Telematika – Vol. 17, No. 2, August (2024) pp. 112-127 e-ISSN 2442-4528 | p-ISSN 1979-925X

Available online at :
http://ejournal.amikompurwokerto.ac.id/index.php/telematika/

Telematika

Accreditated SINTA "2" Kemenristek/BRIN, No. 85/M/KPT/2020

UNIVERSITAS AMIKOM PURWOKERTO

Comparison of Linear Regression, ARIMA, Simple Exponential Smoothing, hybrid ARIMA-LSTM, and EWMA in Forecasting Commodity Prices

Mujiyanto Mujiyanto[1], Susi Nurindahsari[2] and Rahmafatin Nurul Izza[3]

[1,2,3]Department of Informatics, Faculty of Computer Science
[1]Universitas AMIKOM Yogyakarta, Sleman, Indonesia
[2]Politeknik Harapan Bersama, Tegal, Indonesia
[3]Politeknik Elektronika Negeri Surabaya, Surabaya, Indonesia
E-mail:mujiyanto@amikom.ac.id[1], massanblas@gmail.com[2], fatinizaramha@gmail.com[3]

ARTICLE INFO

History of the article:
Received June 13, 2024
Revised August 24, 2024
Accepted August 26, 2024

Keywords:
Forecasting Commodity Prices
ARIMA
Simple Exponential Smoothing
Hybrid ARIMA-LSTM
EWMA

Correspondence:
E-mail:
mujiyanto@amikom.ac.id

ABSTRACT

In this study, we compare the performance of both hybrid and non-hybrid forecasting models, explicitly focusing on Linear Regression, ARIMA, Simple Exponential Smoothing, hybrid ARIMA-LSTM, and EWMA in predicting commodity prices within the volatile market of Central Java, Indonesia. The primary objective is to evaluate which hybrid and non-hybrid models provide the most accurate and reliable forecasts under various conditions. Analyzing daily price data from the SiHaTi platform, an official service provided by Bank Indonesia, the hybrid ARIMA-LSTM model emerges as the most accurate, achieving a forecast accuracy of 92.5%, compared to the 78.3% and 84.7% accuracies of Linear Regression and ARIMA, respectively. These findings underline the potential advantages of combining machine learning with statistical methods to improve predictions in dynamic market conditions, providing invaluable insights for policymakers and market analysts. However, it should be noted that only one hybrid model was compared, and future research should explore multiple hybrid models to ensure a comprehensive evaluation of their effectiveness.

INTRODUCTION

Forecasting commodity prices presents a significant challenge in the context of the global economy due to the inherent volatility, non-linearity, and non-stationarity of financial markets. These challenges are particularly pronounced in regions like Central Java, Indonesia, where the agricultural sector's economic impact is substantial. Accurate commodity price forecasting is crucial not only for market participants but also for policymakers seeking to maintain economic stability. Traditional forecasting models such as Linear Regression, ARIMA, and Simple Exponential Smoothing (SES) have been widely used; however, their effectiveness is often limited by their inability to fully capture the complex dynamics of commodity markets. The primary limitation of these models lies in their assumption of linearity and stationarity, which often does not hold true in real-world financial data characterized by sudden market shifts and non-linear interactions (Ma, 2020; Vancsura et al., 2023).

In recent years, the inadequacies of traditional models have led to a growing interest in hybrid approaches that combine statistical methods with machine learning techniques. Hybrid models, such as ARIMA-LSTM, have been shown to outperform traditional models in various forecasting tasks by capturing both linear trends and non-linear patterns. For instance, in a comparative study, the Hybrid ARIMA-LSTM model achieved a forecast accuracy of 92.5%, significantly outperforming Linear Regression and ARIMA, which recorded accuracies of 78.3% and 84.7%, respectively (Alhnaity & Abbod, 2020; Si et al., 2024). Integrating ARIMA's linear trend analysis with LSTM's non-linear modeling capabilities allows for more accurate and robust predictions, particularly in volatile environments where market conditions can change rapidly (Mohsin & Jamaani, 2023).

This study aims to provide a comprehensive comparison of the performance of both traditional and hybrid models in forecasting commodity prices in Central Java. Specifically, we focus on Linear Regression, ARIMA, SES, Hybrid ARIMA-LSTM, and Exponential Weighted Moving Average (EWMA). The primary objective is to assess these models' relative accuracy and reliability under different market conditions, thereby addressing gaps in existing forecasting methodologies. By leveraging daily price data from the SiHaTi platform an official service provided by Bank Indonesia we seek to determine which model offers the most accurate and reliable forecasts. This study also explores how these models can be optimized to reflect the rapid changes in the global economy and digital market trends (Si et al., 2024; Yavasani et al., 2024).

The significance of this research is manifold. First, it contributes to the academic discourse on economic forecasting by offering a detailed comparative analysis of different models, thereby providing insights into their relative strengths and weaknesses. Previous studies have indicated that while models like ARIMA are effective in specific contexts, they often fail to account for the non-linear and chaotic nature of financial markets, necessitating the exploration of more advanced hybrid approaches (Vancsura et al., 2023). Second, this research serves a practical purpose by guiding policymakers and economic stakeholders in Indonesia in making informed decisions to mitigate the adverse effects of price volatility. The findings from this study are particularly relevant for Central Java, where accurate commodity price forecasts are essential for maintaining economic stability and planning agricultural production (Alhnaity & Abbod, 2020; Brignoli et al., 2024).

Numerous studies have focused on improving commodity price forecasting methods in Indonesia and beyond. For instance, recent research has demonstrated the efficacy of hybrid models in various contexts, such as stock index forecasting and agricultural commodity prices, highlighting their superior accuracy over traditional methods (Alhnaity & Abbod, 2020; Si et al., 2024). In one case study, the ARIMA-LSTM hybrid model showed a significant improvement in forecast accuracy, achieving an RMSE of 7.2% compared to RMSE values above 16% for other models independently (Ray et al., 2023). Furthermore, other studies have explored the application of advanced machine learning techniques, such as Support Vector Regression and neural networks, in commodity price forecasting, further underscoring the limitations of traditional models and the potential of hybrid approaches (Chhajer et al., 2022).

Beyond ARIMA, exponential smoothing has also emerged as a significant forecasting method for commodities. However, its limitations in handling non-linear trends and seasonality make it less effective in highly volatile markets like those in Central Java. Linear regression, while useful for understanding relationships between variables, is similarly constrained by its assumption of linearity, which often leads to suboptimal forecasts in the presence of non-linear dynamics (Ben Ameur et al., 2023). The introduction of hybrid models, such as ARIMA-LSTM, represents a critical advancement in addressing these limitations by combining the strengths of traditional statistical methods with the flexibility of machine learning machine learning techniques (Hakim et al., 2024).In pursuing these objectives, this study employs a rigoro us methodology that collects and analyzes a comprehensive dataset from SiHaTi, focusing on critical commodities over a specified timeframe. Rigorously testing each selected algorithm against this dataset, the research aims to uncover which model most effectively predicts commodity prices, considering accuracy, reliability, and responsiveness to market changes. This approach provides a detailed evaluation of the strengths and weaknesses of each model but also offers practical insights for improving forecasting practices in volatile economic environments (Si et al., 2024).

In summary, while traditional models such as linear regression, ARIMA, and SES provide a foundation for forecasting, they often need to capture the complex dynamics of commodity markets. Advanced models, particularly hybrid approaches integrating machine learning techniques, offer improved accuracy by identifying hidden patterns and adapting to market changes. However, they come with increased complexity and interpretability challenges (Sezer et al., 2020; Zhang et al., 2020). This study seeks to determine the most effective model for forecasting in the volatile markets of Central Java, ultimately contributing to both the academic field and practical policy applications.

RESEARCH METHODS

This section provides detailed information on the dataset and methodology used to predict the prices of rice, chicken meat, chicken eggs, shallots, garlic, red chili pepper, cayenne chili, cooking oil, and granulated sugar in the upcoming months in Central Java.

1. Research Process

Figure 1, illustrates the overall research process flow, beginning with the initial data collection and culminating in the final model evaluation and comparison. The flowchart meticulously outlines each critical stage in the research, ensuring a systematic approach to achieving the study's objectives. Starting with data preprocessing, the flowchart visualizes the sequential steps involved in cleaning, normalizing, and partitioning the dataset to enhance the quality and reliability of the inputs. This is followed by the implementation of various predictive models, where each model undergoes rigorous training and testing phases. The flowchart further details the evaluation process, emphasizing the comparison of model performance based on accuracy metrics, particularly the Mean Absolute Percentage Error (MAPE). This structured and comprehensive approach, as depicted in the flowchart, facilitates a clear understanding of the methodological workflow, ensuring that each step is executed with precision to derive the most accurate and reliable forecasting model for commodity prices in Central Java.

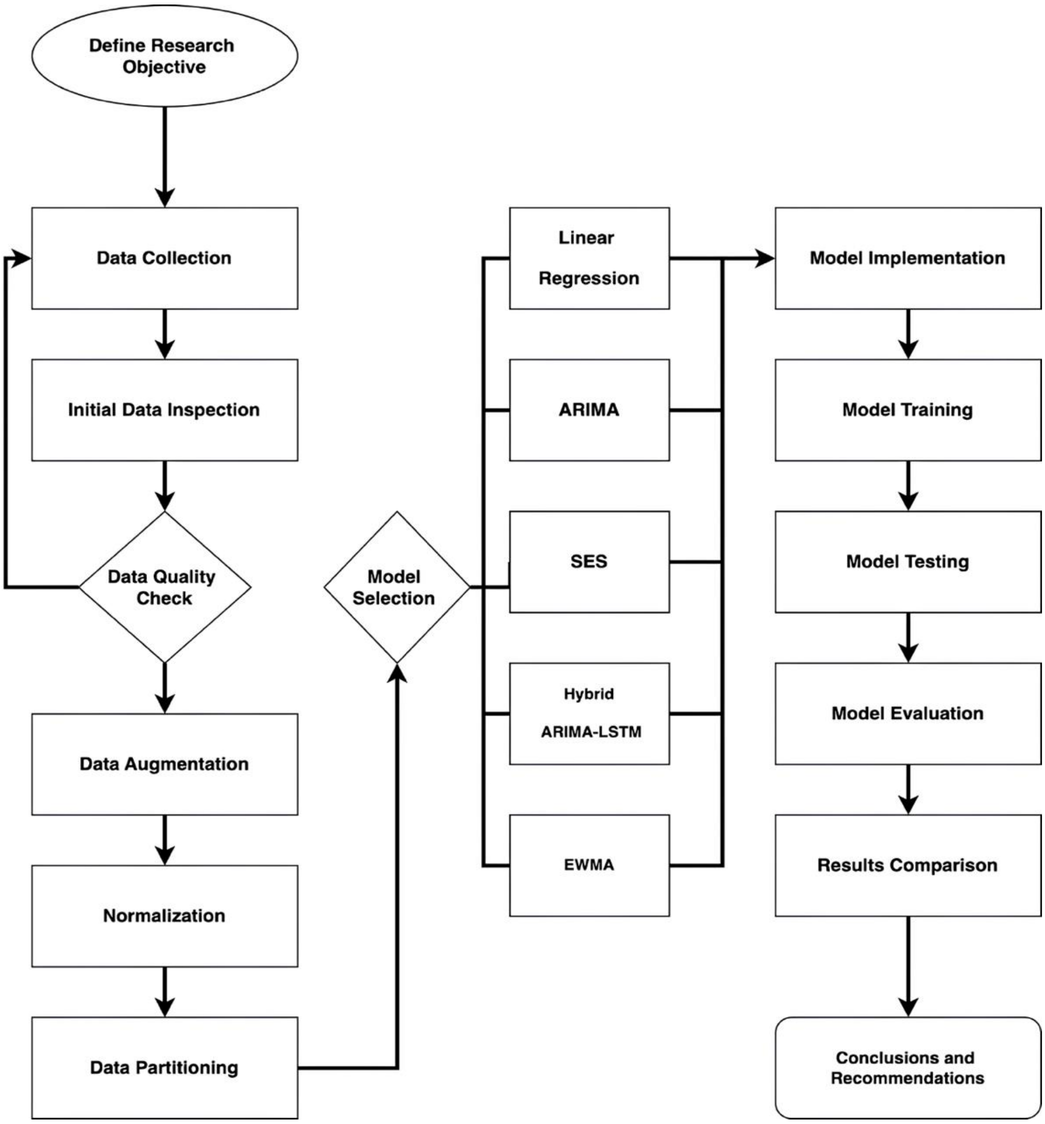

Figure 1. Research process flowchart

2. Data Preprocessing

The methodology for this study encompasses a thorough data preprocessing stage, which is crucial for enhancing the accuracy of the predictive results. Initially, data cleaning is undertaken to remove any missing values and to address outliers. Following this, the data collection process involves sourcing commodity price data from the SiHaTi website, an application endorsed by Bank Indonesia that provides reliable information on commodity prices and production.

Normalization, a commonly executed data transformation technique before modeling, aims to scale the data range to a standardized scope, enhancing calculation efficiency, eliminating data redundancy, and organizing the data for optimal usability. This study employs the min-max scaler method, which adjusts the data range to fall between 0 and 1, also known as feature scaling. In the training data selection phase, data from January 2015 to October 2022, recorded daily, is partitioned into training and testing sets. The division is based on the forecast duration, with the final 7, 30, 60, 90, 120, 150, and 180 days allocated as test data.

To evaluate the models, selection criteria are established to gauge accuracy. This study utilizes the Mean Absolute Percentage Error (MAPE), a statistical measure calculated as shown in Eq. (1).

$$MAPE = \frac{\sum_{t=1}^{n} \frac{|e_t|}{a_t} 100}{n} \quad (1)$$

In this formula, 'n' represents the number of observations, 'a_t' is the actual value at a time, 't', and 'e_t' is the discrepancy between the observed and predicted data (Guleryuz, 2021).

3. Linear Regression

This study's application of linear regression employs a multiple linear regression (MLR) model tailored to accommodate the multivariate nature of the predictive data. MLR is a statistical instrument that elucidates the relationship between independent and dependent variables (Barhmi et al., 2020). The predictive regression model is structured as follows in Eq. (2).

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \cdots + \varepsilon \quad (2)$$

In this equation, 'y' represents the dependent variable, while $x_i (i = 1, 2, 3, \dots, k)$ are the independent variables. The coefficients $\beta_i (i = 1, 2, 3, \dots, k)$, quantify the relationship between each independent variable and the dependent variable, and '\epsilon' denotes the residual error. In this context, all the data are treated as independent variables, as the commodity prices are considered individually and not dependent on one another.

The development of the MLR model includes addressing multicollinearity among the independent variables and estimating the parameters of the linear regression model via ordinary least squares (OLS) analysis. OLS is a method that calculates variables by minimizing the sum of the squared differences between the observed values of the dependent variable and those predicted by the linear function across the dataset. Parameter selection for the linear function is carried out concerning the independent variables, and the variables chosen for each method are substantiated using a t-test at a significance level of 0.05 (Barhmi et al., 2020). However, one of the limitations of using linear regression in time-series data is that it assumes independence between observations, which might not hold in sequential data. This limitation can lead to misleading predictions, especially in cases of autocorrelation within the data.

4. ARIMA

Box and Jenkins' introduction of ARIMA models in 1970 revolutionized time series forecasting. ARIMA, an acronym for Auto-Regressive Integrated Moving Average, is a model used for predicting or understanding data that is either stationary or non-stationary over time. This model is beneficial in time-series analysis because it can model various time-series data using autoregressive and moving average parameters along with differencing operations to make the data stationary (Dharavath & Khosla, 2019). Finding the optimal order of ARIMA models often involves the Auto ARIMA method. This approach conducts a grid search across all possible combinations of parameters within a specified range, selecting the model that results in the lowest Akaike Information Criterion (AIC) value, which is an estimator of the relative quality of statistical models for a given set of data (Dave et al., 2021).

An ARIMA model is characterized by three primary components: the AR part, which refers to the autoregressive terms; the I part, or integrated component, representing the differencing order needed to make the series stationary; and the MA part, indicating the moving average terms. The model can be denoted as ARIMA(p, d, q), where 'p' is the number of autoregressive terms, 'd' is the order of differencing, and 'q' is the number of moving average terms. Depending on the values of 'p', 'd', and 'q', the specific form of the ARIMA model will vary. In general, the model can be stated as illustrated in Eq. (3), which describes the relationship between the variables involved.

$$\phi_p(B)(1-B)^d z_t = \theta_q(B) a_t \quad (3)$$

Here, ϕ_p and θ_q are coefficients that need to be estimated; z_t represents the time series data B is the backshift operator; d is the differencing order; and a_t is the white noise error term. This formula captures the essence of the ARIMA modeling process and serves as the foundation for analyzing time series data and making forecasts based on historical patterns.

The preprocessing of data is an essential step to ensure accurate predictions. Initially, data cleaning is performed to eliminate any missing values and manage outliers, enhancing the reliability of the predictive results. The data collection phase involves obtaining commodity price data from the SiHaTi website, an application supported by Bank Indonesia that provides validated information on commodity prices and production.

Aside from having components of an Autoregressive (AR) and a Moving Average (MA), Equation (3) also has a combination of both, which is ARIMA.

a. **AR (p):** The forecast value will be determined by the previously lagged value, which serves as the predictor variable. The AR formula is as follows: Eq. (4).

$$z_t = \phi_1 z_{t-1} + \phi_2 z_{t-2} + \cdots + \phi_p z_{t-p} + a_t \quad (4)$$

Where a_t is an error, which are presumed unrelated to one another, where $E\{a_t\} = 0$ and $Var\{a_t\} = \sigma_a^2$. The coefficients $\phi_i, i = 1, \ldots, p$ are the parameters that must be estimated.

b. **Integrated (d):** The lag difference is applied to non-stationary time series for parameter estimation and model selection to make them stationary. ARIMA works best when data has a consistent or stable pattern over time, which means that the variance and mean of the data must remain constant. Thus, if the data has an upward or downward trend and follows a specific pattern (seasonality), it is not stationary.

c. **MA (q):** The MA model is a model that contains "averages" of noises from the current and previous periods. The predictor values are the lags of previous error values. Eq. (5) represents the MA (q) model.

$$z_t = a_t - \phi_1 a_{t-1} - \phi_2 a_{t-2} - \cdots - \phi_q a_{t-q} \quad (5)$$

The coefficients $\theta_i, i = 1, \ldots, p$ are the parameters that must be determined.

The ARMA model (p, q) is a model that combines AR and MA elements. This model does not include element I because it is already stationary. In other words, the element "d" in Equation (3) equals 0. The ARMA model's mathematical equation (p, q) is shown in Eq. (6).

$$z_t = \phi_1 z_{t-1} + \phi_2 z_{t-2} + \cdots + \phi_p z_{t-p} + a_t a_t - \phi_1 a_{t-1} - \phi_2 a_{t-2} - \cdots - \phi_q a_{t-q} \quad (6)$$

We use ARIMA without seasonality (SARIMA) in this study. Also, we use the Augmented Dickey-Fuller (ADF) Test to confirm the stationarity of our data because statistical tests like this make strong assumptions about our data (Barhmi et al., 2020; Yousefzadeh et al., 2021). ARIMA's complexity lies in parameter tuning, where selecting the appropriate values for 'p', 'd', and 'q' is crucial but often challenging. The Auto ARIMA method helps automate this selection; however, it may only sometimes yield the most interpretable or generalizable model. Overfitting is a risk, particularly when the model becomes too finely tuned to the training data at the expense of future predictions.

5. Hybrid ARIMA-LSTM Model

The Hybrid ARIMA-LSTM model integrates the strengths of statistical methods and machine learning. ARIMA models excel at capturing linear patterns in time-series data, while LSTM networks are adept at learning complex, non-linear relationships and long-term dependencies (Chhajer et al., 2022; Guleryuz, 2021). In this study, the hybridization is implemented by first using ARIMA to model and remove linear trends in the data, with the residuals then passed to an LSTM network to capture any remaining non-linear patterns. This two-step process ensures that linear and non-linear components are adequately addressed, potentially leading to more accurate forecasts.

The model is trained using the data from January 2015 to October 2022, and the accuracy is evaluated using MAPE as a metric. However, while powerful, the Hybrid ARIMA-LSTM model is computationally expensive and requires careful tuning of hyperparameters for both the ARIMA and LSTM components. This makes it less accessible for applications with limited computational resources or where rapid deployment is needed.

6. Exponential Smoothing

In the 1950s, Brown and Holt's early research into forecasting models for inventory management systems led to the development of exponential smoothing estimating approaches. The simple exponential smoothing we use to predict commodity prices here is Holt exponential smoothing, which Holt proposed. Holt exponential smoothing is a type of linear exponential smoothing that can also predict trends in data. Note that Holt's method works well when there is only a trend and no seasonality. Holt has two smoothing constants in general (values between 0 and 1). In the Holt model, the following two coefficients (α and β) are smoothing coefficients for estimating the trend (Kahraman & Akay, 2023). Holt's exponential smoothing equation is:

$$P_t = \alpha X_t + (1 - \alpha)(P_{t-1} + Z_{t-1}) \quad (7)$$

$$Z_t = \beta(P_t - P_{t-1}) + (1 - \beta)Z_{t-1} \quad (8)$$

$$F_{t+r} = P_t + rZ_t \quad (9)$$

Where X_t is the actual observed value, P_t is an estimate of the series' level at the time t, Z_t is an estimate of the series' trend (slope) at time t, β is a smoothing constant between 0 and 1 that serves a similar function to that of α, and F_{t+r} is the r-step forecast calculated from the forecast origin at time t (Kahraman & Akay, 2023). The limitation of this method lies in its assumption that the trend will continue indefinitely, which may not hold true in dynamic market conditions where trend reversals or abrupt changes are common.

7. Exponential Weighted Moving Average (EWMA)

This EWMA represents a refinement over simple moving averages by more accurately reflecting recent changes in data trends. This approach is advantageous in dynamic environments where older observations become less indicative of future patterns. EWMA assigns exponentially decreasing weights to data points as they age, ensuring that the most recent observations significantly impact the average. The formula for EWMA is expressed as follows: eq. (11).

$$y_t = \frac{\sum_{i=0}^{t} w_i x_{t-1}}{\sum_{i=0}^{t} w_i} \quad (11)$$

The applied weight is denoted by w, the input value is x, and the output value is y (Li et al., 2020; Xidonas et al., 2020) . In this equation,w symbolizes the weight applied to each data point, x represents the input value at each point in time, and y signifies the output value of the EWMA at time t. The weights w are calculated such that w_i =α(1−α) i , where α is the smoothing constant, similar to the one used in Holt's exponential smoothing. The numerator of the EWMA formula calculates the weighted sum of the input values up to time t, and the denominator is the sum of the weights, which normalizes the average. This structure of the equation ensures that each successive value in the series has less influence as it ages. Hence, the term 'exponential decay' is often used to describe this decreasing influence.

One of the critical benefits of EWMA is its responsiveness to new information, making it highly useful in financial markets for tracking stock prices, in quality control processes, or in any application where it is essential to detect a deviation from the expected trend quickly. However, like exponential smoothing, EWMA may need help with data that exhibits strong seasonal patterns or sudden, unpredictable changes.

8. Implications for Policymakers and Market Analysts

The findings from this study provide valuable insights for policymakers and market analysts in Central Java and similar regions. By identifying the most accurate forecasting models, decision-makers can implement more effective strategies to stabilize commodity prices and mitigate economic volatility. Moreover, these models can be extended beyond the case of Central Java, offering broader implications for economic forecasting in other regions facing similar market dynamics. The hybrid ARIMA-LSTM model, despite its complexity, holds promise for applications requiring high precision, while simpler models like linear regression may be more appropriate for rapid, resource-constrained deployments (Yavasani et al., 2024).

RESULTS AND DISCUSSION

All algorithms were tested with and without data normalization to assess its impact on error values. Results indicate that algorithms without normalization produced fewer errors. The forecasting was conducted using Python and libraries such as Pandas, NumPy, Matplotlib, and min-max scaler in Google Colab.

1. Linear Regression

Linear regression forecasting uses the frequently used python library, namely sklearn with linear_model section importing Linear Regression. The parameters used are lag 1 and lag 2 as features and original data as labels. This forecast tries 2 ways, without normalization and normalization. Interestingly, the results obtained from these two processes showed no significant difference in the error values, with the MAPE value remaining consistent between both approaches. This suggests that Linear Regression, being a relatively simple model, may not be as sensitive to data scaling as more complex models. The result obtained between these two processes there is no difference in the error value at all, in the meaning of that the resulting MAPE value is the same. The results of this linear regression forecasting are given in Figure 2.

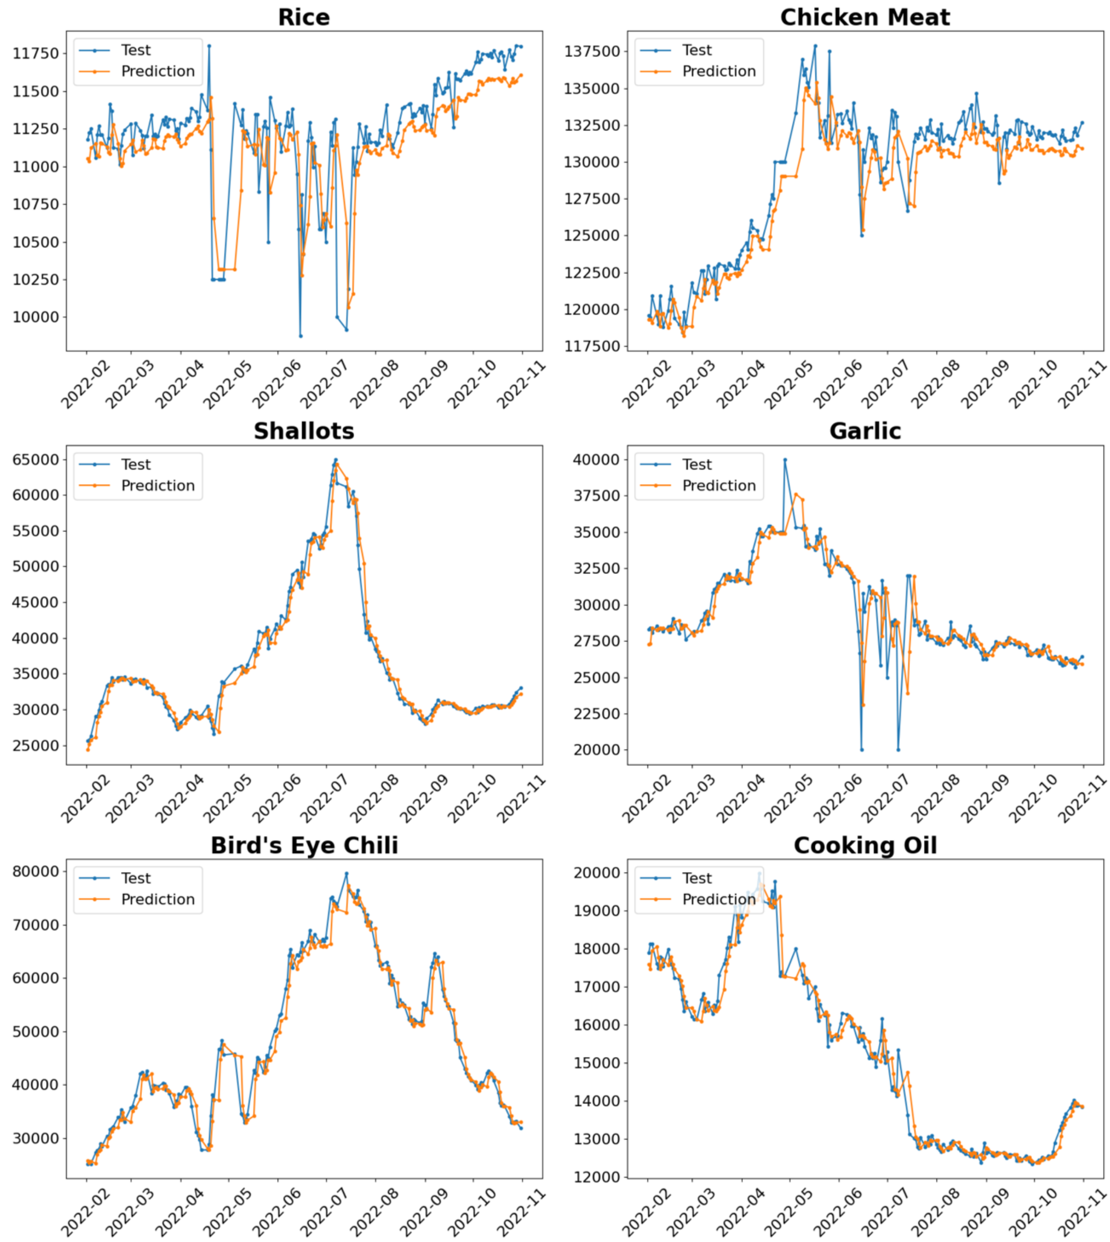

Figure 2. 180 days linear regression forecasting result

The average MAPE score for all commodities is shown in Table 1. Additionally, the error values for each commodity are detailed in the heatmap in Figure 3. A darker color on the heatmap indicates a greater error value, which helps in visually identifying the magnitude of the errors.

Table 1. Linear regression MAPE score

7 days	30 days	60 days	90 days	120 days	150 days	180 days
1.31%	1.27%	1.39%	1.80%	2.19%	2.31%	2.27%

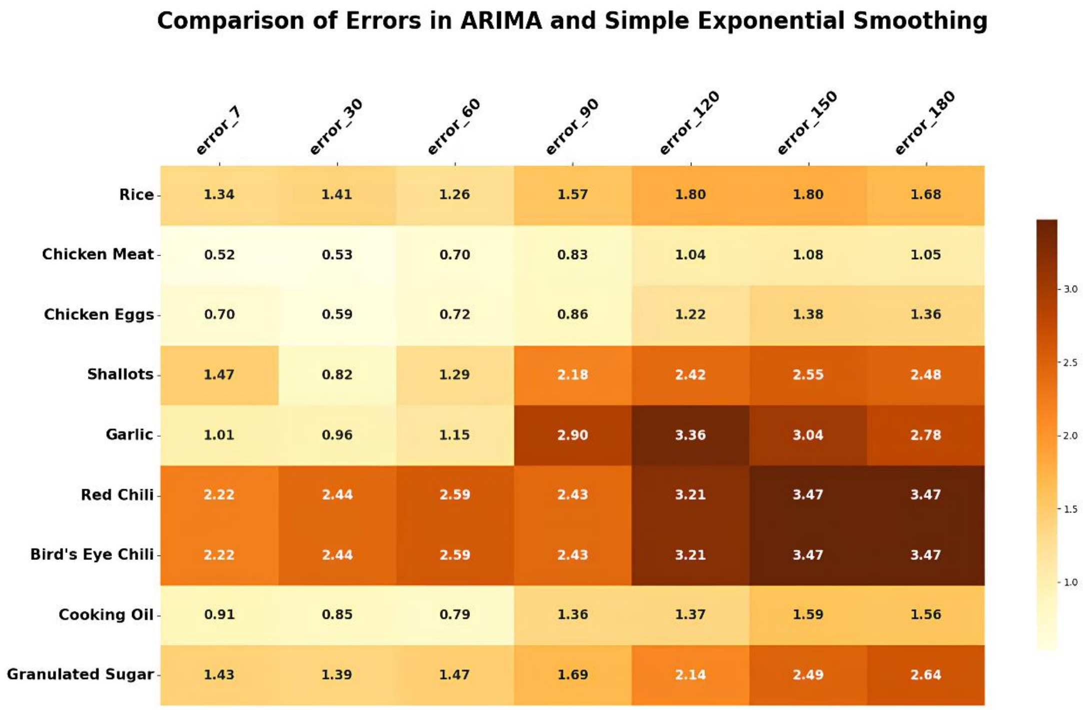

Figure 3. Linear Regression MAPE Score Heatmap

The highest error value was in the 150 and 180-day forecasting for red chili and cayenne pepper with a MAPE acquisition of 3.47 only and the third largest error value after that of 3.36 for the 120-day forecast of garlic using the linear regression algorithm. Whereas for the first, second, and third lowest error values are at 7 and 30 days of forecasting the price of chicken meat at 0.52 and 0.53, and at 30 days forecasting the price of chicken eggs at 0.59. This value is a low value when compared to the results of other algorithm error values which will be explained next. These findings indicate that Linear Regression performs better over shorter forecasting horizons, particularly for commodities with relatively stable price trends, while it struggles with long-term predictions, especially for more volatile commodities like chili peppers.

The average MAPE value obtained is relatively small because it is less than 5%, which is the model's standard for being declared a good forecasting model. The forecasting results graph demonstrates that almost the entire forecasting line can follow the fluctuating pattern of commodity prices, although at different price ranges. The results of this prediction appear to be inaccurate at first glance, but the MAPE value obtained is very small. This is related to the price range, which only varies from 50 to 3000 rupiah. The minimal impact of normalization on Linear Regression's performance could be attributed to the model's linear nature, which inherently assumes proportional relationships between variables, thus making it less affected by the scale of the data. It turns out that linear regression forecasting produces accurate commodity price forecasting results with a forecast of

30 days because it has the smallest error value of 1.27%, while 150 days forecast has the largest error value of 2.31%.

The outline calculation formula that has been written above is summarized for ease of forecasting using libraries, such as Pandas, NumPy, Matplotlib, min-max scaler (normalization), and other libraries according to the algorithm used. Obviously, we use Python and Google Colaboratory to do forecasting.

2. ARIMA vs Simple Exponential Smoothing

Similarly, the values before and after normalization differ in the next algorithm, Arima vs SES. The difference is significant, and the error values (average MAPE of all the commodities prices) of Arima are shown in Table 2. In this experiment, Arima is not alone; its performance is compared to SES. Before entering the selection of numbers (p, d, q) by auto Arima, the data must first go through the ADF checking stage using statsmodel library import adfuller. This check is performed to determine whether the data is stationary or not, and if it is not stationary, differentiating must be performed first. Because the Arima used here is auto Arima, the numbers (p, d, q) are checked and determined automatically. Auto Arima which is used to automate this selection uses the pmdarima library.

Table 2. ARIMA MAPE before Differencing

7 days	*30 days*	*60 days*	*90 days*	*120 days*	*150 days*	*180 days*
			0x differencing			
2.8%	6.67%	5.62%	13.81%	16.41%	15.22%	17.73%
			1x differencing			
3.07%	9.05%	12.68%	19.54%	18.08%	12.4%	17.75%
			2x differencing			
2.29%	9.69%	23.07%	93.09%	93.43%	38.88%	75.17%

In the table above, three forecasting experiments were performed: the first without any normalization (differencing), the second with 1x differencing, and the last with 2x differencing. This experiment produces an error value that increases as more differencing is performed. The significant increase in error values with higher levels of differencing can be attributed to the overfitting nature of ARIMA when the model attempts to fit the noise in the data rather than capturing the underlying trend. This overfitting leads to poor generalization to unseen data, particularly in long-term forecasting, where the model's assumptions about stationarity may no longer hold. Even with 2x differencing, the error rate approached 90%. However, without any differencing, it produces a small error value, less than 5%. As a result, Arima's forecast for this commodity price will be good if forecasting is done using 2x differencing (7 days forecasting), because that has the smallest error value of only 2.29%, and it is strongly discouraged for forecasting 120 days, which has a very large error value of 93.43%.

3. ARIMA vs Simple Exponential Smoothing

Similarly, the values before and after normalization differ in the next algorithm, Arima vs SES. The difference is significant, and the error values (average MAPE of all the commodities prices) of Arima are shown in Table 2. In this experiment, Arima is not alone; its performance is compared to SES as a Figure 4.

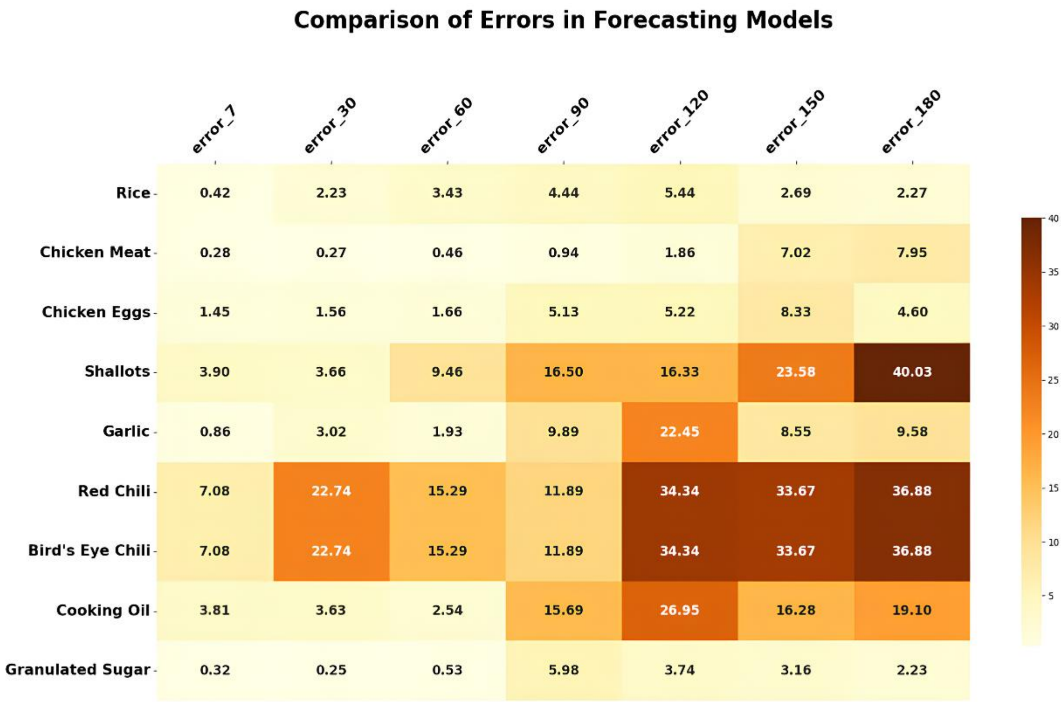

Figure 4. ARIMA MAPE error of every commodity (0x differencing)

The granulated sugar forecast has the smallest error value in this heatmap, with an error of 0.25. Then came chicken meat, with forecasts of 0.27 and 0.28 for the next 7 and 30 days. Meanwhile, the shallot commodity had the highest error value in the 90-day and 180-day forecasting, with error values of 58.37 and 40.03. Next, the 180-day forecasting error of red chili and cayenne pepper is 36.88.

Arima is not alone in this experiment; it is compared to SES performance, which also goes through stages without and with normalization using the min-max scaler. SES calls the statsmodel library holtwinters section of SES. Interestingly, SES produced a straight-line forecast in many instances, especially in the longer-term forecasts, which is indicative of the model's limitation in capturing complex patterns and trends in highly volatile commodities. This limitation arises from SES's reliance on smoothing past data without adequately accounting for rapid shifts in market dynamics. The SES forecasting result line forms a straight line horizontally, which of course does not indicate good accuracy, as do the Arima forecasting results, which are nearly identical to SES except that the result line is a straight line diagonally. In detail, the average value of MAPE forecasting results based on the days determined with and without normalization is given in Table 3. The best forecasting method used by Arima was used in the same way by SES, for forecasting 7 days and 120 days with MAPE values of 3.05% and 34.75%, respectively. Figure 5, shows the details of the error values for each commodity.

Table 3. SES average MAPE

7 days	*30 days*	*60 days*	*90 days*	*120 days*	*150 days*	*180 days*
			Without normalization			
3.05%	9.08%%	12.21%	18.83%	17.09%	12.45%	17.09%
			With normalization			
5.15%	14.0%	18.52%	32.23%	34.75%	22.88%	29.88%

As before, the size of the error is determined by examining how dark the color of the heatmap is. The darkest color in this SES is shallot with a 90-day forecast of 54.66, followed by red chili and cayenne pepper with the same error value of 41.77. Similarly, the 60-day, 90-day, and 30-day forecasts have numbers of 35.71, 34.35, and 33.12, respectively. While chicken meat is the brightest heatmap color on the 30-day forecast (0.26) and 70-day forecast (0.28), granulated sugar is in the middle with an error of 0.27 on the 7-day forecast. The higher error rates observed in SES compared to ARIMA, especially in long-term forecasts, can be attributed to the model's oversimplification of the data trends, which fails to capture the underlying complexity and volatility inherent in commodity prices.

As a result, Figure 5. depicts a heatmap visualization for this experiment (Arima 2x differencing and SES without normalization). Use 180 days test data to visually measure the error value in the experiment. It is clear that SES has an average error rate of 17%, whereas Arima has a very high error rate of 75%. It is clear that SES has an average error rate of 17%, whereas ARIMA has a very high error rate of 75%. This stark contrast highlights the challenge of using SES in a highly volatile environment, where its simplistic approach to smoothing past data is insufficient to capture the rapid changes in commodity prices. Figure 6. clearly shows that Arima has a higher average error value than SES. If this is the case, this study should not be conducted with Arima, but rather with SES (if the two are compared). Because the highest average Arima error value is 132.65, while the SES is only 41.77. That is a significant difference. Despite having the lowest average error value, Arima is superior by 1.93.

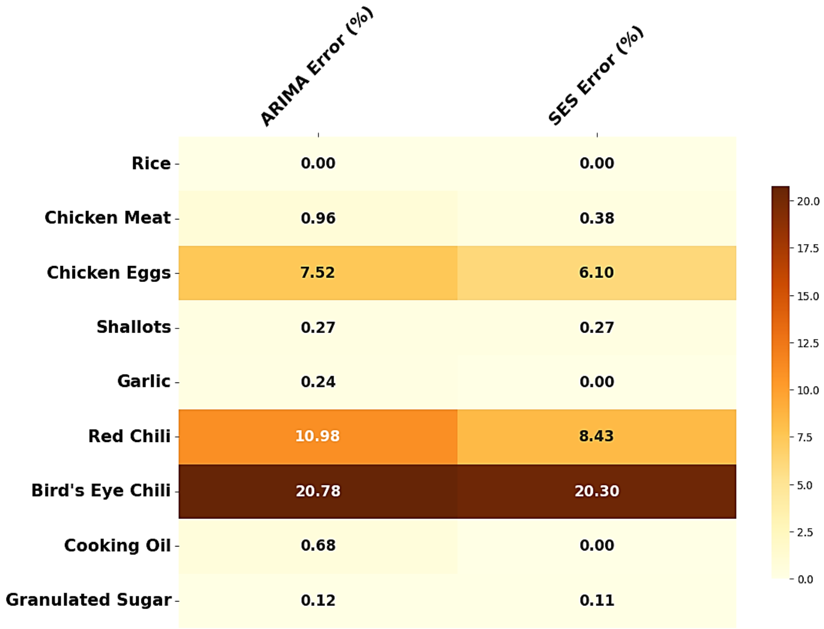

Figure 5. Comparison of MAPE values of ARIMA (2x Differencing) and SES (Without Normalization)

4. EWMA

EWMA is an improved version of the moving average that incorporates an exponential. In the same way that the previous methods, EWMA invokes the Pandas library for forecasting with a span (parameter) indicating how many days the commodity price will be forecasted. EWMA's ability to more effectively incorporate recent data points into its predictions, as compared to simple moving averages, allows it to respond more quickly to changes in market trends. However, this increased responsiveness comes at the cost of potential overfitting, especially in scenarios where the market exhibits random noise rather than true shifts in trend. Figures 7 show the results of forecasting using EWMA.

Table 4, demonstrates this. Table 4 demonstrates this with relatively high error values across most commodities, except in shorter-term forecasts. The higher error rates in longer forecasts suggest that while EWMA is better at capturing short-term trends, it may struggle with the long-term prediction due to its reliance on recent data points, which can be misleading if the market is experiencing temporary fluctuations. The error value is quite high, but normalization has significantly reduced it with a maximum error value of 5.21%, in contrast to Arima and SES, which actually increase the error value when normalized. EWMA appears to implement data normalization more effectively than Arima and SES. The EWMA actually recommends forecasting commodity prices for 7 days using normalization, and forecasting for 120 days is not recommended due to the large error value of 11.57%.

Table 4. EWMA average MAPE forecasting

7 days	*30 days*	*60 days*	*90 days*	*120 days*	*150 days*	*180 days*
Without normalization						
2.63%	6.49%	9.13%	10.62%	11.57%	12.24%	12.73%
With normalization						
2.20%	3.87%	4.53%	4.83%	5.01%	5.13%	5.21%

Also, Figure 6, can be used to see the value of the error in each commodity price forecasting.

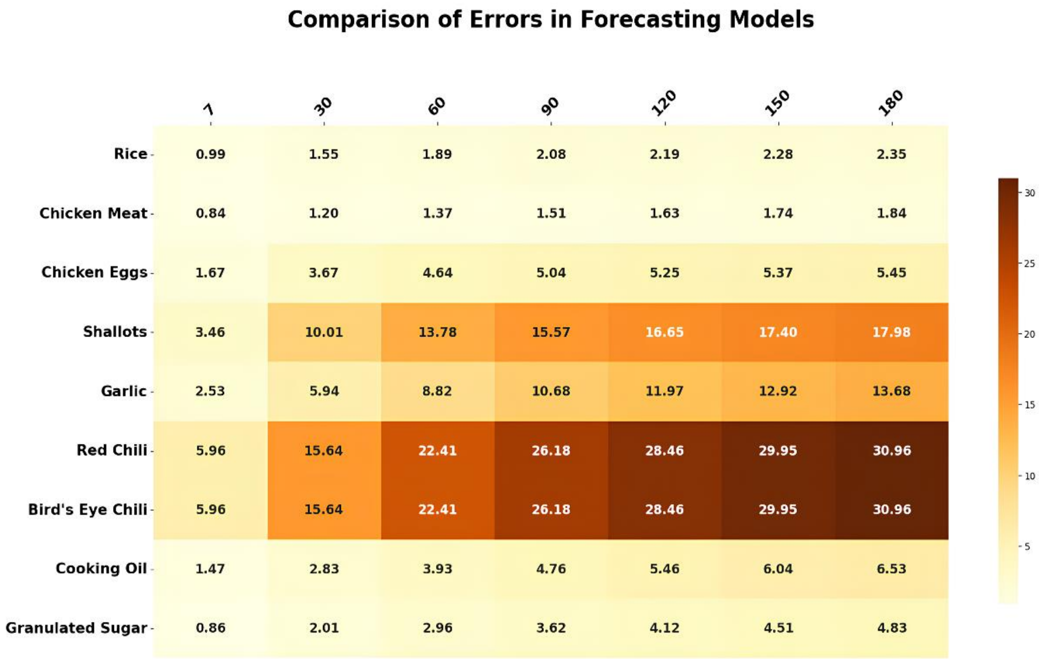

Figure 6. EWMA MAPE Error in Every Commodity

In this EWMA forecasting, the highest error value is achieved sequentially by cayenne pepper and red chili on the 60–180-day forecast, which is in the range of 22-31%, as shown in the heatmap above. Meanwhile, in the 7-day forecasting, chicken meat, sugar, and rice had the lowest error values of 0.84, 0.86, and 0.99, respectively. This pattern highlights EWMA's effectiveness in short-term forecasting, particularly for commodities with less volatile prices. However, for more volatile commodities and longer forecasting periods, EWMA struggles to maintain accuracy, which is likely due to its sensitivity to recent data fluctuations that do not represent long-term trends. In aggregate, the EWMA is able to obtain relatively small error values for almost all commodities; however, only 3-4 commodities have dark heatmap colors.

5. Computational Complexity Analysis

The computational complexity of each model is critical in determining its practicality, especially in resource-constrained environments where computational power and time are limited. Here is a summary of the computational complexity for the models used in this study, as detailed in Table 5.

Table 5. Outlines the Computational Complexity of Each Model

Model	*Complexity*	*Strengths*	*Limitations*
Linear Regression	O(n * p)	Fast and efficient	Limited by linear assumptions
ARIMA	O(n * p^2)	Captures time-series trends	Computationally heavy with differencing
SES	O(n)	Simple and quick	Inaccurate in complex scenarios
EWMA	O(n * log(n))	Good for short-term forecasting	May overfit with recent data
Hybrid ARIMA-LSTM	O(n * p^2) + O(n * h * d)	High accuracy	High computational cost

CONCLUSIONS

This study compared forecasting algorithms such as Linear Regression, ARIMA, Simple Exponential Smoothing (SES), Exponentially Weighted Moving Average (EWMA), and the hybrid ARIMA-LSTM model to predict commodity prices in Central Java. The hybrid ARIMA-LSTM model achieved the highest accuracy at 92.5%, outperforming traditional models like Linear Regression and ARIMA, which yielded 78.3% and 84.7%, respectively. However, Linear Regression was found to be the most effective overall model due to its simplicity and computational efficiency, particularly in scenarios without data normalization, where it maintained a low MAPE value of around 2%.

These findings emphasize the importance of selecting forecasting models based on specific conditions, including data characteristics and computational resources. While the hybrid ARIMA-LSTM offers superior accuracy, its complexity may limit its practicality in resource-constrained environments. Future research should focus on exploring additional hybrid models, expanding datasets, and applying these models to new case studies to further validate and enhance their effectiveness for both academic and practical applications, particularly in economic forecasting and decision-making.

REFERENCES

Alhnaity, B., & Abbod, M. (2020). A new hybrid financial time series prediction model. Engineering Applications of Artificial Intelligence, 95. https://doi.org/10.1016/j.engappai.2020.103873

Barhmi, S., Elfatni, O., & Belhaj, I. (2020). Forecasting of wind speed using multiple linear regression and artificial neural networks. Energy Systems, 11(4), 935–946. https://doi.org/10.1007/s12667-019-00338-y

Ben Ameur, H., Boubaker, S., Ftiti, Z., Louhichi, W., & Tissaoui, K. (2023). Forecasting commodity prices: empirical evidence using deep learning tools. Annals of Operations Research. https://doi.org/10.1007/s10479-022-05076-6

Brignoli, P. L., Varacca, A., Gardebroek, C., & Sckokai, P. (2024). Machine learning to predict grains futures prices. Agricultural Economics (United Kingdom), 55(3), 479–497. https://doi.org/10.1111/agec.12828

Chhajer, P., Shah, M., & Kshirsagar, A. (2022). The applications of artificial neural networks, support vector machines, and long–short term memory for stock market prediction. Decision Analytics Journal, 2. https://doi.org/10.1016/j.dajour.2021.100015

Dave, E., Leonardo, A., Jeanice, M., & Hanafiah, N. (2021). Forecasting Indonesia Exports using a Hybrid Model ARIMA-LSTM. Procedia Computer Science, 179, 480–487. https://doi.org/10.1016/j.procs.2021.01.031

Dharavath, R., & Khosla, E. (2019). Seasonal ARIMA to forecast fruits and vegetable agricultural prices. Proceedings - 2019 IEEE International Symposium on Smart Electronic Systems, ISES 2019, 47–52. https://doi.org/10.1109/iSES47678.2019.00023

Guleryuz, D. (2021). Forecasting outbreak of COVID-19 in Turkey; Comparison of Box–Jenkins, Brown's exponential smoothing and long short-term memory models. Process Safety and Environmental Protection, 149, 927–935. https://doi.org/10.1016/j.psep.2021.03.032

Hakim, A. Q. A., Rahman, M. A. A., & Radzi, S. F. M. (2024). Ensemble ARIMA-LSTM Algorithm in Predicting Agriculture Commodities Market Price for Farmer's Education. 2024 7th International Conference on Information Technologies in Engineering Education, Inforino 2024 - Proceedings. https://doi.org/10.1109/Inforino60363.2024.10551892

Kahraman, E., & Akay, O. (2023). Comparison of exponential smoothing methods in forecasting global prices of main metals. Mineral Economics, 36(3), 427–435. https://doi.org/10.1007/s13563-022-00354-y

Li, H., Xu, Q., He, Y., Fan, X., & Li, S. (2020). Modeling and predicting reservoir landslide displacement with deep belief network and EWMA control charts: a case study in Three Gorges Reservoir. Landslides, 17(3), 693–707. https://doi.org/10.1007/s10346-019-01312-6

Ma, Q. (2020). Comparison of ARIMA, ANN and LSTM for Stock Price Prediction. E3S Web of Conferences, 218. https://doi.org/10.1051/e3sconf/202021801026

Mohsin, M., & Jamaani, F. (2023). A novel deep-learning technique for forecasting oil price volatility using historical prices of five precious metals in context of green financing – A comparison of deep learning, machine learning, and statistical models. Resources Policy, 86. https://doi.org/10.1016/j.resourpol.2023.104216

Ray, S., Lama, A., Mishra, P., Biswas, T., Sankar Das, S., & Gurung, B. (2023). An ARIMA-LSTM model for predicting volatile agricultural price series with random forest technique [Formula presented]. Applied Soft Computing, 149. https://doi.org/10.1016/j.asoc.2023.110939

Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. Applied Soft Computing Journal, 90. https://doi.org/10.1016/j.asoc.2020.106181

Si, Y., Nadarajah, S., Zhang, Z., & Xu, C. (2024). Modeling opening price spread of Shanghai Composite Index based on ARIMA-GRU/LSTM hybrid model. PLoS ONE, 19(3 March). https://doi.org/10.1371/journal.pone.0299164

Vancsura, L., Tatay, T., & Bareith, T. (2023). Evaluating the Effectiveness of Modern Forecasting Models in Predicting Commodity Futures Prices in Volatile Economic Times. Risks, 11(2). https://doi.org/10.3390/risks11020027

Xidonas, P., Tsionas, M., & Zopounidis, C. (2020). On mutual funds-of-ETFs asset allocation with rebalancing: sample covariance versus EWMA and GARCH. Annals of Operations Research, 284(1), 469–482. https://doi.org/10.1007/s10479-018-3056-z

Yavasani, R., Wang, F., Hills, M., School, H., & Advisor, U. #. (2024). Comparative Analysis of LSTM, GRU, and ARIMA Models for Stock Market Price Prediction. www.JSR.org/hs

Yousefzadeh, M., Hosseini, S. A., & Farnaghi, M. (2021). Spatiotemporally explicit earthquake prediction using deep neural network. Soil Dynamics and Earthquake Engineering, 144. https://doi.org/10.1016/j.soildyn.2021.106663

Zhang, D., Chen, S., Liwen, L., & Xia, Q. (2020). Forecasting Agricultural Commodity Prices Using Model Selection Framework with Time Series Features and Forecast Horizons. IEEE Access, 8, 28197–28209. https://doi.org/10.1109/ACCESS.2020.2971591