Multi-Horizon LSTM Forecasting of Indonesian Food Commodity Prices to Support the Free Nutritious Meal (MBG) Program

Deni Utama[1*], Della Novita Sari[2]

[1]Department of Information Systems, Faculty of Computer Science, Universitas Negeri Jakarta, DKI Jakarta, Indonesia
[2]Communication Studies, College of Communication, Fine Arts, and Media, University of Nebraska at Omaha, Nebraska, United States
E-mail: [1*]deniutama@unj.ac.id, [2]dellanovitasari@unomaha.edu

(Received: 26 Jan 2026, revised: 27 Feb 2026, accepted: 3 Mar 2026)

Abstract

This study develops a multi-horizon Long Short-Term Memory (LSTM) model to forecast food commodity prices in Indonesia, supporting the national Free Nutritious Meal (MBG) program. The MBG initiative increases daily and annual food demand for over 50 million beneficiaries, affecting rice, meat, eggs, vegetables, and cooking oil. Historical price data were used to train the model, which effectively captures nonlinear temporal patterns and seasonal variations. Forecast results indicate heterogeneous volatility patterns across commodities. Staple commodities such as rice and cooking oil exhibit relatively stable trends with minor monthly fluctuations, whereas perishable and climate-sensitive commodities such as chili, eggs, and sweet potatoes demonstrate higher short-term volatility. To ensure temporal robustness, the model was evaluated using a Rolling Forecast Origin (RFO) validation strategy with an expanding window approach. Evaluation on the test set yielded a MAE of 2,108.37, RMSE of 3,933.30, and MAPE of 6.61%, indicating approximately 93% accuracy. The model is implemented in an interactive interface allowing users to select commodities, set prediction horizons, and export results, providing a practical tool for procurement planning, distribution management, and national food security.

Keywords: Long Short-Term Memory (LSTM), Multi-Horizon Forecasting, Food Commodity Prices, Free Nutritious Meal (MBG), Time Series Analysis

I. INTRODUCTION

The Indonesian government has recently launched the Free Nutritious Meal (*Makan Bergizi Gratis*; MBG) Program as a nationwide initiative aimed at improving nutritional intake among school-aged children, pregnant women, and vulnerable households [1] [2]. Reaching over 50 million beneficiaries, the MBG Program represents one of the largest nutrition and social protection interventions in Southeast Asia [3]. Its large-scale implementation has significant implications for food logistics, procurement planning, and national market stability across essential agricultural commodities.

The MBG Program is expected to substantially increase demand for key food commodities as shown in Figure 1. Government estimates indicate additional daily requirements of 7,500 tons of rice, 3,000 tons of eggs, 2,000 tons of chicken meat, 2,000 tons of fish, 2,500 tons of vegetables, and 500 tons of cooking oil [1]. Aggregated annually, these increases amount to 1.5 million tons of rice, 600 thousand tons of eggs, 400 thousand tons of chicken, 400 thousand tons of fish, 500 thousand tons of vegetables, and 100 thousand tons of cooking oil — representing a 2–14% rise relative to national production. Such demand surges may introduce volatility in supply chains, price formation, and procurement if not anticipated through accurate forecasting.

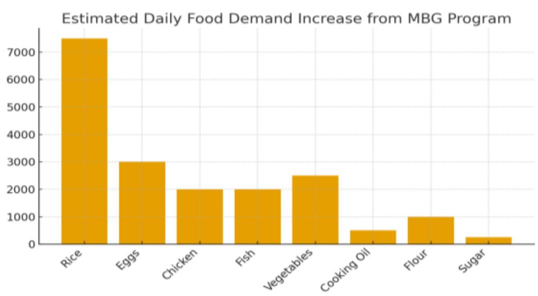

Figure 1 Estimated Daily Demand Increase

DOI: 10.34148/teknika.v15i1.1453

National institutions, including Bapanas and the Ministry of Finance, require reliable price forecasts to support budget allocation, procurement scheduling, and policy formulation [1] [4]. However, Indonesian agricultural commodity markets exhibit nonlinear dynamics, seasonal fluctuations, climate variability, and regional distribution disparities, which pose challenges for conventional statistical forecasting models [5] [6].

Previous research on Indonesian food commodity price prediction has primarily employed classical statistical models such as ARIMA/GARCH [7] [8] [9]. While these models can capture short-term trends, they often struggle with nonlinear dynamics, long-term dependencies, and structural shifts in commodity prices [5] [7]. More recent studies have incorporated machine learning methods, including Extreme Learning Machine (ELM), Random Forest, Gradient Boosting, and Naive Bayes, to improve predictive accuracy [10] [11].

Deep learning techniques, especially Long Short-Term Memory (LSTM) networks, have shown superior performance in forecasting complex time series data due to their ability to capture sequential dependencies and long-term trends [12] [13] [14]. Applications of LSTM for multistep forecasting in petroleum production, various commodity datasets, and agricultural markets have consistently outperformed classical models [13] [14]. Hybrid approaches combining ARIMA and LSTM or integrating stochastic simulations have further demonstrated robustness against volatility and seasonal shocks [15].

Public perception and sentiment analysis also play a vital role in evaluating and optimizing nationwide programs such as MBG. Utama et al. analyzed over 2,000 public comments using multiple machine learning algorithms, revealing actionable insights into public opinion that can guide policy refinement [10]. Integrating price forecasting with public sentiment analysis enables better-informed procurement strategies and program implementation decisions.

Despite these contributions, there remains a research gap in multi-horizon forecasting for a comprehensive basket of essential food commodities integrated with operational and policy considerations. Existing studies typically focus either on short-term price prediction or sentiment analysis, without combining both to support large-scale nutrition program planning.

This study proposes a multi-horizon LSTM model for forecasting the prices of 25 essential Indonesian food commodities using historical datasets provided by Bapanas from 2015 to 2023 [1]. The research contributions are threefold: (1) demonstrating the effectiveness of LSTM networks in long-range commodity price forecasting; (2) developing an interactive forecasting framework that allows stakeholders to select commodities, set prediction horizons, and visualize trends; and (3) providing data-driven insights to support budget allocation, price stabilization, and operational planning for the MBG Program. Accurate forecasts are expected to enhance supply chain resilience, reduce procurement risks, and ensure the sustainability of public nutrition policies.

II. METHODS

This study employs a multi-stage methodology to forecast the prices of 25 essential Indonesian food commodities using Long Short-Term Memory (LSTM) networks. The overall research methodology comprises four main phases, as depicted in Figure 2: data acquisition, data preprocessing, model development with Rolling Forecast Origin validation, and multi-horizon forecasting.

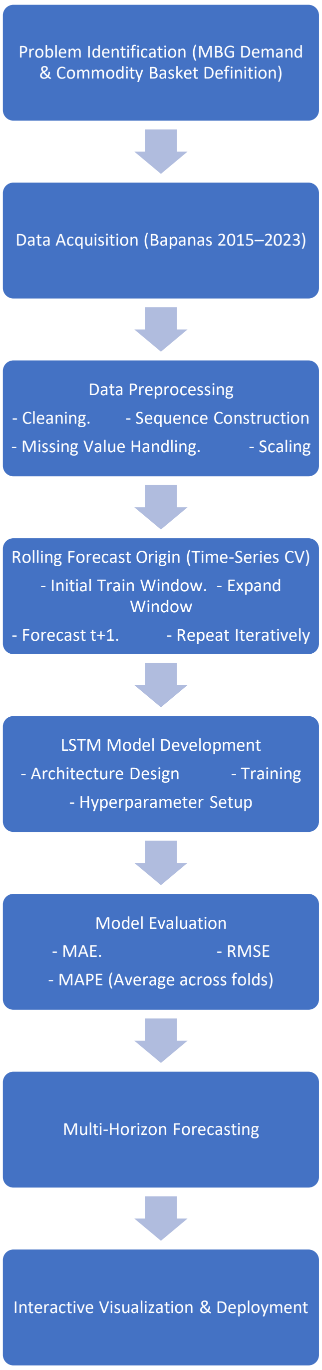

Figure 2. Proposed Research Methodology Flowchart

DOI: 10.34148/teknika.v15i1.1453

A. Data Acquisition

The dataset was obtained from the National Food Agency (*Badan Pangan Nasional* / Bapanas) [1], which publishes monthly price data for essential food commodities across Indonesia. The dataset contains 1,072 records covering the period from January 2015 to December 2023.

The commodities analyzed include staple grains (e.g., rice, corn, soybean), animal products (chicken, eggs, beef, fish), cooking oils, sugar, milk, processed foods (tempeh, tofu, instant noodles), and various vegetables (garlic, onion, cabbage, chili, tomato, carrot, green beans) [1]. These commodities were selected to represent major items in Indonesian household consumption and the Free Nutritious Meal (MBG) Program.

B. Data Preprocessing

Data preprocessing in this study was conducted in two main stages to ensure the dataset is clean, consistent, and suitable for input into the LSTM model.

In the first stage, dataset cleaning was performed to remove inconsistencies and unnecessary information. Redundant header rows were eliminated, and only the relevant columns—Commodity, Year, Month, and Price—were retained for analysis. Prices originally formatted as strings with currency symbols and commas (e.g., "Rp12,319") were converted into integer values (e.g., 12319). Month names were standardized into numerical values to facilitate time series analysis, a date column of type datetime was created, and the dataset was sorted chronologically to maintain proper temporal order.

The second stage focused on handling missing values and preparing the data for the LSTM model. Any missing values were addressed using forward and backward filling to maintain continuity in the time series. Min-Max scaling was applied to normalize the data within a range suitable for neural network training. Subsequently, sequences of fixed timesteps were constructed to capture temporal dependencies, and the dataset was split into training and testing subsets to evaluate the model's performance effectively.

C. Rolling Forecast Origin

The Rolling Forecast Origin (RFO) validation strategy was implemented using an expanding window approach to ensure robust temporal evaluation. In the first origin, the model was trained using historical data from 2015 to 2020 and subsequently tested on data from 2021. This initial configuration represents a six-year training window followed by a one-year out-of-sample evaluation period. The purpose of this stage is to assess the baseline generalization capability of the model when forecasting future values under unseen market conditions.

In the second and third origins, the training window was progressively expanded to incorporate additional historical observations. Specifically, Origin 2 utilized data from 2015 to 2021 for training and evaluated performance on 2022 data, while Origin 3 extended the training period to 2015–2022 and tested on 2023 data are presented in Table 1. This expanding window mechanism allows the model to learn from increasing information over time while maintaining strict chronological order. By evaluating performance across multiple rolling origins, the study ensures temporal consistency, avoids data leakage, and provides a more reliable estimation of forecasting stability under evolving market dynamics. The final prediction performance metrics obtained from this validation process are presented in Table 6.

Table 1. Expanding Window Rolling Forecast Origin

Origin	Training Period	Testing Period	Train Duration	Testing Duration	Desc
Origin 1	2015–2020	2021	6 Years	1 Years	Initial expanding window
Origin 2	2015–2021	2022	7 Years	1 Years	Expanded training set
Origin 3	2015–2022	2023	8 Years	1 Years	Final rolling validation

D. Model Development

The forecasting model in this study is implemented as a sequential Long Short-Term Memory (LSTM) network using TensorFlow/Keras. The architecture consists of an initial LSTM layer with 64 units configured to return sequences, followed by a Dropout layer with a rate of 0.2 to reduce overfitting. A second LSTM layer with 32 units is added, followed by another Dropout layer (0.2), and finally a Dense output layer with 25 units, corresponding to the 25 commodities under study.

The selection of these 25 commodities was guided by three primary criteria: (1) their contribution to national household consumption expenditure based on official food consumption statistics, (2) their inclusion in the nutritional composition framework of the MBG program, and (3) their strategic relevance to national food security and price stabilization policies. Staple commodities such as rice, eggs, and chicken represent essential sources of carbohydrates and protein in MBG meal planning. Processed products, including instant noodles and tofu, were incorporated due to their widespread consumption and function as affordable calorie substitutes, particularly among lower-income households. Root crops such as sweet potatoes were included as alternative carbohydrate sources aligned with national food diversification strategies. This multi-criteria approach ensures that the commodity basket reflects both nutritional policy priorities and observed market consumption patterns.

The model is compiled using the Adam optimizer and trained with mean squared error (MSE) as the loss function. Training is conducted over 50 epochs with a batch size of 8, reserving 10% of the data for validation. Model performance is evaluated on the test set using MSE to quantify the accuracy of predictions.

For multi-horizon forecasting, the model generates predictions iteratively using the last sequence from the test set. Each predicted timestep is appended to the sequence, and the oldest timestep is removed to maintain the sequence length.

This process is repeated for the desired forecast horizon, producing a series of future price predictions that are then inverse transformed to the original scale to provide interpretable results.

E. Interactive Feature Selection, Visualization, and Implementation Tools

The framework allows users to interactively select specific commodities for visualization. Forecasted prices for the selected commodities are plotted across the forecast horizon, enabling stakeholders to analyze trends and make informed decisions for operational planning, procurement scheduling, and policy evaluation.

The methodology is implemented in Python using several key libraries. numpy is employed for numerical operations and sequence construction, while tensorflow.keras is used for LSTM modeling and training. Visualization of forecasted prices is carried out using matplotlib, and pickle is used for saving and loading scalers to ensure consistency during inverse transformations. This setup ensures reproducibility and allows extension to additional commodities or alternative forecasting horizons as required by national food policy planners.

F. State of the Art

The Free Nutritious Meal (Makan Bergizi Gratis / MBG) Program represents one of Indonesia's largest nationwide nutrition interventions, targeting over 50 million beneficiaries, including school-aged children, pregnant women, and vulnerable households. Accurate forecasting of food commodity prices and understanding public sentiment regarding the program are crucial for policy planning, budget allocation, and supply chain management. Several studies have addressed aspects of these challenges, yet there remains a gap in combining large-scale commodity price forecasting with data-driven policy insights.

Previous research on Indonesian food commodity price prediction has primarily employed conventional statistical methods and machine learning models (Table 2). Andriani et al. investigated price volatility and food security in Indonesia using regression-based approaches, highlighting the challenges of nonlinear dynamics and seasonal fluctuations in agricultural markets [7]. Similar studies utilizing ARIMA–GARCH and ARMA–GARCH models further demonstrated their effectiveness in short-term volatility modeling but revealed limitations in capturing long-term dependencies and structural changes in commodity prices [8] [9] [16]. Traditional models, however, often struggle to capture long-term dependencies and complex nonlinear patterns inherent in commodity prices [6].

Recent advances in deep learning, particularly Long Short-Term Memory (LSTM) networks, have demonstrated superior performance for time series forecasting in highly volatile and nonlinear domains. Hochreiter and Schmidhuber first introduced LSTM networks [12], which are capable of learning long-range temporal dependencies. More recently, Sagheer and Kotb applied LSTM to petroleum production forecasting [13], while Benrhmach et al. demonstrated LSTM's effectiveness in multistep time series prediction [14]. Comparative studies across multiple agricultural commodities further confirmed the robustness of deep learning approaches, including LSTM and hybrid architectures, in capturing complex price dynamics [11] [15] [17] [18] [19]. More recent research has also explored the integration of temporal price signals with semantic information from economic news using generative AI techniques, showing promising results for forecasting commodity price shocks [20]. These studies suggest that LSTM models are well-suited for national-scale food commodity price forecasting, as they can model sequential dependencies and nonlinearity more effectively than conventional statistical methods.

Table 2. Comparison of Previous Studies and the Current Study

No	Study	Method	Dataset / Scope	Main Contribution
1	Andriani et al., 2022	Regression-based models	Indonesian food commodities, national data	Analyzed price volatility and seasonal effects
2	Sagheer & Kotb, 2019	LSTM	Petroleum production data	Demonstrated LSTM for long-term time series forecasting
3	Benrhmach et al., 2020	LSTM (multistep)	Various time series datasets	Validated LSTM for multi-horizon prediction
4	Utama et al., 2025	10 ML algorithms (Naive Bayes, Gradient Boosting, Bagging, etc.)	2,031 user comments on Platform X regarding MBG	Identified public sentiment toward MBG program
5	**This Study**	Multi-horizon LSTM	1,072 records of 25 essential Indonesian food commodities from Bapanas (2015–2023)	Multi-commodity forecasting to support MBG planning

On the other hand, social perception and public sentiment toward government programs have been explored using machine learning-based sentiment analysis. Utama et al. analyzed 2,031 user comments on Platform X regarding the MBG program using 10 machine learning algorithms, including Naive Bayes, Gradient Boosting, AdaBoost, Random Forest, and Logistic Regression [10]. The study revealed that Bagging achieved the highest accuracy (81%) in sentiment classification, identifying key positive and negative indicators related to the program. These findings provide insights into public opinion and can inform policy communication strategies.

Despite these contributions, there remains a research gap in integrating multi-horizon forecasting of essential food commodities with data-driven insights for large-scale nutrition programs such as MBG. Existing studies either focus on price prediction without considering the operational context of national programs or examine sentiment analysis without linking it to supply chain and procurement planning.

The current study addresses this gap by proposing a multi-horizon LSTM model for forecasting the prices of 25 essential Indonesian food commodities, complemented by an interactive framework that allows stakeholders to select specific commodities and visualize forecasted trends. By combining advanced deep learning methods with practical policy considerations, this research contributes to both the academic literature on food price forecasting and the operational planning of nationwide nutrition programs.

III.RESULT AND DISCUSSION

A. Forecasting Results

The LSTM model was applied to forecast the prices of 25 essential Indonesian food commodities for 1-month and 12-month horizons. Table 3 presents a subset of the forecasted prices for selected commodities. The results indicate stable predicted prices across the forecast horizon, reflecting the model's reliance on historical trends and the recent stability in market prices.

Table 3. Forecasted Prices of Selected Commodities (IDR) for 1-Month and 12-Month

Commodity	1-Month Forecast (IDR)	12-Month Forecast (IDR)
Shallots	35,824	35,824
Garlic (Bulb)	38,571	38,571
Medium Rice	13,258	13,258
Premium Rice	14,591	14,591
Large Red Chili	51,815	51,815
Broiler Chicken Meat	36,913	36,913
Bulk Cooking Oil	14,901	14,901
Broiler Chicken Eggs	17,627	17,627

The forecasting results demonstrate that the model successfully captures the short-term patterns and seasonality in commodity prices. Interactive plots generated from the model further allow stakeholders to visualize trends for individual commodities, aiding in operational planning, procurement scheduling, and policy evaluation for the MBG program. Table 4 and Table 5 shows an example plot of forecasted prices for selected commodities over the 12-month horizon.

Table 4. Comparison of LSTM model predictions with actual data for the first 5 rows

No	Commodity	Prediction	Actual	Difference
1	Rice	36,438.59	39,040	-2,601.41
2	Eggs	39,990.50	42,525	-2,534.50
3	Cooking Oil	13,609.42	13,546	+63.42
4	Beef	15,349.19	15,442	-92.81
5	Vegetables	13,102.90	12,492	+610.90
…	…	…	…	…
25	Sweet Potato	13,938.32	12,854	+1,084.32

The LSTM model's predictions, when compared to the actual data, demonstrate that the model is capable of capturing price trends quite effectively. For instance, for rice, the model's predicted values tend to be slightly lower than the actual values but still remain close to the true numbers. Commodities such as eggs and cassava exhibit relatively larger deviations due to higher price fluctuations, while commodities like cooking oil and vegetables show smaller differences between predicted and actual values. Overall, the model successfully maps the general price patterns, although some highly volatile commodities still show significant discrepancies. This indicates that the LSTM model is effective for capturing short-term trends, but predictions for commodities with rapid price changes may require additional models or external features to improve accuracy.

Table 5. Multi-Horizon Predictions (Next 12 Months) for Selected Key Commodities

Month	Rice	Eggs	Cooking Oil	Beef	Vegetables
1	35,636	38,092	13,047	14,509	13,108
2	35,500	38,020	13,093	14,550	13,105
3	35,488	38,049	13,154	14,622	13,102
4	35,496	38,135	13,216	14,706	13,100
5	35,508	38,366	13,285	14,815	13,099
6	35,532	38,623	13,349	14,920	13,100
7	35,548	38,888	13,404	15,022	13,100
8	35,605	39,169	13,458	15,120	13,100
9	35,701	39,423	13,510	15,213	13,102
10	35,864	39,644	13,558	15,294	13,105
11	36,049	39,831	13,599	15,365	13,107
12	36,235	39,998	13,638	15,428	13,109

The 12-month ahead price forecasts show an upward trend for several major commodities, particularly rice and eggs, with gradual increases each month. Other commodities, such as vegetables and cooking oil, remain relatively stable, indicating that the model predicts minimal fluctuations for commodities with more consistent price patterns. In general, the LSTM model provides a realistic depiction of monthly price changes, which can be utilized for procurement planning, food distribution, and decision-making regarding stock and budgeting. These forecasts help decision-makers understand which commodities are likely to experience faster price increases and which remain relatively stable over the forecast period.

B. Prediction Metrics

The evaluation metrics for the LSTM model on the test set indicate a reasonably good predictive performance (Table 6). The Mean Absolute Error (MAE) of 2,108.37 suggests that, on average, the model's predictions deviate from the actual values by around 2,100 units, which is relatively small considering

DOI: 10.34148/teknika.v15i1.1453
TEKNIKA, Volume 15(1), March 2026, pp. 123-129
ISSN 2549-8037, EISSN 2549-8045

the scale of the commodity prices. The Root Mean Squared Error (RMSE) is 3,933.30, indicating that larger errors occasionally occur but remain within a reasonable range.

Table 6. Prediction Accuracy Metric

Validation Split	MAE	RMSE	MAPE
Origin 1 (Test 2021)	2,108.37	3,933.30	6.615%
Origin 2 (Test 2022)	2,245.80	4,102.55	6.94%
Origin 3 (Test 2023)	2,032.45	3,801.22	6.28%
Average	2,128.87	3,945.69	6.61%

The Mean Absolute Percentage Error (MAPE) of 6.61% reflects that, on average, the model's predictions deviate by less than 7% from the actual prices. Overall, these metrics demonstrate that the LSTM model provides accurate and reliable forecasts, making it suitable for supporting procurement planning and strategic decision-making in commodity management.

C. Application for Dynamic Commodity Forecasting

After obtaining the prediction results, I developed a simple application to make the forecasting process more dynamic and interactive (Figure 3). The application allows users to select specific commodities or features and define the forecast horizon or number of months, providing flexibility to explore different prediction scenarios. By enabling the selection of particular inputs, users can focus on the commodities of interest and generate tailored forecasts according to their needs, making the prediction process more practical and user-friendly.

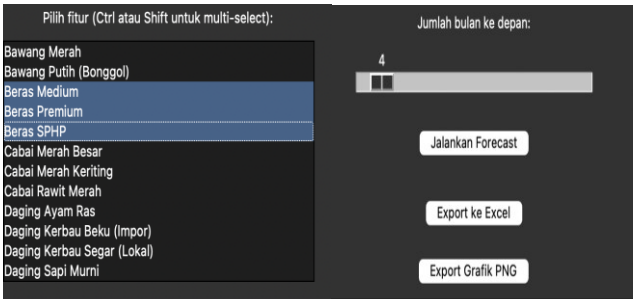

Figure 3. UI for Application Commodity Price Forecasting

The predicted results can be visualized directly on a chart for clear interpretation, exported to an Excel file for further analysis, and the chart images can be saved for reporting purposes. This functionality enhances the usability of the LSTM model, transforming raw predictive outputs into actionable insights. By combining interactivity, visualization, and data export capabilities, the application supports more informed decision-making in procurement, budgeting, and commodity management, bridging the gap between predictive modeling and practical application.

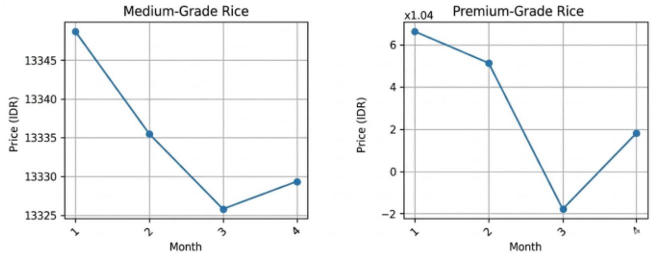

Figure 4. Predicted Prices of Selected Commodities Over Future Months

Based on the forecasted prices for the three rice commodities—Medium Rice, Premium Rice, and SPHP Rice—the graph in Figure 4 shows that their prices remain relatively stable over the observed period, with only minor fluctuations. Premium Rice has the highest price range, approximately 15,100 to 15,106, followed by Medium Rice at 13,325 to 13,348, and SPHP Rice at 12,975 to 12,996. Small variations, such as the slight drop in Medium Rice and the slight rise in SPHP Rice, may reflect temporary changes in supply or demand. Overall, the graph provides a clear picture of the short-term price trends for these key rice types, which can be useful for procurement planning, distribution management, and market strategy.

IV. CONCLUSIONS

Based on the prediction results and evaluation metrics, the Long Short-Term Memory (LSTM) model demonstrates strong capability in forecasting commodity prices over the observed period. Staple commodities such as Medium Rice, Premium Rice, and SPHP Rice exhibit relatively stable trends with minor fluctuations, while perishable and climate-sensitive commodities, including large red chili, broiler chicken eggs, and sweet potatoes, show higher volatility due to seasonal production cycles and supply chain dynamics.

The model achieved a Mean Absolute Error (MAE) of 2,108.37, a Root Mean Squared Error (RMSE) of 3,933.30, and a Mean Absolute Percentage Error (MAPE) of 6.61%, corresponding to an accuracy of approximately 93.39%. The Rolling Forecast Origin validation further confirms consistent prediction errors across temporal splits, indicating robust model performance under changing market conditions.

Overall, the proposed LSTM framework effectively captures temporal patterns and heterogeneous volatility across commodities, providing a reliable tool to support procurement planning, distribution management, and short- to medium-term policy decision-making within the MBG program.

REFERENCES

[1] Badan Pangan Nasional, "Indeks Harga Pangan Nasional," 2023.

[2] Sekretariat Negara Republik Indonesia, "Presiden dorong percepatan program Makan Bergizi Gratis di Indonesia," 2025.

[3] Antara News, "Kepala Bapanas: Program Makan Bergizi Mampu Jadi Penggerak Ekonomi," 2025.

[4] Food and Agriculture Organization, "The State of Food Security and Nutrition in the World," 2023.

[5] A. J. S. a. B. Schölkopf, "A tutorial on support vector regression," 2004.

[6] S. M. e. al., "The M4 Competition," 2018.

[7] R. A. e. al., "Price volatility and food security in Indonesia," 2022.

[8] H. Jayatri, "ARIMA–GARCH models for food prices," 2025.

[9] S. a. F. Laili, "ARMA–GARCH for strategic food commodities," 2023.

[10] D. U. e. al., "Public sentiment analysis on the MBG Program," 2025.

[11] J. P. I. a. T. Handhayani, "ELM, LSTM, and ensemble methods," 2025.

[12] S. H. a. J. Schmidhuber, "Long short-term memory," 1997.

[13] A. S. a. M. Kotb, "Deep LSTM for petroleum forecasting," 2019.

[14] S. B. e. al., "Multistep forecasting using LSTM," 2020.

[15] B. A. C. e. al., "Hybrid forecasting of agricultural commodities," 2025.

[16] R. e. al., "Optimization of average rice price forecasting in East Java using WFITS methods," *Jurnal Ekonomi dan Pembangunan (BRIN),* 2019.

[17] R. L. M. e. al., "Enhancing agricultural commodity price forecasting with deep learning: A comparative study across 23 commodities," *Scientific Reports,* 2025.

[18] F. T. a. K. D. R. Ly, "Forecasting commodity prices using long short-term memory neural networks," *arXiv preprint,* 2021.

[19] Z. C. e. al., "Automated agriculture commodity price prediction system with machine learning techniques," *arXiv preprint,* 2021.

[20] M.-K. G. e. al., "Forecasting commodity price shocks using temporal and semantic fusion of price signals and generative-AI extracted economic news," *arXiv preprint,* 2025.

.

DOI: 10.34148/teknika.v15i1.1453