), and Gated Recurrent Units (GRUs) [8]-[11]. Ulina et al. (2020) demonstrated improved prediction accuracy by integrating CEEMDAN and IFA-LSTM, achieving superior results compared to using standalone LSTM and IFA-LSTM models [12]. Islam et al. (2020) analyzed the existing literature, noting that GRU and LSTM models were the most commonly utilized by researchers [8]. Similarly, Dautel et al. (2020) employed a basic Recurrent Neural Network (RNN) model but reported that LSTM and GRU yielded better results [13]. Fisichella (2021) affirmed these general findings [9]. However, Hu et al. (2021), conducting a similar comparative analysis that year, presented contrasting results, finding that the LSTM and DNN models delivered the best performance compared to other models [10]. This period also marked the beginning of hybrid model development, notably with Islam’s (2021) GRU-LSTM model, which consistently demonstrated improved prediction values [7].

The pursuit of optimal Forex prediction models continued actively into 2022. Junior et al. introduced an LSTM model that utilizes a two-layer stacked (TLS) architecture, which demonstrated better performance compared to the CEEMDAN-IFALSTM hybrid [14]. At the same time, Panda et al. (2022) discovered that the CNN-Random Forest (RF) hybrid model performed better than more conventional options, including MLP, ARIMA, and Linear Regression (LR) [15]. Research persisted in 2023, the same researcher Junior et al., conducting a comparative analysis, similar to Islam and Hu, confirming that LSTM and Artificial Neural Networks (ANN) remained the most frequently utilized models in predictions [11]. In the same year, Caroline et al. (2023) used a hybrid LSTM and Temporal Convolutional Network (TCN) model to study stock price prediction; the findings showed that the LSTM-TCN configuration was better for a 90% training and 10% testing split, while the TCN-LSTM configuration performed well with an 80% training and 20% testing data split [16].

In a 2024 study, Mahmud et al. conducted Forex price predictions using various single models, finding that the Random Forest Regressor (RFR) outperformed other machine learning algorithms. Conversely, Bidirectional LSTM was identified as the best performer among the tested deep learning models, including LSTM, CNN, and GRU [17]. Also in 2024, Zitis et al. conducted Forex research that emphasized sophisticated data preprocessing, first utilizing the Hurst Exponent (H) and Fuzzy Entropy (FuzzyEn) to process input data, which was subsequently fed into a deep learning architecture, ultimately finding that LSTM and GRU models delivered superior results [18]Saghafi et al. recently published a Dual-Input LSTM model in 2025 that was created especially for foreign exchange forecasting. When compared to models that exclusively used technical data (T-LSTM), fundamental data (F-LSTM), or a mix of both in a single input stream (FT-LSTM), this model significantly reduced the Root Mean Square Error (RMSE) score by 24% [19]. Observing the variety of successful models applied to Forex to date, the authors are motivated to contribute to this field by testing other hybrid architectures that have previously demonstrated strong performance in related financial domains, specifically forex predictions.

Applying deep learning models to increase the precision of intricate and erratic Forex price forecasts is the aim of this research. The LSTM model is employed because it can identify long-term trends in sequential data [20], [21]. The TCN model, conversely, was chosen for its efficiency in handling long-term dependencies compared to traditional methods [10]-[12][10]-[12]. To leverage the respective strengths of each model, this study proposes a hybrid approach comprising LSTM-TCN and TCN-LSTM architectures. The LSTM-TCN architecture first employs the LSTM layer to capture long-term temporal patterns, which are then passed to the TCN as an additional feature extractor [10], [13]-[15]. Conversely, the TCN-LSTM model uses the TCN for initial feature extraction before the LSTM further processes these features to understand complex market dynamics [16], [22], [27]. Furthermore, these models will be tested on the major currency pairs namely, EUR/USD, GBP/USD, and AUD/USD to assess the effectiveness of the proposed approach [12]. Model evaluation was conducted using the Root Mean Squared Error (RMSE) metric to quantify prediction accuracy. This approach is expected to yield a more accurate and efficient model for predicting Forex prices, thereby providing new insights into the application of deep learning within the financial sector.

2. RESEARCH METHOD

The research methodology proceeds through distinct stages, which are detailed in Figure 1. This process begins with the foundational steps of data collection, preprocessing, and dataset partitioning. The work then advances to hybrid model creation and testing, which ultimately leads to the analysis of results and formulation of conclusions. This process begins with the foundational steps of data collection, preprocessing, and dataset partitioning. The work then advances to hybrid model creation and testing, which ultimately leads to the analysis of results and formulation of conclusions.

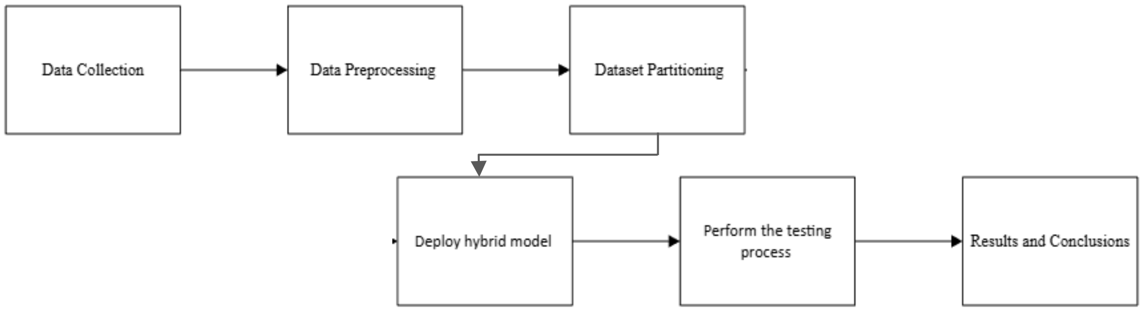

Figure 1. Research Flow Diagram

2.1. Data Collection

In this research, the authors utilize historical price data of major Forex currency pairs (EUR/USD, GBP/USD, and AUD/USD), sourced from dukascopy.com. The dataset spans from January 1, 2010, to December 30, 2019, and includes 2,623 daily records for EUR/USD, 2,627 for AUD/USD, and 2,622 for GBP/USD. The specific period (January 1, 2010 – December 30, 2019) and the data source were selected to ensure direct comparability with the dataset used in previous foundational studies[12]. Following prior research, the authors restricted their analysis to the four essential attributes: the opening, closing, high, and low prices. A sample of the EUR/USD data utilized is presented in Table 1.

Table 1. EUR/ USD Historical Data

Date	Open	High	Low	Close
01.01.2010	1.43216	1.43356	1.43178	1.43335
04.01.2010	1.43024	1.44556	1.42559	1.4412
05.01.2010	1.44112	1.44834	1.43445	1.43639
06.01.2010	1.4363	1.44342	1.42807	1.44062
07.01.2010	1.44053	1.44432	1.42976	1.43045
08.01.2010	1.43063	1.44382	1.42616	1.44083
11.01.2010	1.44265	1.45539	1.44049	1.45111
12.01.2010	1.45108	1.45479	1.44507	1.44816
13.01.2010	1.44832	1.45797	1.4455	1.45082
14.01.2010	1.45089	1.45546	1.4445	1.44969
15.01.2010	1.4498	1.45111	1.43347	1.43849
….	….	….	….	….
….	….	….	….	….
20.12.2019	1.11194	1.11248	1.10662	1.10769
23.12.2019	1.1081	1.10958	1.10699	1.10895
24.12.2019	1.10895	1.10938	1.1069	1.10879
25.12.2019	1.10876	1.10885	1.1073	1.10821
26.12.2019	1.10885	1.11088	1.10821	1.10974
27.12.2019	1.10973	1.11883	1.10964	1.11713
30.12.2019	1.11736	1.12207	1.11718	1.11985
31.12.2019	1.11985	1.12391	1.11974	1.12076

2.2. Data Preprocessing

Normalization, an essential preprocessing step, was performed on the dataset once it was acquired. The Scikit-learn library's MinMaxScaler function was used to carry out this procedure. MinMaxScaler is a feature scaling method that consistently converts the data into a predetermined range, namely 0–1. This standardization is primarily intended to prevent disparities in the absolute amount of currency prices from unduly affecting the forecasting model [28].

2.3. Data partitioning

The dataset was then split, with 80% used for training and the remaining 20% reserved for testing. This split ratio was chosen to maintain consistency and direct comparability with the data partitioning scheme employed in the foundational prior research [12]. Specifically, the training subset is used to optimize the model and find the best parameters, which are then tested on the separate, independent testing subset. Using identical data division and parameters across studies is necessary to ensure the results are robustly comparable.

2.4. Deploy hybrid model

The models deployed in this study are the hybrid LSTM-TCN and TCN-LSTM architectures. Seven layers make up the selected architecture, which effectively combines two LSTM layers, two Dense levels, and three TCN layers. TCN uses a filter number of 200 neurons, kernel size 3 with a dropout of 0.25 and activates the ReLU function while LSTM uses 100 neurons then enters a dense screen and is finally equipped with an Adam optimizer. This configuration was selected for consistency with prior research [16]. These two deep learning models will be briefly described as follows:

2.4.1. Long Short-Term Memory (LSTM)

To preserve important information in time-series data, the LSTM network employs an advanced internal architecture. Three specialized logic gates, the input, output, and forget gates, that control information flow and retention are found at the heart of each LSTM cell. Each of these gates utilizes a sigmoid neural network layer that produces a value between 0 and 1 to manage the information. This value determines the extent to which information is allowed to pass through, where 0 signifies complete blockage and 1 signifies full passage. Specifically, the forget gate controls the retention of past cell information, allowing for either complete deletion (0) or full preservation (1). The sigmoid input gate determines which values need to be adjusted during the cell state update process, and a tanh layer generates a vector with new candidate values. The current state of the cell is then selectively incorporated with these candidate values. Lastly, the output gate creates the layer's final output by combining the updated and filtered cell state [16], [22].

The LSTM architecture's primary advantage is its inherent capacity to preserve long-term relationships in time-series data, a feat that overcomes the vanishing gradient problem common to traditional RNNs. Furthermore, LSTMs can natively handle input sequences of variable length and facilitate efficient end-to-end data processing, often mitigating the need for extensive manual feature engineering or preprocessing steps, which contributes to efficiency [16], [22].

Drawbacks of LSTMs stem mainly from their complex architecture. This makes them inherently more challenging to train and fine-tune compared to simpler deep learning models, often necessitating rigorous selection and tuning of hyperparameters. Due to their complexity, LSTMs can be limited in their usage in real-time applications because they require a lot of computation, especially when dealing with huge datasets. Last but not least, the intricate internal structure makes it difficult to properly comprehend and communicate the model's prediction decisions to others, hence impeding interpretability [16], [22].

As illustrated in Figure 2, the LSTM model is designed with specific functions that allow for the retention of long-term sequence dependencies [16], [22].

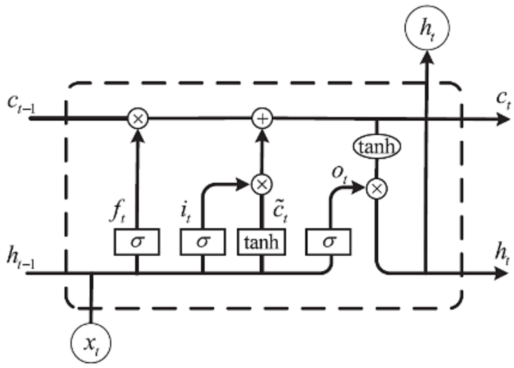

Figure 2. LSTM Model Architecture

2.4.2. Temporal Convolutional Network (TCN)

The TCN represents a hybrid deep learning model, seamlessly integrating foundational concepts from both Recurrent Neural Networks (RNNs) and CNNs. Two fundamental design principles primarily define the TCN architecture [16], [29]. First, its convolutions are strictly causal, which prevents any information leakage from future timesteps into past predictions, thereby ensuring that predictions for a given time depend only on known inputs up to that point. Secondly, the design makes it easier to translate a length of input sequence to an identically lengthened output sequence. TCN achieves a large and flexible receptive field by utilizing dilated convolutions. Furthermore, it incorporates residual connections, which are vital for enabling the network to learn and modify identity mappings easily; this is crucial for deep networks as it helps mitigate the risk of exploding gradients during training [30].

The TCN architecture offers several key benefits. Its core causality prevents look-ahead bias by ensuring strictly forward-temporal information flow. The model's convolutional nature also allows for straightforward parallel processing of inputs, leading to faster and more efficient training compared to inherently sequential models. The receptive field size is highly tunable, its value depending on the filter size, the dilation factors, and the network's number of layers. This flexibility grants specific control over the model's memory capacity for various domain requirements. Additionally, the complexity and number of network kernels depend solely on the number of layers, independent of the input sequence length. Despite these advantages, the TCN has two primary drawbacks. It typically demands more data storage compared to models like RNNs, and it may face challenges with proper domain transfer, limiting its effectiveness when applied to a significantly different type of sequential data than it was originally trained on [16], [30]. The complete TCN model architecture is visually represented in Figure 3.

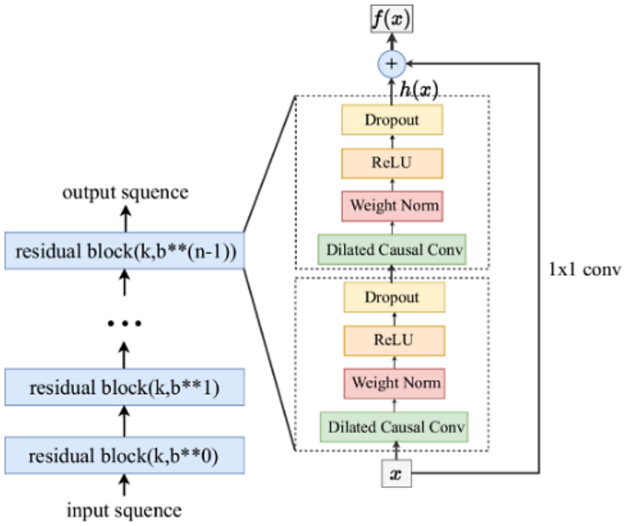

Figure 3. TCN Model Architecture

2.5. Perform the testing process

Improved prediction accuracy in financial time series has led to the development of hybrid models. LSTM model has long been the benchmark for time series prediction thanks to its ability to overcome vanishing gradient problems and maintain long-term dependencies. The effectiveness of LSTM is confirmed by previous studies, in which it and GRU became the most widely used model by Forex researchers which can be seen in the previous introduction. However, to cope with extreme volatility, hybrid models have become standard. For example, Ulina et al. [12] demonstrated an improvement in the accuracy of Forex predictions by integrating CEEMDAN and Improved FA-LSTM, outperforming standalone models.

TCN model, which utilizes dilated causal convolutions to achieve large receptive fields and efficient parallel processing, offers a unique advantage over traditional repetitive models (RNNs). TCN has been successfully implemented in a variety of time series tasks, including network traffic prediction [22]. To combine the sequential modeling power of LSTM with the extraction efficiency of TCN features, a hybrid architecture of TCN-LSTM and LSTM-TCN has been explored. The study by Caroline et al. was the main motivation for this study, where such hybrid configurations were used for stock price predictions, and it has been proven that architectural sequences greatly influence performance. Similar hybrid implementations have also been shown to be effective in other contexts, such as load forecasting by Heng et al. [27]. With this literature, the study focuses on the systematic testing of the TCN-LSTM and LSTM-TCN hybrid models, aimed at determining the optimal configuration for Forex price prediction.

The authors perform testing on the trained model in this process. For the purpose of robust comparison, RMSE was calculated and utilized as the principal data analysis metric. The proper hyperparameters decided upon and used to train the model in this step determine the performance and results.

3. RESULTS AND ANALYSIS

The results of the close price prediction will be compared with the real price of the close column to compare the prediction and actual. A comparison graph between the predicted value close and the real value of the AUDUSD with the TCNLSTM model can be seen in Figure 4.

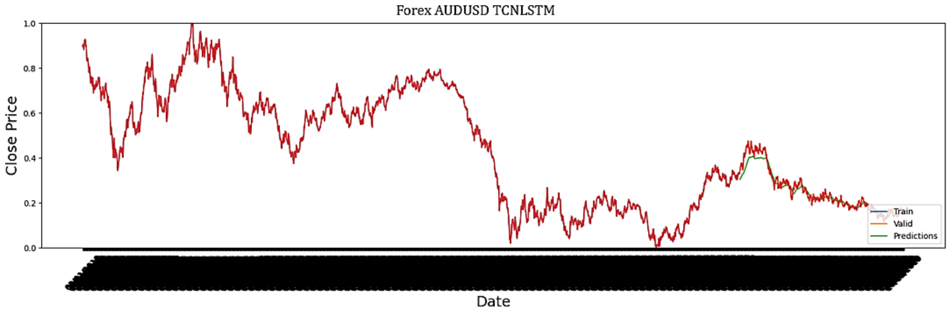

Figure 4. Forex AUD/ USD TCNLSTM

The x-axis indicates the time period, while the y-axis indicates the closing value. The TCN-LSTM and LSTM-TCN hybrid models were tested on the EUR/USD, AUD/USD, and GBP/USD currency pairs using identical parameters. The best-performing hybrid model was identified using the RMSE as the comparison statistic. Table 2 displays the complete test results.

Table 2. TCN-LSTM and LSTM-TCN Test Results

Training data, Testing data		80,20			
Epochs		2			
Pair of Currency		EUR/ USD	AUD/ USD	GBP/ USD	Average
Batch size		256			
Experiment	TCN-LSTM	0.000267	0.008197	0.004078	0.004181
	LSTM-TCN	0.004161	0.004336	0.003237	0.003911
	Average	0.002213	0.006266	0.003657	0.004046

The configuration and training of the hybrid models (TCN-LSTM and LSTM-TCN) were executed with rigorous attention to maintaining uniformity and fairness across both architectures. This rigorous approach was crucial in ensuring that the resultant difference in performance (measured by RMSE) could be attributed solely to the sequential configuration, rather than to discrepancies in hyperparameter tuning. The parameters used were selected based on standard practices in time-series forecasting and are considered sufficient and representative for the objective of architectural comparison.

The dataset was partitioned into an 80% training segment and a 20% testing segment, a ratio standard in deep learning time-series studies that ensures a robust validation process. The models utilized the four essential technical attributes (Open, High, Low, Close prices) as input features, focusing the study on the architectural performance rather than complex feature engineering. For training, an Adam Optimizer was utilized, known for its adaptive learning rate capabilities. Furthermore, a Batch Size of 256 was employed to balance computational efficiency and learning stability, and 2 Epochs were deemed sufficient for convergence, given the considerable volume of the input data. In summary, the implementation of these uniform, representative parameters ensures that the comparison between the TCN-LSTM and LSTM-TCN architectures is robust and conclusive, allowing the findings regarding the superiority of the LSTM-TCN sequence to be validated reliably.

The comparative testing of the hybrid models demonstrated a clear distinction in performance, with the LSTM-TCN architecture proving to be the superior predictor. The model achieved a lower average Root Mean Squared Error (RMSE) of 0.003911, compared to the TCN-LSTM's 0.004181. This suggests that the LSTM-TCN architecture is more effective at capturing both the long-term dependencies (via LSTM) and the local feature patterns (via TCN) essential for accurate financial time series forecasting. Furthermore, an analysis of the individual currency pairs revealed significant differences in predictability. The EUR/USD currency pair exhibited the lowest RMSE value (0.002213 average), indicating that this pair is the easiest to predict using the developed models compared to the other two pairs. Conversely, the AUD/USD currency pair exhibited the largest RMSE value (0.006266 average), indicating that its price movements contain greater volatility or are influenced by more complex factors, making it the most challenging pair to predict accurately.

The primary finding that the LSTM-TCN sequence is superior highlights the importance of architectural sequence in hybrid deep learning for complex time series data. The LSTM-TCN model first utilizes the LSTM layer to capture long-term temporal patterns, which is critical in Forex data that exhibits strong chaotic and non-linear characteristics. These extracted sequential features are then passed to the TCN for more detailed local feature extraction and pattern recognition.This approach is more effective than the TCN-LSTM sequence, which performs initial feature extraction via TCN before the LSTM attempts to process these pre-processed features for complex market dynamics. The superior result suggests that prioritizing the modeling of long-term sequence dependencies is more crucial for minimizing prediction error in volatile Forex markets. This observation is consistent with related financial domain research, where Caroline et al. [16] demonstrated that configuration sequence (LSTM-TCN vs. TCN-LSTM) significantly influences stock price prediction performance.

The performance of the current hybrid model can be benchmarked against prior research, specifically the CEEMDAN-IFA-LSTM method, as detailed in Table 3.

Table 3. RMSE Results Comparison

	EUR/ USD	AUD/ USD	GBP/ USD
CEEMDAN-IFA-LSTM	0.0792	0.0071	0.0103
TCN-LSTM	0.0003	0.0082	0.0041
LSTM-TCN	0.0042	0.0043	0.0032

The performance of the proposed hybrid architecture demonstrates a highly significant improvement over prior research benchmarks. Specifically, the LSTM-TCN model's RMSE for EUR/USD (0.0042) represents a massive reduction compared to the 0.0792 RMSE reported by the CEEMDAN-IFA-LSTM method. The substantial decrease in error across all currency pairs confirms that the synergistic combination of TCN and LSTM is a highly effective strategy for this deep learning task. This combination successfully maximizes the unique strengths of each component, the LSTM’s ability to handle sequence memory and the TCN's efficiency in extracting hierarchical features, resulting in a model that drastically outperforms the previously established method.The findings also align with Mahmud et al. (2024) [17], who found that the Bidirectional LSTM was the best single deep learning model. Our study advances this finding by demonstrating that the LSTM-TCN hybrid can further leverage these sequential strengths by incorporating a complementary convolutional component (TCN), resulting in greater predictive accuracy.

The primary contribution (strength) of this research is the empirical demonstration of the architectural superiority of the LSTM-TCN sequence over its inverted counterpart, TCN-LSTM, in Forex price forecasting. This finding, coupled with the resulting low prediction errors across all major currency pairs, confirms the high potential of hybrid deep learning strategies for this task. Furthermore, the analysis revealed valuable insights into market behavior, demonstrating that the models' effectiveness varies significantly by currency pair; specifically, EUR/USD was the easiest to predict while AUD/USD was the most difficult, a variance that highlights inherent differences in market volatility and complexity.

Despite the strong performance, the research acknowledges certain limitations. First, in terms of Data Scope, the models were trained solely on four essential technical attributes (Open, Close, High, and Low prices). Price prediction can often be significantly improved by incorporating complex fundamental and macroeconomic indicators, which were omitted in this analysis. Second, the Model Complexity inherent in the hybrid LSTM and TCN architectures requires substantial computational resources for both training and deployment. This high resource demand may potentially hinder the model's application in real-time, high-frequency trading scenarios where computational latency is critical. A deeper limitation is that they may not be the absolute optimal settings for maximum performance on all currency pairs (e.g., AUD/USD's higher RMSE of 0.004336 suggests potential for further tuning). This sets the stage for future work.

Based on these findings, future work should involve further developing these models by incorporating additional mechanisms, such as the attention mechanism, which has been successful in related time series research. Furthermore, exploring other effective model combinations or integrating the model with fundamental economic data will be crucial to push the boundaries of Forex predictability.

4. CONCLUSION

This research successfully achieved its aim of applying hybrid deep learning models to increase the precision of intricate and erratic Forex price forecasts for the EUR/USD, AUD/USD, and GBP/USD pairs. The study confirmed that the hybrid LSTM-TCN architecture provides a notable increase in prediction accuracy compared to previous single-model approaches and the established CEEMDAN-IFA-LSTM benchmark. Crucially, the LSTM-TCN proved to be the superior configuration, recording a lower average Root Mean Squared Error (RMSE) of 0.003911. This outcome highlights the importance of sequencing the LSTM layer first to effectively capture long-term temporal patterns before the TCN performs local feature extraction, directly answering the research question regarding the optimal hybrid sequence.

However, the study acknowledges a key weakness: the limitation of input data to only technical price attributes (Open, Close, High, Low). The current model does not account for the influence of fundamental or macroeconomic factors, which are known to drive significant, large-scale price changes in the Forex market.

Based on the strong results obtained from the hybrid architecture, the potential for further development is substantial. Future research should focus on developing the models further by integrating a Dual-Input strategy that includes fundamental and macroeconomic indicators to create a more comprehensive feature set. Additionally, exploring advanced mechanisms such as the Attention Mechanism is essential to selectively weigh the importance of time steps or features, potentially improving performance beyond the current hybrid framework and continually pushing the boundaries of Forex predictability.

REFERENCES

[1] A. Chaboud, D. Rime, and V. Sushko, "The foreign exchange market," in Research Handbook of Financial Markets, Edward Elgar Publishing, 2023, pp. 253–275. doi: 10.4337/9781800375321.00020.
[2] M. Ayitey Junior, P. Appiahene, O. Appiah, and C. N. Bombie, "Forex market forecasting using machine learning: Systematic Literature Review and meta-analysis," J Big Data, vol. 10, no. 1, p. 9, Jan. 2023, doi: 10.1186/s40537-022-00676-2.
[3] R. C. Michie, Forex Forever. Oxford University PressOxford, 2024. doi: 10.1093/oso/9780198903697.001.0001.
[4] G.-J. Wang, L. Wan, Y. Feng, C. Xie, G. S. Uddin, and Y. Zhu, "Interconnected multilayer networks: Quantifying connectedness among global stock and foreign exchange markets," International Review of Financial Analysis, vol. 86, p. 102518, Mar. 2023, doi: 10.1016/j.irfa.2023.102518.
[5] J. E. Sandubete, L. Beleña, and J. C. García-Villalobos, "Testing the Efficient Market Hypothesis and the Model-Data Paradox of Chaos on Top Currencies from the Foreign Exchange Market (FOREX)," Mathematics, vol. 11, no. 2, p. 286, Jan. 2023, doi: 10.3390/math11020286.
[6] B. H. A. Khattak et al., "A Systematic Survey of AI Models in Financial Market Forecasting for Profitability Analysis," IEEE Access, vol. 11, pp. 125359–125380, 2023, doi: 10.1109/ACCESS.2023.3330156.
[7] M. S. Islam and E. Hossain, "Foreign exchange currency rate prediction using a GRU-LSTM hybrid network," Soft Computing Letters, vol. 3, p. 100009, Dec. 2021, doi: 10.1016/j.socl.2020.100009.
[8] Md. S. Islam, E. Hossain, A. Rahman, M. S. Hossain, and K. Andersson, "A Review on Recent Advancements in FOREX Currency Prediction," Algorithms, vol. 13, no. 8, p. 186, Jul. 2020, doi: 10.3390/a13080186.
[9] M. Fisichella and F. Garolla, "Can Deep Learning Improve Technical Analysis of Forex Data to Predict Future Price Movements?," IEEE Access, vol. 9, pp. 153083–153101, 2021, doi: 10.1109/ACCESS.2021.3127570.
[10] Z. Hu, Y. Zhao, and M. Khushi, "A Survey of Forex and Stock Price Prediction Using Deep Learning," Applied System Innovation, vol. 4, no. 1, p. 9, Feb. 2021, doi: 10.3390/asi4010009.
[11] M. Ayitey Junior, P. Appiahene, O. Appiah, and C. N. Bombie, "Forex market forecasting using machine learning: Systematic Literature Review and meta-analysis," J Big Data, vol. 10, no. 1, p. 9, Jan. 2023, doi: 10.1186/s40537-022-00676-2.
[12] M. Ulina, R. Purba, and A. Halim, "Foreign Exchange Prediction using CEEMDAN and Improved FA-LSTM," in 2020 Fifth International Conference on Informatics and Computing (ICIC), IEEE, Nov. 2020, pp. 1–6. doi: 10.1109/ICIC50835.2020.9288615.
[13] A. J. Dautel, W. K. Härdle, S. Lessmann, and H.-V. Seow, "Forex exchange rate forecasting using deep recurrent neural networks," Digit Finance, vol. 2, no. 1–2, pp. 69–96, Sep. 2020, doi: 10.1007/s42521-020-00019-x.
[14] M. Ayitey Junior, P. Appiahene, and O. Appiah, "Forex market forecasting with two-layer stacked Long Short-Term Memory neural network (LSTM) and correlation analysis," Journal of Electrical Systems and Information Technology, vol. 9, no. 1, p. 14, Dec. 2022, doi: 10.1186/s43067-022-00054-1.
[15] M. M. Panda, S. N. Panda, and P. K. Pattnaik, "Forecasting Foreign Currency Exchange Rate using Convolutional Neural Network," International Journal of Advanced Computer Science and Applications, vol. 13, no. 2, 2022, doi: 10.14569/IJACSA.2022.0130272.
[16] C. Caroline, R. Purba, and M. F. Pasha, "Stock Price Prediction using Long-Short Term Memory and Temporal Convolutional Network," in 2023 Fifth International Conference on Informatics and Computing (ICIC), Malang: IEEE, Dec. 2023.
[17] T. Mahmud, T. Akter, S. Anwar, M. T. Aziz, M. S. Hossain, and K. Andersson, "Predictive Modeling in Forex Trading: A Time Series Analysis Approach," in 2024 Second International Conference on Inventive Computing and Informatics (ICICI), IEEE, Jun. 2024, pp. 390–397. doi: 10.1109/ICICI62254.2024.00070.
[18] P. I. Zitis, S. M. Potirakis, and A. Alexandridis, "Forecasting Forex Market Volatility Using Deep Learning Models and Complexity Measures," Journal of Risk and Financial Management, vol. 17, no. 12, p. 557, Dec. 2024, doi: 10.3390/jrfm17120557.
[19] A. Saghafi, M. Bagherian, and F. Shokoohi, "Forecasting Forex EUR/USD Closing Prices Using a Dual-Input Deep Learning Model with Technical and Fundamental Indicators," Mathematics, vol. 13, no. 9, p. 1472, Apr. 2025, doi: 10.3390/math13091472.
[20] D. M. Ahmed, M. M. Hassan, and R. J. Mstafa, "A Review on Deep Sequential Models for Forecasting Time Series Data," Applied Computational Intelligence and Soft Computing, vol. 2022, pp. 1–19, Jun. 2022, doi: 10.1155/2022/6596397.
[21] M. Mohammadi, S. Jamshidi, A. Rezvanian, M. Gheisari, and A. Kumar, "Advanced fusion of MTM-LSTM and MLP models for time series forecasting: An application for forecasting the solar radiation," Measurement: Sensors, vol. 33, p. 101179, Jun. 2024, doi: 10.1016/j.measen.2024.101179.

[22] J. Bi, X. Zhang, H. Yuan, J. Zhang, and M. Zhou, “A Hybrid Prediction Method for Realistic Network Traffic With Temporal Convolutional Network and LSTM,” IEEE Transactions on Automation Science and Engineering, vol. 19, no. 3, pp. 1869–1879, Jul. 2022, doi: 10.1109/TASE.2021.3077537.
[23] J. Fan, K. Zhang, Y. Huang, Y. Zhu, and B. Chen, “Parallel spatio-temporal attention-based TCN for multivariate time series prediction,” Neural Comput Appl, vol. 35, no. 18, pp. 13109–13118, Jun. 2023, doi: 10.1007/s00521-021-05958-z.
[24] H. Xue, X. Gui, G. Wang, X. Yang, H. Gong, and F. Du, “Prediction of gas drainage changes from nitrogen replacement: A study of a TCN deep learning model with integrated attention mechanism,” Fuel, vol. 357, p. 129797, Feb. 2024, doi: 10.1016/j.fuel.2023.129797.
[25] C.Y. Hsu, Y.W. Lu, and J.H. Yan, “Temporal Convolution-Based Long-Short Term Memory Network With Attention Mechanism for Remaining Useful Life Prediction,” IEEE Transactions on Semiconductor Manufacturing, vol. 35, no. 2, pp. 220–228, May 2022, doi: 10.1109/TSM.2022.3164578.
[26] X. Wei and Y. Xu, “Research on carbon emission prediction and economic policy based on TCN-LSTM combined with attention mechanism,” Front Ecol Evol, vol. 11, Oct. 2023, doi: 10.3389/fevo.2023.1270248.
[27] L. Heng, C. Hao, and L. C. Nan, “Load forecasting method based on CEEMDAN and TCN-LSTM,” PLoS One, vol. 19, no. 7, p. e0300496, Jul. 2024, doi: 10.1371/journal.pone.0300496.
[28] M. Diqi, “StockTM: Accurate Stock Price Prediction Model Using LSTM,” International Journal of Informatics and Computation, vol. 4, no. 1, p. 1, Oct. 2022, doi: 10.35842/ijicom.v4i1.50.
[29] S. Bai, J. Z. Kolter, and V. Koltun, “An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling,” Mar. 2018.
[30] P. Janssen, “Attention based Temporal Convolutional Network for Stock Price Prediction,” Utrecht University, 2022.

BIBLIOGRAPHY OF AUTHORS

Caroline is a dedicated Lecturer in the Information Systems major at Universitas Mikroskil. She earned her Master of Computer Science (M.Kom.) degree in Information Technology and her Bachelor's degree (S.Kom.) from Universitas Mikroskil. Her primary research focus lies in Information System Development and the application of deep learning methodologies for complex prediction and classification tasks.

Wulan Sri Lestari, S.Kom., M.Kom., is a lecturer in the Information Technology Study Program, Faculty of Informatics, Universitas Mikroskil, Medan. She earned her Bachelor's degree in Informatics Engineering and her Master's degree in Information Technology from STMIK Mikroskil Medan. She has authored several books, including Introduction to Databases and Fundamentals of Programming. In addition to her publications, her research interests encompass Cryptography, Data Mining, and Artificial Intelligence.

Mustika Ulina, lecturer in the Information Technology Study Program, Faculty of Informatics, Mikroskil University, Medan, has been specializing in artificial intelligence research for the past four years. The author's research focuses on areas such as predictive modeling, cybersecurity, and machine learning optimization.