Forecasting the Highest and Lowest Prices in Financial Markets Using a VMD-LSTM Hybrid Model

I Made Adi Purwantara*[1], Kusrini[2], Arief Setyanto[3], Ema Utami[4]
[1,2,3,4]Department of Informatics Doctorate, Universitas Amikom Yogyakarta, Indonesia
e-mail: ***[1]adipurwa@students.amikom.ac.id**, [2]kusrini@amikom.ac.id, [3]arief@amikom.ac.id, [4]ema.u@amikom.ac.id

Abstract

Accurate forecasting of the lowest and highest prices in financial markets poses a considerable challenge due to the inherent nonlinear behaviour, non-stationarity, and high noise levels of financial time series data. Most prior studies focus only on closing prices, with limited attention to the simultaneous prediction of high and low prices. Yet, predicting the lowest and highest prices is essential for investors to make informed trading decisions. To address this gap, this study proposes a hybrid DL framework that integrates VMD and LSTM networks for predicting daily high and low prices simultaneously. This study used 12 years of daily data from three diverse assets: AUD/USD, TLKM, and XAU/USD. The data underwent preprocessing, VMD-based decomposition, and were input into the LSTM. The dataset was split 80% for training and 20% for testing. Experiments varied the number of decomposition modes (K = 7, 10, 12) and sliding window sizes (5, 15, 30, 45, 60, 90). Results show that the VMD-LSTM model exhibits improved performance in most of the tested scenarios compared to traditional LSTM. These findings underscore that the use of VMD decomposition can help enhance the accuracy of forecasting the highest and lowest prices in the financial market.

Keywords— VMD, LSTM, Decomposition, Deep Learning, Financial Market

1. INTRODUCTION

Recent advancements in machine learning (ML) and deep learning (DL) have significantly improved the forecasting accuracy of financial time series. Studies consistently show that ML and DL approaches outperform traditional statistical models—such as ARMA, ARIMA, VAR, and GARCH—in predicting market behavior [1], [2], [3]. These traditional models are often constrained by rigid statistical assumptions, making them less effective in capturing the nonlinear and multiscale patterns inherent in real-world financial data [2]. In contrast, artificial intelligence-based models, particularly DL methods, overcome these limitations by adapting to nonlinearity and non-stationarity in financial signals [4]. The superior performance of ML over classical models like the Random Walk and ARIMA has established it as a powerful tool for financial forecasting [5]. Among DL architectures, Long Short-Term Memory (LSTM) networks are widely recognized for their ability to model complex dynamics and long-range dependencies in sequential [6], and have become increasingly popular in stock price prediction tasks [7].

While many statistical and intelligent models have been developed, relying solely on a single model often proves insufficient in effectively denoising financial data [8]. The introduction of neural network-based architectures such as LSTM has greatly enhanced model accuracy [9], leveraging past market activity to make informed predictions [10]. The LSTM model demonstrated superior performance compared to the SVR model in prediction accuracy [11]] and remains one of the most effective DL methods for learning from sequential time series data [12], particularly in the financial domain [13], [14], [15].

Nevertheless, predicting financial market prices remains a major challenge due to the strong presence of nonlinearity, non-stationarity, and noise [8], [16], [17], [18], [19]. This highlights the need for a robust denoising technique during preprocessing before applying DL models [12]. Several signal decomposition methods have been used, including Empirical Mode Decomposition (EMD), wavelet transform, and VMD [20]. While wavelet transform requires prior identification of signal components—often resulting in inaccurate frequency analysis—adaptive methods like EMD and VMD allow data-driven decomposition based on local signal characteristics [21]. However, EMD suffers from issues such as mode mixing and lacks solid theoretical foundations [22]. To address these limitations, Dragomiretskiy and Zosso (2014) introduced VMD as a theoretically grounded and non-recursive decomposition technique with strong robustness to noise and sampling [23].

VMD has demonstrated strong performance in enhancing feature extraction and reducing noise in non-stationary time series, including financial market data. It decomposes complex sequences into simpler, more stationary Intrinsic Mode Functions (IMFs), making meaningful patterns more accessible to predictive models [1], [12], [24]. Recent studies have begun to explore hybrid forecasting frameworks that combine signal decomposition methods like VMD with DL models to enhance prediction accuracy [4]. The integration of signal decomposition and AI models significantly impacts forecasting performance [2], [4], [24]. For example, the CNN-TLSTM hybrid model successfully predicted the next-day closing price of the USD/CNY exchange rate [3]. Likewise, a hybrid LSTM-DNN model produced consistent and accurate stock price predictions across multiple datasets [10]. More refined techniques, such as GA-VMD, further optimize the number of decomposed components using genetic algorithms, offering a deeper understanding of financial signals [2], [4].

Most of the existing studies have focused on predicting opening and closing prices [2], [3], [5], [10], [16], [24], [25], [26], [27]. However, [1], [28] argue that the highest and lowest prices provide a more comprehensive representation of market characteristics and risk levels compared to closing prices. Forecasting these price extremes is crucial for effective risk management and offers critical insights to both market participants and regulators [1].

Despite their importance, few studies address the simultaneous prediction of both highest and lowest prices in financial time series. This paper fills that gap by proposing a hybrid deep learning model—VMD-LSTM—designed to forecast both the daily high and low prices simultaneously. VMD is used to decompose the raw price series into a set of IMFs, reducing noise and isolating multiscale patterns. These refined components are then used as inputs for the LSTM network, which captures long-term dependencies in the data. This dual-stage design leverages the strengths of both VMD and LSTM to provide more accurate, reliable forecasts.

This study makes four key contributions: (1) it presents a unified VMD-LSTM model for simultaneous prediction of high and low prices, offering richer insights than single-target models; (2) it leverages VMD's denoising and multiscale decomposition capabilities to improve input quality; (3) it applies LSTM's sequential learning capabilities to capture temporal dynamics; and (4) it validates the model using three diverse datasets—AUD/USD (currency), TLKM (stock), and XAU/USD (commodity)—demonstrating the model's adaptability and robustness across markets. Practically, this approach offers substantial value to stakeholders. For traders and investors, accurate high and low price forecasts improve entry and exit strategies. For regulators and policymakers, they offer better visibility into market volatility and systemic risk [1], [28].

2. RESEARCH METHODS

2.1. *Data Collection*

To ensure the robustness and real-world relevance of the developed forecasting model, this study uses historical time series data from three different financial asset classes, each representing unique market characteristics. The assets were carefully selected to encompass

variations in volatility levels, market participant types, and trading volume. This strategy aims to test the model's generalizability across a range of complex and dynamic market conditions.
The three assets used in this study are:
1. AUD/USD, a currency pair from the foreign exchange (forex) market,
2. TLKM, a telecommunications company stock from the Indonesian domestic stock market,
3. XAU/USD, a gold commodity traded in US dollars on the global market.

These three assets were selected based on three main criteria:
1. High volatility, to evaluate the model's resilience to extreme price fluctuations,
2. Availability of complete and continuous daily historical data,
3. Representation of different types of financial markets (forex, stock, and commodity markets) to allow for testing the model's performance in different economic contexts.

All datasets were retrieved from Investing.com, a widely used and reliable financial data provider. The data span a comprehensive 12 years, from early 2013 to the end of 2024. By leveraging real and high-resolution daily data over a long horizon, this study ensures both temporal depth and diversity in the input sequences used for model development and testing.

2.2. Data Pre-Processing

To prepare the datasets for effective modeling, a structured data pre-processing workflow was implemented.
1. Missing Value Handling
   Each dataset was first cleaned to handle missing values and ensure time alignment across the series.
2. Outlier Detection and Adjustment
   Outliers were identified and smoothed where necessary to avoid distortion in model training.
3. Data Normalization
   All data were scaled using Min-Max normalization to transform the price values into a standardized range between 0 and 1.
4. Data Splitting
   An 80:20 ratio was employed to divide the normalized dataset into training and testing segments. The initial 80% was utilised for training purposes, while the remaining 20% was set aside to assess the model's performance. This temporal split guarantees that the model is evaluated on unseen data, reflecting real-world forecasting scenarios where future prices must be predicted based solely on past observations.
5. Sliding Window Formation
   To generate time series inputs that can be utilized by the model, a sliding window approach is used. Several window lengths are used, including 5, 15, 30, 45, 60, and 90 days, to capture both short-term and long-term price patterns. This technique also enriches the feature representation provided to the model, thereby improving accuracy and generalization capabilities.

2.3. Varational Mode Decomposition (VMD)

VMD algorithm proposed by [23]. VMD was developed as a non-recursive and theoretically grounded approach that decomposes signals into adaptive modes simultaneously, offering greater robustness to noise and sampling compared to previous methods [23].

Let the input signal be $f(t)$, and assume we want to decompose it into K modes $\{u_k(t)\}_{k=1}^{K}$ with corresponding center frequencies$\{\omega_k\}_{k=1}^{K}$. Each mode is expected to be band-limited after demodulation to baseband.

$$\min_{\{u_k\},\{\omega_k\}} \left\{ \sum_{k=1}^{K} | \partial_t \left(u_k(t) \cdot e^{-j\omega_k t} \right) |_2^2 \right\} \tag{1}$$

Subject to $\sum^{K}{}_{k=1} u_k(t) = f(t)$ (2)

This formulation minimizes the total bandwidth of all modes by penalizing the gradient of each demodulated mode in time (which corresponds to concentration in frequency). The constraint ensures that the sum of the modes reconstructs the original signal.

To solve this constrained optimization, an augmented Lagrangian approach is applied:

$$\mathcal{L}(\{u_k\},\{\omega_k\},\lambda) = \alpha \sum_{k=1}^{K} | \partial_t \backslash big\left(u_k(t) \cdot e^{-j\omega_k t} \backslash big \right) |_2^2 + |f(t) - \sum_{k=1}^{K} u_k(t) |_2^2 + \left\langle \lambda(t), f(t) - \sum_{k=1}^{K} u_k(t) \right. \tag{3}$$

Where:

- α is a penalty parameter controlling the bandwidth constraint,
- $\lambda(t)$ is the Lagrange multiplier

In this study, VMD was used to decompose the daily high and low price time series into IMFs. These IMFs were then used as multivariate inputs for an LSTM model. The VMD parameters used in this experiment were as follows:

1. Number of modes (K): 7, 10, and 12
2. Convergence tolerance (ε): 1e-7
   This value is used as the stopping criterion in the VMD iterative process.
3. Alpha penalty (α): 2000
   This is the default value used to maintain consistency between modes.
4. Initialized center frequencies (Init): 1
   Represent the starting frequency estimates around which each mode is formed.

The choice of VMD over other decomposition methods such as Wavelet and EMD is based on the following advantages:

1. Robustness to noise
   VMD exhibits stable denoising performance even on highly volatile financial signals.
2. Avoids mode mixing
   Unlike EMD, which often produces mixed modes, VMD explicitly limits the bandwidth of each mode [22].
3. Robust mathematical formulation
   VMD is built on variational theory, making it more controllable and analytically tractable [23].
4. Compatibility with Deep Learning
   VMD decomposition results are well-suited for use as input features in DL architectures such as LSTMs, as they simplify temporal complexity and improve prediction accuracy [24].

2.4. LSTM (Long Short-Term Memory)

LSTM is a specialized type of RNN intended to learn long-term dependencies and address the limitations of traditional RNNs, especially issues associated with vanishing and exploding gradients [16]. LSTM is a widely used and highly effective DL approach within RNNs for time series and sequence prediction tasks [15], [16], [29]. LSTM introduces a memory cell and a series of gates that control the flow of information, enabling it to retain relevant information over extended time steps. The architecture of LSTM is shown in Figure 1.

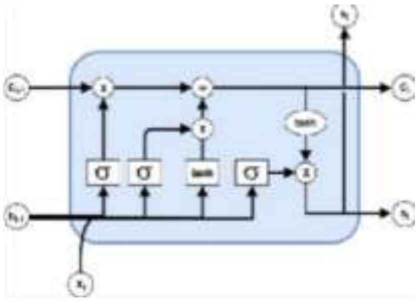
Figure 1. LSTM Architecture

Each LSTM cell contains three core gates: input, forget, and output gates [16], which manage the retention, update, and transmission of information across time steps. Let x_t be the input vector at time step t, C_{t-1} the previous cell state, and h_{t-1} the hidden state from the previous time step. The computations of an LSTM cell are defined as follows:

Forget Gate

Determines which information to remove from cell state:

$$f_t = \sigma\left(W_f \cdot [h_{\{t-1\}}, x_t] + b_f\right) \quad (4)$$

Input Gate

Updates the cell state with new candidate values:

$$i_t = \sigma\left(W_i \cdot [h_{\{t-1\}}, x_t] + b_i\right) \quad (5)$$

$$\tilde{C}_t = \tanh\left(W_C \cdot [h_{t-1}, x_t] + b_C\right) \quad (6)$$

Cell State Update

Combines the previous cell state and new information:

$$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t \quad (7)$$

Output Gate

Determines the next hidden state:

$$o_t = \sigma\left(W_o \cdot [h_{t-1}, x_t] + b_o\right) \quad (8)$$

$$h_t = o_t * \tanh\left(C_t\right) \quad (9)$$

In this study, the LSTM model is applied to learn temporal patterns from historical financial time series data to predict the next-day lowest and highest prices of a financial market. The sequential nature of LSTM enables the model to capture both short-term and long-term dependencies in the price fluctuations, providing a robust framework for accurate forecasting in volatile financial environments.

2.5. The Proposed Model (VMD-LSTM)

This study proposes a hybrid DL framework combining VMD and LSTM networks—referred to as VMD-LSTM—to simultaneously predict daily lowest and highest prices in financial markets. This approach is motivated by the need to enhance the model's ability to capture both the nonlinear and non-stationary characteristics of financial time series, which are often difficult to model using conventional methods. VMD functions as an effective signal decomposition technique that breaks down a time series into a collection of IMFs, each capturing oscillatory behavior within distinct frequency ranges. These IMFs reflect different underlying components in the original data, such as trend, noise, or short-term fluctuations. By decomposing the input data before training the model, VMD helps isolate meaningful features and reduce noise, allowing the LSTM network to focus on relevant temporal patterns. LSTM networks are then employed to learn the sequential dependencies within the decomposed signals. Their internal memory mechanism allows them to retain valuable historical information, making them particularly effective for time series prediction. This hybrid framework leverages the strengths of both VMD and LSTM: VMD enhances signal clarity, while LSTM captures long-term temporal dependencies.

The proposed model architecture, as shown in Figure 2, follows these sequential steps:

1. Input Data
   The process begins with financial time series data of daily highest and lowest prices from the AUD/USD, TLKM, and XAU/USD datasets.
2. VMD Decomposition
   Using VMD, the input sequence is separated into multiple IMFs, each corresponding to a unique frequency element of the initial time series.
3. Concatenation of IMFs
   The resulting IMFs are concatenated in a predefined order (e.g., from highest to lowest frequency), forming a new input sequence enriched with decomposed features.
4. LSTM Layer
   The concatenated sequence is fed into the LSTM layer, which learns the temporal patterns and dependencies within the multi-resolution data.
5. Dense Layer
   The LSTM layer's output is subsequently fed into one or more fully connected (dense) layers, which transform the learned features into the target prediction domain.
6. Output Layer
   The final prediction consists of the next-day highest and lowest prices, estimated from the processed temporal signals.

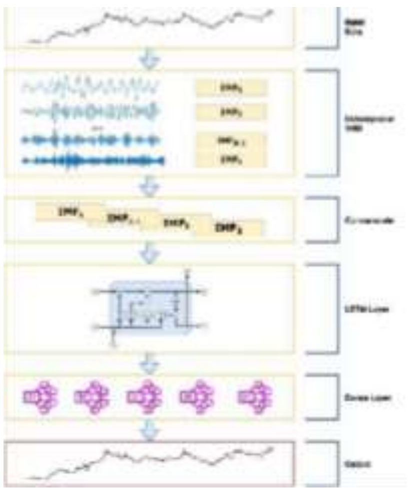

Figure 2. VMD-LSTM Model Architecture

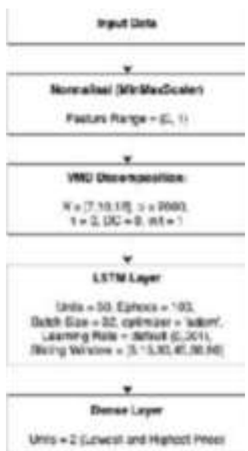

Figure 3. VMD-LSTM Model Technical Configuration

To provide a comprehensive overview of the proposed prediction system, Figure 3 presents the block diagram of the technical configuration of the VMD-LSTM model. This diagram outlines the end-to-end process, including preprocessing, signal decomposition via VMD, sequence construction using sliding windows, model architecture, and training parameters. This hybrid VMD-LSTM model is designed to improve forecasting accuracy by integrating signal decomposition and deep sequential learning in a unified framework. The layered architecture ensures that both low-level fluctuations and high-level trends are captured effectively, enabling the model to adapt to the complex nature of financial markets.

The LSTM model was constructed using a sequential architecture, starting with an input layer derived from the VMD decomposition signal. All IMFs of the highest and lowest prices are arranged and consolidated into a multivariate input. The model architecture consists of one LSTM layer with 50 hidden units, which is responsible for capturing temporal dependencies in financial time series data. After the LSTM, a Dense output layer with two units is used to simultaneously predict the highest and lowest prices. This model was compiled using the Adam optimizer, a widely used adaptive gradient-based optimization algorithm known for its efficiency in training deep learning models. The training process was conducted over 100 epochs with a batch size of

32.

2.6. Evaluation Metrics

To evaluate the predictive performance of the proposed model in forecasting the highest and lowest prices in financial markets, several widely accepted statistical metrics are employed. In this study, the performance evaluation is conducted using MSE, RMSE, MAE, MAPE, and R^2. Each metric captures different aspects of prediction quality, and using a combination of them ensures a comprehensive assessment of the model's forecasting capability.

1. MSE (Mean Squared Error)
   By averaging the squared differences between predictions and actual outcomes [29], MSE places greater emphasis on larger errors, thereby increasing its sensitivity to outliers.

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \quad (10)$$

2. RMSE (Root Mean Squared Error)
   As the square root of MSE, RMSE expresses error in the predicted variable's unit, allowing for clearer interpretation [16].

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \quad (11)$$

3. MAE (Mean Absolute Error)
   MAE measures the mean of the absolute deviations between predicted and actual values [29]. In contrast to MSE, it assigns equal weight to all errors and demonstrates greater robustness to outliers.

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \quad (12)$$

4. MAPE (Mean Absolute Percentage Error)
   As a percentage-based metric, MAPE allows consistent evaluation of prediction accuracy across datasets with different value ranges [29]

$$\text{MAPE} = \frac{100\%}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \quad (13)$$

5. R^2 (Coefficient of Determination)
   The R^2 value spans from 0 to 1, where 0 indicates that the model fails to explain any variance in the target variable and 1 indicates perfect prediction [29].

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \quad (14)$$

These metrics collectively offer insights into the model's effectiveness regarding accuracy, error magnitude, relative error, and explanatory strength. By analyzing these indicators, the reliability and generalization ability of the proposed VMD-LSTM model for forecasting financial market prices can be effectively validated. To determine the effect of applying the VMD decomposition method, the performance of the proposed model was compared with that of a baseline model (a single LSTM) across various window size parameters and *K* values. To determine whether the effect was significant, a statistical test was performed using a paired t-test between the models.

3. RESULT AND DISCUSSION

3.1. Decomposition Results

The selection of the number of modes (*K*) in VMD is a critical factor that directly influences the quality of the decomposition process and the subsequent predictive performance of the model. Since VMD requires the value of *K* to be predefined [23], an inappropriate choice can degrade model effectiveness. A small *K* may lead to under-decomposition, conversely, a value of *K* that is too large can cause over-decomposition [2]. To address this, Guo et al. (2022) conducted empirical evaluations using *K* values ranging from 7 to 13 and found that models with *K* = 9, 10, and 12 yielded superior forecasting accuracy compared to other configurations [22]. Informed by these findings, this study explores *K* values of 7, 10, and 12 to identify effective decomposition settings across multiple financial time series.

Figure 4 presents a comprehensive visualization of the VMD decomposition results for three distinct financial datasets: AUD/USD, TLKM, and XAU/USD, with varying values of the mode parameter *K* set to 7, 10, and 12. Each image illustrates the resulting modes for both the highest and lowest price series, offering a detailed view of how different values of *K* influence the granularity and frequency separation of the extracted components. This visual representation serves as a foundational step for understanding the intrinsic structures within each dataset, enabling a deeper exploration of price dynamics through signal decomposition.

Each subplot illustrates one of the IMFs. This decomposition aims to separate the original signals into distinct frequency bands, enabling a clearer representation of hidden patterns and reducing noise prior to modeling. By isolating these components, the model gains access to more structured and meaningful features that may enhance forecasting performance. By examining these multi-resolution modes, researchers and practitioners are better equipped to design more adaptive and insightful forecasting models that can accommodate the unique characteristics of diverse financial instruments. This structured breakdown helps in uncovering latent temporal dynamics that may be otherwise obscured in the raw data, facilitating a more nuanced analysis of market behavior.

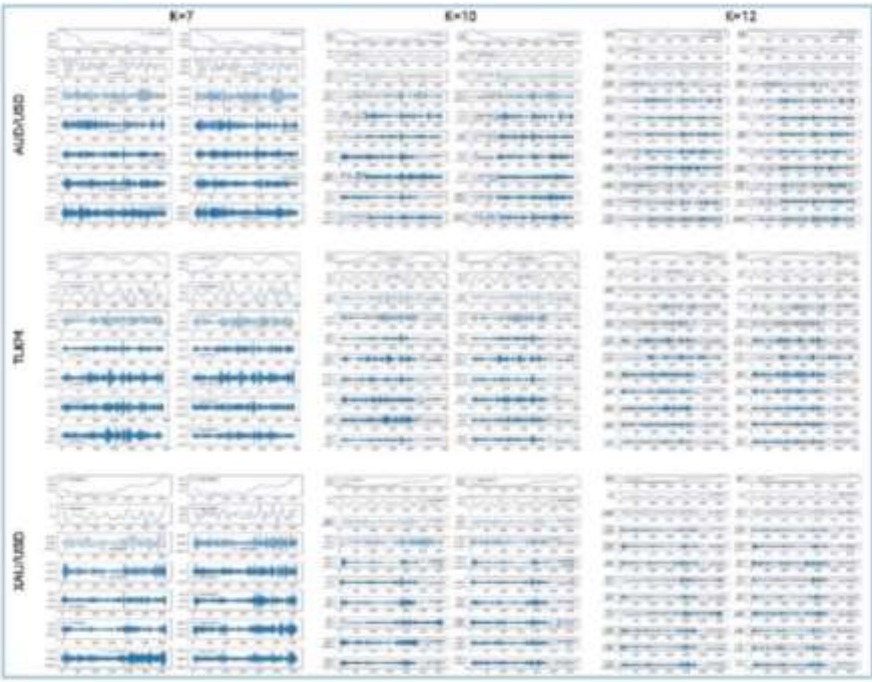

Figure 4. VMD Results on Dataset

Upon closer inspection, increasing the number of modes leads to a finer separation of signal components. For instance, when *K*=7, most of the informative structure appears concentrated in the first few modes, while the remaining modes carry residual noise. As *K* increases to 10 and 12, the later modes become increasingly granular and attenuated, suggesting that the additional components capture less significant, high-frequency fluctuations or noise. This observation underlines the importance of selecting an optimal *K*, where too few modes risk under-representation of complex features, while too many introduce redundancy and potential overfitting.

3.2. Training Performance Evaluation

Table 1 presents an evaluation metric of model performance during the training phase across three financial datasets: AUD/USD, TLKM, and XAU/USD. These evaluations compare the conventional LSTM model with the VMD-LSTM hybrid model across different configurations of mode decomposition (*K*) and window sliding parameters. Each configuration is assessed using several performance metrics: RMSE, MSE, MAPE, MAE, and R^2, offering a multidimensional perspective on prediction accuracy and consistency. By providing a systematic breakdown of these results, the tables serve as a foundational reference for understanding the impact of VMD pre-processing and parameter tuning on model learning behavior across diverse market characteristics. This comparison allows for a deeper understanding of how decomposition and temporal windowing impact model performance.

Table 1. Training Performance Evaluation AUD/USD Dataset

Dataset				AUD/USD	TLKM	XAU/USD
No	Model	K	Sliding Window	RMSE	RMSE	RMSE
1	LSTM		5	0,004442515	53,71704294	12,19181067
2	LSTM		15	0,004031258	54,15265478	11,97839877
3	LSTM		30	0,004208575	54,60715476	11,69768836
4	LSTM		45	0,004177324	55,41750242	12,13575095
5	LSTM		60	0,00398611	53,99657907	12,70419147
6	LSTM		90	0,004096135	57,45842819	11,28761051
7	VMD-LSTM	7	5	0,002409807	25,22548282	6,378307003
8	VMD-LSTM	7	15	0,002327581	28,41748175	6,306351768
9	VMD-LSTM	7	30	0,002360489	26,43825719	6,948995959
10	VMD-LSTM	7	45	0,002125096	28,35296239	6,633423255
11	VMD-LSTM	7	60	0,002261227	27,9316685	6,193917283
12	VMD-LSTM	7	90	0,002297384	28,77239484	6,154127654
13	VMD-LSTM	10	5	0,001965094	21,78152356	5,697663474
14	VMD-LSTM	10	15	0,002372911	21,67915819	6,399940615
15	VMD-LSTM	10	30	0,001861668	19,73209835	5,568940101
16	VMD-LSTM	10	45	0,002362396	20,46459545	4,99930748
17	VMD-LSTM	10	60	0,001904868	21,70196694	5,085203781
18	VMD-LSTM	10	90	0,001825123	23,24276674	6,579942954
19	VMD-LSTM	12	5	0,002322122	19,09675989	5,94250028
20	VMD-LSTM	12	15	**0,00173972**	17,04845191	5,008176365
21	VMD-LSTM	12	30	0,001815474	17,99433365	4,823725969
22	VMD-LSTM	12	45	0,001745456	17,29201142	5,198055305
23	VMD-LSTM	12	60	0,001775374	**14,76778659**	5,03297987
24	VMD-LSTM	12	90	0,001785629	15,76554868	**4,722103917**

The comparative evaluation of the LSTM and VMD-LSTM models across the AUD/USD, TLKM, and XAU/USD datasets suggests a general trend in which the integration of VMD tends to improve predictive accuracy. In the AUD/USD dataset, the VMD-LSTM model with *K*=12 and a sliding window of 15 consistently achieved the lowest error rates, MSE (3,02662E-06), RMSE (0,00173972), MAE (0,001341893), and MAPE (0,172350141), along with the highest R-squared (0,99962635). Similarly, for the TLKM dataset, the model with *K*=12 and window size 60 produces the lowest MSE (218,0875209), RMSE (14,76778659), MAE (11,34078889), and MAPE (0,344460634), with a strong R^2 value of 0,999568616. Meanwhile, for the XAU/USD dataset, the model with *K*=12 and window size 90 produces the lowest MSE (22,29826541), RMSE (4,722103917), MAE (3,55485696), and MAPE (0,247810299), with a strong R^2 value of 0,999659819. These findings illustrate the effectiveness of VMD in decomposing complex and erratic patterns into more interpretable subcomponents, which enhances the LSTM's ability to generalize across highly volatile time series. Notably, the results also suggest that during model training, the VMD-LSTM configuration with *K*=12 consistently yields the best predictive performance among the tested settings.

3.3. Testing Performance Evaluation

Table 2 presents an evaluation metric of model performance during the testing phase for three different financial datasets: AUD/USD, TLKM, and XAU/USD. This table systematically summarizes the performance metrics of two predictive models—standard LSTM and the hybrid VMD-LSTM—across various values of *K* (number of decomposed modes) and window sizes. For each configuration, performance is evaluated using five key indicators: RMSE, MSE, MAPE, MAE, and R^2. This structured presentation allows for a clear and consistent comparison of how well each model captures the patterns within different market instruments under varying experimental setups, setting the stage for a deeper analysis of predictive accuracy and generalization capability.

The evaluation results across the three datasets—AUD/USD, TLKM, and XAU/USD—indicate that the VMD-LSTM hybrid model generally performs better than the conventional LSTM model under the tested conditions. For the AUD/USD and TLKM datasets, the best performing configurations are VMD-LSTM with *K*=12 and a sliding window of 60. For the AUD/USD dataset, the VMD-LSTM with *K*=12 and a sliding window of 60 achieves the lowest MSE (2,07056E-06), RMSE (0,001438943), MAE (0,001097814), and MAPE (0,165546573), along with the highest R-squared (0,99313337). Similarly, for the TLKM dataset, the model with *K*=12 and window size 60 produces the lowest MSE (183,3456417), RMSE (13,54051852), MAE (10,64432145), and MAPE (0,29567552), with R^2 value of 0,999360302. Meanwhile, in the XAU/USD dataset, the configuration with *K*=12 and window size 45 results in the lowest MSE (144,4699331), RMSE (12,0195646), MAE (9,326852137), and MAPE (0,42259835), with the highest R-squared of 0,998409082.

Table 2. Testing Performance Evaluation AUD/USD Dataset

Dataset				AUD/USD	TLKM	XAU/USD
No	Model	K	Sliding Window	R2	R2	R2
1	LSTM		5	0,939114227	0,991540903	0,996798535
2	LSTM		15	0,956511775	0,991475363	0,996829899
3	LSTM		30	0,948308884	0,990575183	0,989006788
4	LSTM		45	0,95039475	0,990113165	0,996481738
5	LSTM		60	0,954434364	0,991498244	0,993272673
6	LSTM		90	0,953293655	0,990628635	0,993566515
7	VMD-LSTM	7	5	0,981232805	0,997820724	0,990755804
8	VMD-LSTM	7	15	0,986102654	0,996644742	0,993941185
9	VMD-LSTM	7	30	0,985834085	0,997940737	0,99817653
10	VMD-LSTM	7	45	0,98628833	0,997064178	0,997109787
11	VMD-LSTM	7	60	0,985942706	0,99739871	0,989493295
12	VMD-LSTM	7	90	0,985152042	0,997715056	0,992981576
13	VMD-LSTM	10	5	0,98548931	0,997630342	0,994046246
14	VMD-LSTM	10	15	0,988471665	0,99813777	0,996070642
15	VMD-LSTM	10	30	0,989985191	0,998825598	0,997370753
16	VMD-LSTM	10	45	0,989317073	0,998728667	0,997693018
17	VMD-LSTM	10	60	0,989078113	0,998272042	0,996357696
18	VMD-LSTM	10	90	0,990840136	0,998454507	0,997159577
19	VMD-LSTM	12	5	0,99274339	0,997574532	0,994603621
20	VMD-LSTM	12	15	0,991651702	0,998714009	0,995828544
21	VMD-LSTM	12	30	0,985636443	0,998718297	0,993816056
22	VMD-LSTM	12	45	0,989352367	0,99922435	**0,998409082**
23	VMD-LSTM	12	60	**0,99313337**	**0,999360302**	0,988004503
24	VMD-LSTM	12	90	0,989701724	0,998766653	0,994646076

.

This suggests not only a potential improvement in predictive precision but also a more stable model fit, highlighting the advantages of decomposing complex price signals into IMFs prior to time-series modeling. In comparison, the standard LSTM models tend to produce lower R^2 scores and higher error values, which may reflect the challenges of modeling raw, unprocessed sequences when capturing detailed price dynamics. These findings are further illustrated in Figure 5, Figure 6, and Figure 7, which clearly demonstrates the superior and more consistent R-squared performance of the VMD-LSTM models across different sliding window sizes.

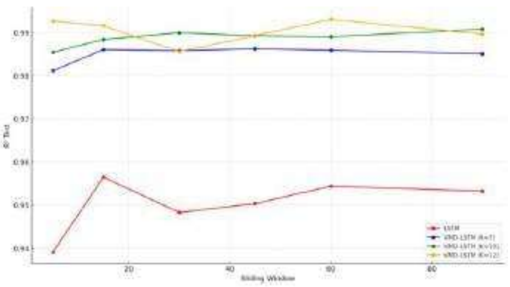

Figure 5. R^2 Performance on the AUD/USD Dataset

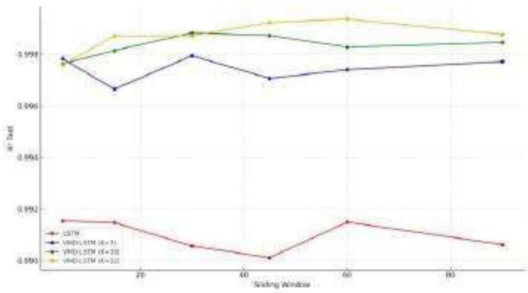

Figure 6. R^2 Performance on the TLKM Dataset

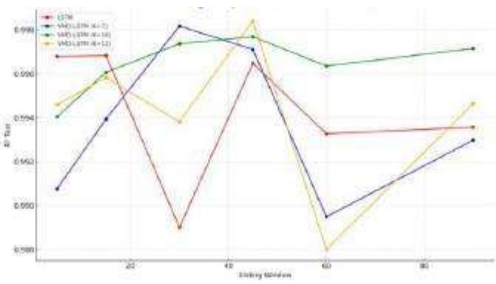

Figure 7. R^2 Performance on the XAU/USD Dataset

These findings provide supportive evidence for the role of VMD in improving LSTM's ability to learn from multi-scale features in financial data. By decomposing the original series into orthogonal sub-signals, VMD may help the LSTM process cleaner and more informative inputs, potentially contributing to reduced prediction errors. Nonetheless, the results also show that model performance is sensitive to the choice of decomposition parameters such as the number of modes (K) and window size. Therefore, while the hybrid approach appears promising, careful tuning of hyperparameters remains essential for achieving optimal performance.

However, the testing results on the XAU/USD dataset indicate that the VMD-LSTM model does not consistently outperform the baseline LSTM, particularly under configurations with shorter sliding windows (5 and 15), where the conventional LSTM exhibits better predictive accuracy. This discrepancy may be due to the non-stationarity of the XAU/USD series, which poses greater challenges for modeling. Nonetheless, the VMD-LSTM model achieves better performance across most other sliding window configurations.

The experimental results indicate that the sliding window size is vital in determining the predictive performance of both LSTM and VMD-LSTM models across different financial time series. On the AUD/USD and TLKM datasets, which exhibit relatively stable and periodic behavior, window sizes of 60 yielded the best performance, particularly for VMD-LSTM with $K = 12$. Conversely, on the XAU/USD dataset, a shorter window (e.g., window = 45) performed better for VMD-LSTM with $K = 12$. This study reaffirms that window size must be aligned with the temporal granularity of the signal, where medium windows are optimal for detecting seasonal patterns, and shorter windows offer better responsiveness in environments with rapid shifts. Thus, sliding window tuning should be considered a critical hyperparameter in the design of hybrid DL forecasting frameworks for financial applications.

To assess whether the integration of VMD into the LSTM model significantly enhances predictive performance, we conducted a paired T-test. This statistical method compares the prediction accuracy of the baseline LSTM model to the VMD-LSTM model across multiple commonly used performance metrics. The comparison was carried out on three different financial time series datasets, namely AUD/USD, TLKM, and XAU/USD. The results, including p-values and interpretations, are presented as Table 3 below.

Table 3. Paired T-Test Results

AUD/USD			TLKM			XAU/USD		
Metric	p-value	Interpretation	Metric	p-value	Interpretation	Metric	p-value	Interpretation
MSE	2,24E-01	Significant	MSE	1,06E-05	Significant	MSE	0.92861	Not Significant
RMSE	2,02E-03	Significant	RMSE	1,70E-01	Significant	RMSE	0.97712	Not Significant
MAE	2,76E-03	Significant	MAE	1,95E-01	Significant	MAE	0.97545	Not Significant
MAPE	2,80E-03	Significant	MAPE	2,38E-01	Significant	MAPE	0.76719	Not Significant
R^2	2,06E-01	Significant	R^2	9,96E-05	Significant	R^2	0.95537	Not Significant

For the AUD/USD and TLKM datasets, across all performance metrics, the p-values are significantly below the 0.05 threshold. This indicates that the VMD-LSTM model for the AUD/USD and TLKM datasets outperforms the standard LSTM model with statistically significant improvements. In contrast, the results from the XAU/USD dataset show a different trend. All p-values are greater than 0.05, indicating that the VMD-LSTM model for the XAU/USD dataset outperforms the standard LSTM model with statistically not significant improvements.

A potential reason for this lack of statistical significance in the XAU/USD dataset lies in the stationarity characteristics of the data. According to our Augmented Dickey-Fuller (ADF) test, the XAU/USD highest price series has an ADF statistic of 1.21 with a p-value of 0.996, indicating strong non-stationarity. Similarly, the lowest price series shows an ADF statistic of 0.61 with a p-value of 0.987. In contrast, the AUD/USD dataset shows borderline stationarity in the highest price series (ADF statistic = -2.81, p-value = 0.057) and weak non-stationarity in the lowest price series (ADF statistic = -2.72, p-value = 0.070). The TLKM dataset also exhibits non-stationarity, with ADF statistics of -2.46 (p = 0.124) for the highest prices and -2.42 (p = 0.137) for the lowest prices.

3.4. Discussion

The experimental results suggest that the VMD-LSTM hybrid model tends to provide improved predictive performance compared to the traditional LSTM in most tested scenarios. The integration of VMD effectively decomposes complex financial time series into smoother IMFs, which help the LSTM model focus on more informative temporal features. This finding supports the hypothesis that preprocessing with VMD reduces noise and highlights essential patterns, ultimately resulting in more accurate and reliable predictions [8], [20], [21], [22], [24].

Support for the reliability of the proposed model is indicated by its relatively consistent performance across three distinct asset classes: currency (AUD/USD), domestic stocks (TLKM), and commodities (XAU/USD). These results suggest that the VMD-LSTM model shows potential to perform well across different types of financial data, indicating a degree of adaptability to varying volatility patterns and structural characteristics in financial markets. These findings are in agreement with earlier studies by [24].

Unlike most prior studies that concentrate solely on predicting closing prices, this study focuses exclusively on the prediction of highest and lowest prices, which represent the daily price range and are crucial for risk management, volatility analysis, and intraday trading strategies. Previous works such as [24] and [27] demonstrated improvements in closing price prediction using hybrid VMD-ICSS-BiGRU and 1D-CapsNet-LSTM models, yet none of them explicitly addressed or evaluated the predictive modeling of high and low price levels. This research fills that gap by presenting a DL framework that is specifically optimized for predicting price extremes rather than closing trends. By focusing on the lowest and highest prices, the proposed model contributes new insights into market behavior that are frequently neglected in conventional forecasting approaches.

Research conducted by [1] proposed a combined DL model to predict the highest and lowest prices of the S&P 500 Index and the SSEC Index, using a dataset spanning approximately 12 years. The performance of the proposed model in that study yielded the lowest MAPE of 0.3105, while in this study, the lowest MAPE for a representative stock market (TLKM) was 0.2957.

Nevertheless, there are several limitations worth noting. First, the determination of the optimal number of VMD modes (K) in this study was based on experiments to get the best value. Developing a more systematic and adaptive method for selecting K remains an important direction for future research. Second, while LSTM networks are powerful in modeling sequential dependencies, they often struggle to capture fine-grained local patterns and are susceptible to noise—especially when dealing with non-stationary and highly volatile financial signals. To address this limitation, integrating CNNs as a complementary component presents a promising avenue for enhancing the model's ability to extract localized temporal features more effectively.

3.5. Future Research Recommendation

Building on the promising results of the VMD-LSTM hybrid approach, future research is encouraged to extend this architecture by integrating additional DL components to further enhance predictive performance and model robustness. Although VMD-LSTM effectively captures temporal dependencies in decomposed financial time-series data, it presents several limitations. LSTM models, while powerful for sequential learning, often struggle to extract fine-grained local patterns and are susceptible to noise, especially when handling non-stationary and volatile financial signals. Furthermore, LSTMs may face computational challenges and gradient degradation over long sequences. To address these limitations, incorporating CNNs as a complementary component is a promising direction. CNNs are well-suited for capturing local temporal-spatial features and filtering noise from VMD-derived sub-signals through their hierarchical feature extraction capability. Therefore, adding a CNN model to the proposed model can provide a more comprehensive framework, where VMD handles signal decomposition, CNN captures local patterns, and LSTM models sequential dependencies across time.

Additionally, future studies may also explore the integration of Transformer-based attention mechanisms to capture long-range dependencies and dynamic temporal interactions beyond what LSTM can achieve. Researchers are also encouraged to investigate adaptive decomposition strategies, in which VMD parameters such as the number of modes (K) are automatically optimized based on the data characteristics, potentially via metaheuristic or learning-based optimization. Finally, it is essential to validate the generalizability of the proposed model across a wider range of financial instruments, including cryptocurrencies, commodities, and sector-specific indices.

4. CONCLUSION

This study has developed a hybrid model combining VMD and LSTM to forecast the lowest and highest prices in the financial markets. VMD is used to decompose complex price signals into IMFs before feeding them into the LSTM, while LSTM learns patterns from these signals. Across three distinct datasets—AUD/USD, TLKM, and XAU/USD—the VMD-LSTM model exhibits improved performance in most of the tested scenarios compared to traditional LSTM, although its effectiveness may vary depending on the dataset's characteristics. These findings underscore that the use of VMD decomposition can help enhance the accuracy of forecasting the highest and lowest prices in the financial market.

However, there are several limitations worth noting. The number of modes (K) in VMD was selected empirically, rather than using automated methods. Furthermore, the model is not always consistent, especially when applied to highly non-stationary data, such as XAU/USD.

Because this study only covered three asset classes, the results cannot be generalized to all types of financial instruments, such as cryptocurrencies or other assets.

In summary, the VMD-LSTM model has potential, but requires more extensive testing. Future research could attempt to optimize parameters automatically, or add other components, such as CNNs or attention mechanisms, to improve performance in more dynamic market conditions.

REFERENCES

[1] H. Gong and H. Xing, "Predicting the highest and lowest stock price indices: A combined BiLSTM-SAM-TCN deep learning model based on re-decomposition," *Appl Soft Comput*, vol. 167, Dec. 2024, doi: 10.1016/j.asoc.2024.112393.

[2] Y. Song, J. Huang, Y. Xu, J. Ruan, and M. Zhu, "Multi-decomposition in deep learning models for futures price prediction," *Expert Syst Appl*, vol. 246, Jul. 2024, doi: 10.1016/j.eswa.2024.123171.

[3] J. Wang, X. Wang, J. Li, and H. Wang, "A Prediction Model of CNN-TLSTM for USD/CNY Exchange Rate Prediction," *IEEE Access*, vol. 9, pp. 73346–73354, 2021, doi: 10.1109/ACCESS.2021.3080459.

[4] S. Li, G. Tang, X. Chen, and T. Lin, "Stock Index Forecasting Using a Novel Integrated Model Based on CEEMDAN and TCN-GRU-CBAM," *IEEE Access*, vol. 12, pp. 122524–122543, 2024, doi: 10.1109/ACCESS.2024.3452426.

[5] Y. Li, L. Chen, C. Sun, G. Liu, C. Chen, and Y. Zhang, "Accurate Stock Price Forecasting Based on Deep Learning and Hierarchical Frequency Decomposition," *IEEE Access*, vol. 12, pp. 49878–49894, 2024, doi: 10.1109/ACCESS.2024.3384430.

[6] A. Mardjo and C. Choksuchat, "HyBiLSTM: Multivariate Bitcoin Price Forecasting Using Hybrid Time-Series Models with Bidirectional LSTM," *IEEE Access*, vol. 12, pp. 50792–50808, 2024, doi: 10.1109/ACCESS.2024.3386029.

[7] Z. Nourbakhsh and N. Habibi, "Combining LSTM and CNN methods and fundamental analysis for stock price trend prediction," *Multimed Tools Appl*, vol. 82, no. 12, pp. 17769–17799, May 2023, doi: 10.1007/s11042-022-13963-0.

[8] C. Cui, P. Wang, Y. Li, and Y. Zhang, "McVCsB: A new hybrid deep learning network for stock index prediction," Dec. 01, 2023, *Elsevier Ltd*. doi: 10.1016/j.eswa.2023.120902.

[9] K. Gajamannage, Y. Park, and D. I. Jayathilake, "Real-time forecasting of time series in financial markets using sequentially trained dual-LSTMs," *Expert Syst Appl*, vol. 223, Aug. 2023, doi: 10.1016/j.eswa.2023.119879.

[10] K. Alam, M. H. Bhuiyan, I. ul Haque, M. F. Monir, and T. Ahmed, "Enhancing Stock Market Prediction: A Robust LSTM-DNN Model Analysis on 26 Real-Life Datasets," *IEEE Access*, 2024, doi: 10.1109/ACCESS.2024.3434524.

[11] Q. Q. Abuein, M. Q. Shatnawi, E. Y. Aljawarneh, and A. Manasrah, "Time Series Forecasting Model for the Stock Market using LSTM and SVR," *International Journal of Advances in Soft Computing and its Applications*, vol. 16, no. 1, pp. 169–185, 2024, doi: 10.15849/IJASCA.240330.10.

[12] Y. Zhang, B. Yan, and M. Aasma, "A novel deep learning framework: Prediction and analysis of financial time series using CEEMD and LSTM," *Expert Syst Appl*, vol. 159, Nov. 2020, doi: 10.1016/j.eswa.2020.113609.

[13] H. Rezaei, H. Faaljou, and G. Mansourfar, "Stock price prediction using deep learning

.

and frequency decomposition,” *Expert Syst Appl*, vol. 169, May 2021, doi: 10.1016/j.eswa.2020.114332.

[14] F. Li, H. Zhou, M. Liu, and L. Ding, “A Medium to Long-Term Multi-Influencing Factor Copper Price Prediction Method Based on CNN-LSTM,” *IEEE Access*, vol. 11, pp. 69458–69473, 2023, doi: 10.1109/ACCESS.2023.3288486.

[15] A. Q. Md *et al.*, “Novel optimization approach for stock price forecasting using multi-layered sequential LSTM,” *Appl Soft Comput*, vol. 134, Feb. 2023, doi: 10.1016/j.asoc.2022.109830.

[16] H. N. Bhandari, B. Rimal, N. R. Pokhrel, R. Rimal, K. R. Dahal, and R. K. C. Khatri, “Predicting stock market index using LSTM,” *Machine Learning with Applications*, vol. 9, p. 100320, Sep. 2022, doi: 10.1016/j.mlwa.2022.100320.

[17] D. Wen, T. Zhao, L. Fang, C. Zhang, and X. Li, “MWDINet: A multilevel wavelet decomposition interaction network for stock price prediction,” *Expert Syst Appl*, vol. 238, Mar. 2024, doi: 10.1016/j.eswa.2023.122091.

[18] Y. Sun, S. Mutalib, N. Omar, and L. Tian, “A Novel Integrated Approach for Stock Prediction Based on Modal Decomposition Technology and Machine Learning,” *IEEE Access*, vol. 12, pp. 95209–95222, 2024, doi: 10.1109/ACCESS.2024.3425727.

[19] A. Kanwal, M. F. Lau, S. P. H. Ng, K. Y. Sim, and S. Chandrasekaran, “BiCuDNNLSTM-1dCNN— A hybrid deep learning-based predictive model for stock price prediction,” *Expert Syst Appl*, vol. 202, Sep. 2022, doi: 10.1016/j.eswa.2022.117123.

[20] Z. J. Peng, C. Zhang, and Y. X. Tian, “Crude Oil Price Time Series Forecasting: A Novel Approach Based on Variational Mode Decomposition, Time-Series Imaging, and Deep Learning,” *IEEE Access*, vol. 11, pp. 82216–82231, 2023, doi: 10.1109/ACCESS.2023.3301576.

[21] E. Antwi, E. N. Gyamfi, K. A. Kyei, R. Gill, and A. M. Adam, “Modeling and Forecasting Commodity Futures Prices: Decomposition Approach,” *IEEE Access*, vol. 10, pp. 27484–27503, 2022, doi: 10.1109/ACCESS.2022.3152694.

[22] W. Guo, Q. Liu, Z. Luo, and Y. Tse, “Forecasts for international financial series with VMD algorithms,” *J Asian Econ*, vol. 80, Jun. 2022, doi: 10.1016/j.asieco.2022.101458.

[23] K. Dragomiretskiy and D. Zosso, “Variational mode decomposition,” *IEEE Transactions on Signal Processing*, vol. 62, no. 3, pp. 531–544, Feb. 2014, doi: 10.1109/TSP.2013.2288675.

[24] Y. Li, S. Wang, Y. Wei, and Q. Zhu, “A New Hybrid VMD-ICSS-BiGRU Approach for Gold Futures Price Forecasting and Algorithmic Trading,” *IEEE Trans Comput Soc Syst*, vol. 8, no. 6,pp. 1357–1368, Dec. 2021, doi: 10.1109/TCSS.2021.3084847.

[25] S. Wang, “A Stock Price Prediction Method Based on BiLSTM and Improved Transformer,”*IEEE Access*, vol. 11, pp. 104211–104223, 2023, doi: 10.1109/ACCESS.2023.3296308.

[26] T. Liu, X. Ma, S. Li, X. Li, and C. Zhang, “A stock price prediction method based on meta- learning and variational mode decomposition,” *Knowl Based Syst*, vol. 252, Sep. 2022, doi: 10.1016/j.knosys.2022.109324.

[27] C. Zhang, N. N. A. Sjarif, and R. Ibrahim, “1D-CapsNet-LSTM: A deep learning-based model for multi-step stock index forecasting,” *Journal of King Saud University - Computer and Information Sciences*, vol. 36, no. 2, Feb. 2024, doi: 10.1016/j.jksuci.2024.101959.

[28] J. Baruník and S. Dvořáková, “An empirical model of fractionally cointegrated daily

high and low stock market prices," *Econ Model*, vol. 45, pp. 193–206, 2015, doi: https://doi.org/10.1016/j.econmod.2014.11.024.

[29] B. Gülmez, "Stock price prediction with optimized deep LSTM network with artificial rabbits optimization algorithm," *Expert Syst Appl*, vol. 227, Oct. 2023, doi: 10.1016/j.eswa.2023.120346.

.