Journal of Artificial Intelligence and Engineering Applications

Website: https://ioinformatic.org/

June 2026. Vol. 5. No. 3; e-ISSN: 2808-4519

Comparison of Xgboost, Lstm, and Neural Prophet Models for Red Cayenne Pepper Price Prediction in East Java

Hafid Alfa Anamsyah[1]*, I Gede Susrama Mas Diyasa[2], Andreas Nugroho Sihananto[3]

[1,2,3]*Informatics, UPN "Veteran" Jawa Timur*
20081010189@student.upnjatim.ac.id[1*], igsusrama.if@upnjatim.ac.id[2]

Abstract

The price of red cayenne pepper in Indonesia, particularly in East Java Province, frequently experiences significant fluctuations, affecting food inflation and economic stability. Accurate price forecasting is therefore essential to support decision-making in food supply chain management and price stabilization policies. This study compares the forecasting performance of three models, namely XGBoost, LSTM, and Neural Prophet, using historical price data from September 2024 to August 2025 obtained from the official Siskaperbapo website of the East Java Provincial Department of Industry and Trade (DISPERINDAG). A quantitative time series forecasting approach was employed, and model performance was evaluated using Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). The results indicate that XGBoost achieved the highest prediction accuracy, with an RMSE of 0.52, MAE of 0.40, and MAPE of 1.42%, outperforming LSTM and Neural Prophet. The findings also show that model selection, dataset length, and training-test split proportion significantly influence forecasting performance, with longer datasets and larger training sets generally improving prediction accuracy. Overall, XGBoost proved to be the most accurate and stable model for forecasting red cayenne pepper prices, providing valuable support for AI-based food price prediction and agricultural policy decision-making.

***Keywords**: Price prediction, Red cayenne pepper, XGBoost, LSTM, Neural Prophet, Time series forecasting*

1. Introduction

Red chili peppers are a strategic horticultural commodity in Indonesia, playing a vital role in people's daily lives. Not only are they a staple ingredient in various dishes, but they also directly impact economic stability, particularly in the food expenditure group. High chili consumption makes price fluctuations highly sensitive and have a broad impact on public welfare.

The phenomenon of price fluctuations in red chili peppers has long been a national issue, particularly in production centers like East Java. As one of the largest chili-producing provinces in Indonesia, East Java frequently faces unstable price dynamics. Spikes in red chili pepper prices can drive food inflation and suppress purchasing power, while drastic price drops harm farmers because selling prices are not commensurate with production costs [1].

According to data from the Central Statistics Agency (BPS), food inflation in Indonesia is often triggered by rising prices of red chili peppers, particularly during the rainy season and when distribution disruptions occur. This situation demonstrates the close relationship between chili prices and national inflation stability. In fact, in 2020–2022, the price of cayenne pepper spiked by more than 100% in a short period of time due to extreme weather and inefficient distribution [2].

The price of red cayenne pepper in East Java exhibits a fairly sharp seasonal fluctuation pattern. When production is abundant, prices plummet, resulting in losses for farmers. Conversely, when supply is low, prices can soar to over IDR 100,000 per kilogram at the consumer level. This situation creates economic uncertainty for both consumers and producers, necessitating an accurate price prediction system to anticipate market changes [3].

The XGBoost, LSTM, and Neural Prophet methods are widely used approaches in time series data forecasting due to their ability to capture nonlinear and complex patterns. XGBoost, as an implementation of extreme gradient boosting, offers advantages in computational efficiency, regularization capabilities, and stable prediction performance across various types of structured data. Research by B. W. Sari and D. Prabowo in Intellect: Indonesian Journal of Learning and Technological Innovation shows that XGBoost produces better accuracy than Random Forest and Gradient Boosting in house price prediction, confirming its strength in handling both numerical and categorical variables simultaneously [4]. This advantage makes XGBoost relevant for research based on tabular and time series data that has undergone feature engineering.

However, XGBoost has limitations in directly modeling long-term temporal dependencies because it is not inherently a sequential model. Its reliance on feature engineering means that its performance is heavily influenced by the quality of the lag features and derived variables created by the researchers. Furthermore, model interpretation can become complex as the number of trees and model depth increase. This differs from recurrent neural network-based approaches such as LSTM, which are inherently designed to capture temporal dynamics.

LSTM, as an extension of Recurrent Neural Networks, has a memory mechanism through cell states and gates that can retain long-term information. A study by A. F. Alkayes and T. Sugihartono in the Indonesian Journal of Education and Technology showed that LSTM can provide competitive, even superior, performance in predicting Tesla stock prices compared to XGBoost on highly fluctuating data [5]. The main advantage of LSTM lies in its ability to capture seasonal patterns and complex trends without the need for extensive feature engineering. However, this model requires large amounts of data, takes longer to train, and is susceptible to overfitting if not accompanied by adequate regularization techniques.

Neural Prophet is an extension of the Prophet model, based on a decomposable time series model combined with neural network components to increase nonlinear flexibility. Research by A. Primawati, F. A. Mustika, and Y. Wibawanti in Simetris: Jurnal Teknik Mesin, Elektro dan Ilmu Komputer (Simetris: Jurnal Teknik Mesin, Elektro dan Ilmu Komputer) shows that the Prophet and LSTM approaches have different performance characteristics in predicting gold prices, with the decomposition-based model excelling in capturing clear trends and seasonality [6]. Neural Prophet extends these capabilities by incorporating autoregression and neural components, making it more adaptive to nonlinear patterns. Its advantage lies in the interpretability of the trend and seasonality components, but its performance may decrease on data with extreme volatility or high noise.

Based on these characteristics, the research hypothesis can be formulated that there are significant differences in predictive performance between XGBoost, LSTM, and Neural Prophet on specific time series data. The first hypothesis states that LSTM will provide higher accuracy on data with long-term dependencies and complex fluctuations. The second hypothesis states that XGBoost will excel on data with structured features and nonlinear patterns that can be represented through feature engineering. The third hypothesis states that Neural Prophet will demonstrate stable performance on data with strong trends and seasonality and relatively structured data.

The selection of these three methods was based on theoretical and empirical considerations. Theoretically, XGBoost represents a boosting-based ensemble approach, LSTM represents a sequential deep learning architecture, and Neural Prophet combines statistical decomposition with neural learning. Empirically, all three have competitive performance in the context of price and economic time series prediction. By comparing these three approaches within a systematic experimental framework, the research can contribute comprehensively to identifying the most optimal method based on specific data characteristics and prediction objectives.

Unfortunately, there is still very little empirical research in Indonesia that directly compares these three models in the context of red cayenne pepper price prediction, particularly in East Java. The majority of previous research has focused on conventional methods or on commodities other than red cayenne pepper. This gap indicates a research gap that needs to be filled to make chili price modeling in Indonesia more comprehensive and accurate [7].

In addition to methodological factors, another challenge is the availability of sufficiently long and clean historical data. Daily price fluctuations available from the East Java Department of Industry and Trade (DISPERINDAG) from September 2024 to August 2025 provide an opportunity to build a predictive model with a high level of accuracy. With this data, the study can test the performance of various algorithms and determine the best model for red cayenne pepper in East Java.

The urgency of this research is also supported by the fact that spikes in red cayenne pepper prices have a significant impact on household consumption and government policy. Fauzi and Andriani (2023) emphasized that rising chili prices can drastically affect supply and demand, thus triggering market imbalances [8]. Therefore, accurate price prediction is not only academic but also practical in supporting decision-making.

Based on the description above, this study aims to compare the performance of three modern models (XGBoost, LSTM, and Neural Prophet) in predicting the price of red chili peppers in East Java. This research is expected to provide new contributions to the field of food commodity price forecasting, addressing the limitations of previous studies that predominantly used conventional methods, and providing practical benefits for farmers, traders, consumers, and the government in anticipating price dynamics.

2. Literature Review

2.1. XGBoost (Extreme Gradient Boosting)

Extreme Gradient Boosting (XGBoost) is a supervised machine learning algorithm based on gradient boosting decision trees developed by Chen and Guestrin [10]. XGBoost is widely recognized for its computational efficiency, scalability, and predictive performance. The model improves prediction accuracy by iteratively combining weak learners into a stronger ensemble model while minimizing prediction errors through gradient optimization.

XGBoost possesses several advantages in time series forecasting. First, the algorithm effectively handles nonlinear relationships and high-dimensional datasets. Second, regularization mechanisms reduce overfitting risks, which commonly occur in volatile data. Third, XGBoost can process missing values and noisy data efficiently. These characteristics make XGBoost suitable for forecasting commodity prices characterized by irregular fluctuations.

However, XGBoost also has limitations. The algorithm does not inherently model sequential temporal dependencies because it primarily relies on engineered lag features rather than internal memory mechanisms. As a result, the forecasting performance highly depends on feature engineering quality. In highly sequential and long-term temporal datasets, deep learning approaches may outperform tree-based models because they can directly learn temporal patterns from sequential inputs.

Several studies demonstrated the effectiveness of XGBoost in financial and commodity forecasting tasks. Ulya et al. compared XGBoost and LSTM for Bitcoin price prediction and reported that XGBoost achieved stable forecasting performance in nonlinear datasets involving sentiment and trend variables [11]. Similarly, Yuan et al. emphasized that gradient boosting models remain competitive in volatile energy price forecasting due to their robustness and interpretability [12].

2.2. LSTM (Long Short-Term Memory)

Long Short-Term Memory (LSTM) is a type of Recurrent Neural Network (RNN) introduced by Hochreiter and Schmidhuber to overcome the vanishing gradient problem in sequential learning [13]. Unlike conventional neural networks, LSTM contains memory cells and gating mechanisms that enable the model to retain long-term dependencies within sequential data.

LSTM consists of three main gates, namely the input gate, forget gate, and output gate. These components regulate information flow within the network, allowing the model to preserve relevant historical information while discarding irrelevant patterns. Such capability makes LSTM highly effective for time series forecasting involving long-range temporal dependencies.

In commodity price prediction, LSTM demonstrates superior performance because agricultural prices often contain seasonal cycles and complex nonlinear interactions. The model can capture hidden sequential relationships without requiring extensive manual feature engineering. Moreover, LSTM adapts effectively to dynamic fluctuations and irregular market movements.

Nevertheless, LSTM also presents several weaknesses. The model requires substantial computational resources and large datasets for optimal performance. Training complexity and hyperparameter tuning can also increase implementation difficulty. In datasets with limited observations, LSTM may suffer from overfitting if regularization strategies are not properly implemented.

Research conducted by Mehtab et al. showed that LSTM-based forecasting models produced higher predictive accuracy compared to conventional regression models in stock market prediction tasks [14]. Similarly, studies on commodity and metal price forecasting revealed that LSTM effectively captures nonlinear temporal structures and long-term dependencies within volatile datasets [15].

2.3. Neural Prophet

Neural Prophet is a modern forecasting framework developed as an extension of Facebook Prophet by integrating neural network components into traditional additive forecasting models [16]. The framework combines trend decomposition, seasonality modeling, autoregression, and neural network learning to improve forecasting flexibility and interpretability.

Neural Prophet maintains the explainability advantages of Prophet while introducing deep learning-based enhancements. The model can capture local temporal contexts using autoregressive neural components and external regressors. Furthermore, Neural Prophet supports trend change detection, multiple seasonality patterns, and missing data handling.

Compared with traditional deep learning models, Neural Prophet provides easier implementation and more interpretable forecasting outputs. This characteristic is beneficial in economic and agricultural forecasting contexts where transparency is important for policy analysis. However, Neural Prophet may underperform in highly volatile datasets containing abrupt fluctuations because additive decomposition assumptions can limit flexibility in capturing extreme nonlinear dynamics.

Triebe et al. reported that Neural Prophet outperformed Prophet in various real-world forecasting datasets while maintaining interpretability and scalability [16]. The framework achieved significant accuracy improvements for short- and medium-term forecasting tasks, making it increasingly relevant for economic and commodity forecasting applications.

3. Research Methods

3.1. Flowchart

Figure 1 below shows the flowchart of the research conducted, starting from the literature review and data collection stages to data predictions from the most optimal model.

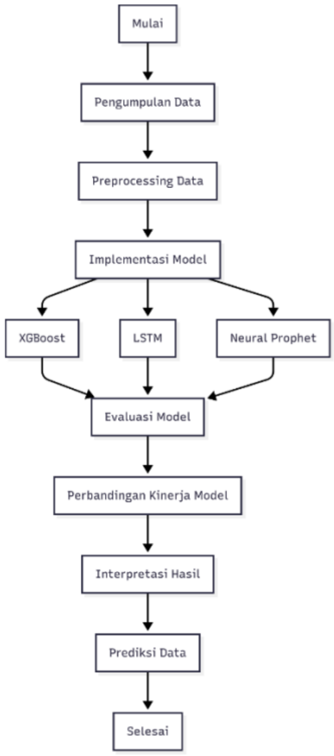

Fig. 1: Research Flowchart

Figure 1 explains the systematic research flowchart as follows:

1. Data Collection
   This stage collects daily red chili price data obtained from siskaperbapo.jatimprov.go.id.
2. Data Preprocessing
   This stage includes data cleaning (handling missing values), normalization, and dividing the data into training and testing sets.
3. Model Implementation
   This stage builds three predictive models (XGBoost, LSTM, Neural Prophet) with appropriate parameters.
4. Model Evaluation
   This stage uses RMSE, MAE, and MAPE metrics.
5. Results Comparison
   This stage identifies the best-performing model.
6. Data Prediction
   This stage predicts or forecasts data from the most optimal model.

3.2. Testing Scenario

The testing scenario is the research step in designing a test scenario for a prediction model. The variable used in this study is the price of red chili peppers from September 2024 to August 2025. Each model (XGBoost, LSTM, and Neural Prophet) will be tested on the red chili pepper dataset for 3 months (June - August 2025), 6 months (March - August 2025), and 1 year (September 2024 - August 2025), with train:test data splits of 60:40, 70:30, and 80:20, respectively.

4. Research Results

4.1. Model Implementation Results

4.1.1. XGBoost Model Results

The XGBoost model was trained using the following parameters: n_estimators=500, learning_rate=0.05, max_depth=5, subsample=0.8, colsample_bytree=0.8, and random_state=42. This model demonstrated excellent ability to handle short-term nonlinear patterns, but had limitations in capturing long-term dependencies.

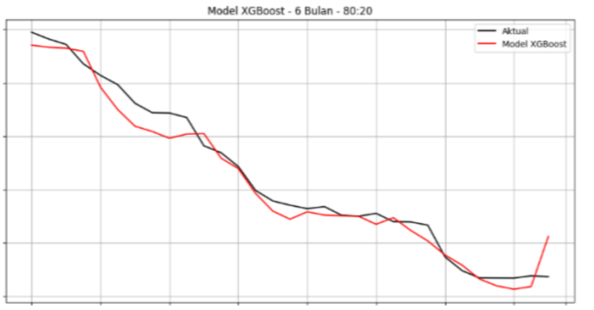

Fig. 2: XGBoost model implementation with the best evaluation

Table 1: XGBoost best evaluation model

Metric	Value
RMSE	0.52
MAE	0.4
MAPE	1.42%

Fig 2 and Table 1 show that the implementation of the XGBoost model on the 6-month dataset with a 60:40 data split yielded an RMSE of 2.65, an MAE of 1.87, and a MAPE of 3.91%. These results indicate that the medium dataset produc-es a more optimal model, and the small training data split makes the model unable to capture patterns.

4.1.2. LSTM Model Results

The LSTM model uses an architecture with 3 hidden layers and 64 neuron units, a tanh activation function, and an adam optimizer. The model was trained for 50 epochs with a batch size of 16 and a verbose setting of 1.

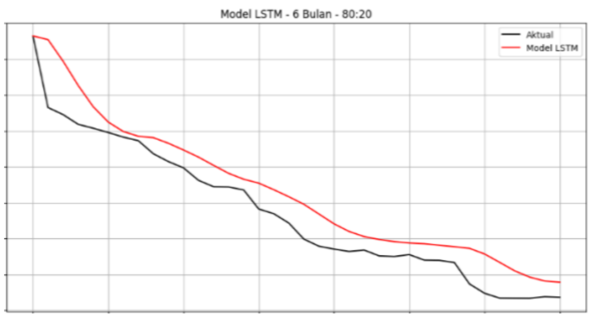

Fig. 3: LSTM model implementation with the best evaluation

Table 2: LSTM best evaluation model

Metric	Value
RMSE	1.45
MAE	1.25
MAPE	4.3%

Fig 3 and Table 2 show that the implementation of the LSTM model on a 6-month dataset with an 80:20 data split yielded an RMSE of 1.45, MAE 1.25, and MAPE of 4.3%. These results indicate that a moderate dataset produces a more optimal model, and a large training data split allows the model to adequately capture patterns.

4.1.3. Neural Prophet Model Results

The Neural Prophet model was developed from Meta's (Facebook) Prophet, with the addition of neural network regression to handle complex seasonal patterns. Neural Prophet uses yearly seasonality, weekly seasonality, and daily seasonality.

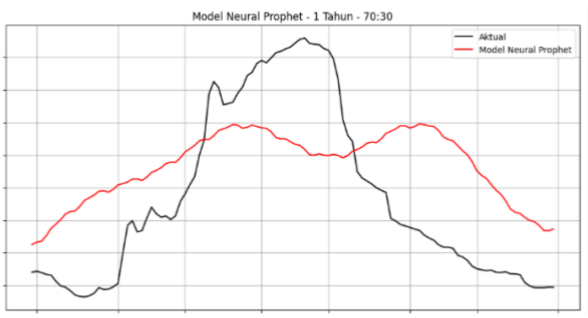

Fig. 4: Neural Prophet model implementation with the best evaluation

Table 3: Neural Prophet best evaluation model

Metric	Value
RMSE	11.61
MAE	10.69
MAPE	31.65%

Fig 4 and Table 3 show that the implementation of the Neural Prophet model on a 1-year dataset with a 70:30 data split yielded an RMSE of 11.61, an MAE of 10.69, and a MAPE of 31.65%. These results indicate that a large dataset produces a fairly optimal model, and a moderate training data split makes the model more capable of captur-ing patterns.

4.2. Model Performance Comparison

After the model is implemented according to the test scenario, the next step is to compare its performance. The model performance comparison is outlined in Table 4 below.

Table 4: Model Performance Comparison

Model	Dataset	Split	RMSE	MAE	MAPE
XGBoost	3 Months	60:40	6.18	5.46	20.84%
		70:30	3.34	3.06	11.86%
		80:20	1.95	1.62	6.47%
	6 Months	60:40	2.65	1.87	3.91%
		70:30	1.58	1.05	2.76%
		80:20	0.52	0.4	1.42%
	1 Year	60:40	2.55	1.71	4.30%
		70:30	2.13	1.49	3.39%
		80:20	1.96	1.32	2.76%
LSTM	3 Months	60:40	6.39	6.36	22.28%
		70:30	4.04	4.01	14.70%
		80:20	1.74	1.66	6.38%
	6 Months	60:40	3.38	2.77	6.61%
		70:30	3.87	3.04	8.58%
		80:20	1.45	1.25	4.30%
	1 Year	60:40	4.1	3.13	7.66%
		70:30	2.56	1.88	4.79%
		80:20	2.73	2.11	5.10%
Neural Prophet	3 Months	60:40	429.08	362.66	1309.90%
		70:30	461.44	367.99	1413.25%
		80:20	399.37	285.6	1132.39%
	6 Months	60:40	35.35	29.2	66.87%
		70:30	14.41	12.65	40.01%
		80:20	19.48	17.55	59.35%
	1 Year	60:40	51.77	44.82	140.95%
		70:30	11.61	10.69	31.65%
		80:20	21.75	20.05	58.49%

Based on the overall analysis results as table 4 above, it can be concluded that model selection significantly influences price prediction accuracy. XGBoost proved to be the best model in this study, providing the most accurate and stable prediction results. Furthermore, variations in data length and splitting proportion also significantly impacted model performance, with longer datasets and a larger proportion of training data tending to yield better performance.

4.3. Discussion of Research Results

The results show that dataset size has a significant impact on model performance. On the 3-month dataset, almost all models performed less than optimally, especially Neural Prophet. On the 6-month dataset, there was a significant improvement in performance, particularly for XGBoost. On the 1-year dataset, model performance tended to be stable, although not always better than the 6-month dataset.

This phenomenon suggests that sufficient data is important for improving a model's ability to capture patterns. However, too much data does not always guarantee improved performance, especially if the data contains noise or inconsistent patterns.

The data split ratio also impacts prediction results. In general, an 80:20 split yielded the best results in most experiments. This is due to the greater amount of training data used to build the model.

However, in some cases, such as with Neural Prophet on the 6-month dataset, a 70:30 split actually produced the best performance. This suggests that a balance between training and testing data is also important to avoid overfitting.

The results of this study align with previous studies that show XGBoost has superior performance in various prediction scenarios, particularly on tabular and non-linear data. On the other hand, LSTM is often reported to excel in time series prediction with large and complex data. However, this study shows that on a relatively limited dataset, LSTM performance is not yet able to outperform tree-based models like XGBoost. Meanwhile, Neural Prophet, developed as a hybrid model, does not always produce better results. This indicates that model complexity is not always directly proportional to prediction accuracy.

4.4. Data Prediction Based on the Best Model

Based on the research results above, several predictions were obtained for each test scenario: for the 3-Month Split 60:40 Dataset, the 3-Month Split 70:30 Dataset, the 3-Month Split 80:20 Dataset, the 6-Month Split 60:40 Dataset, the 6-Month Split 70:30 Dataset, the 6-Month Split 80:20 Dataset, the 1-Year Split 60:40 Dataset, the 1-Year Split 70:30 Dataset, and the 1-Year Split 80:20 Dataset. These predictions can be used by local governments to formulate price stabilization policies, such as providing reserve stocks and increasing distribution efficiency between regions. The figures are as follows.

- **3-Month Dataset Split 60:40 - XGBoost**

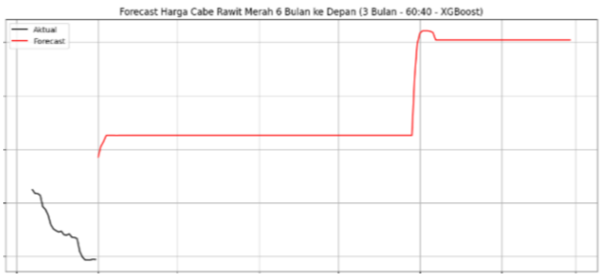

Fig. 5: Red Chili Pepper Price Forecast for the Next 6 Months (3 Months - 60:40 - XGBoost)

- **3-Month Dataset Split 70:30 - XGBoost**

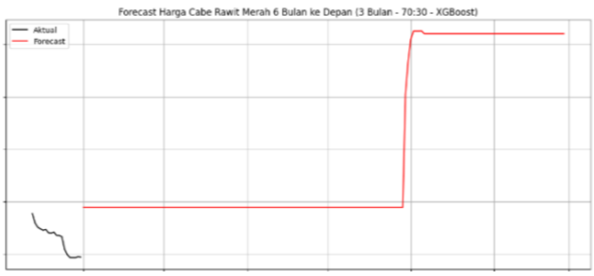

Fig. 6: Red Chili Pepper Price Forecast for the Next 6 Months (3 Months - 70:30 - XGBoost)

- **3-Month Dataset Split 80:20 - LSTM**

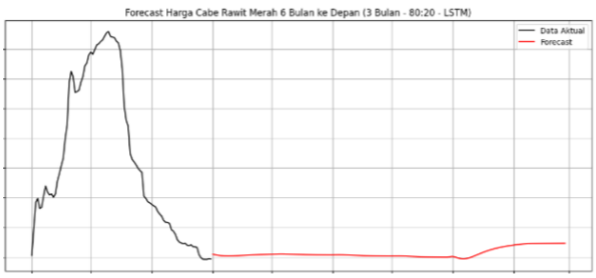

Fig. 7: Red Chili Pepper Price Forecast for the Next 6 Months (3 Months - 80:20 - LSTM)

- **6-Month Dataset Split 60:40 - XGBoost**

Fig. 8: Red Chili Pepper Price Forecast for the Next 6 Months (6 Months - 60:40 - XGBoost)

- **6-Month Dataset Split 70:30 - XGBoost**

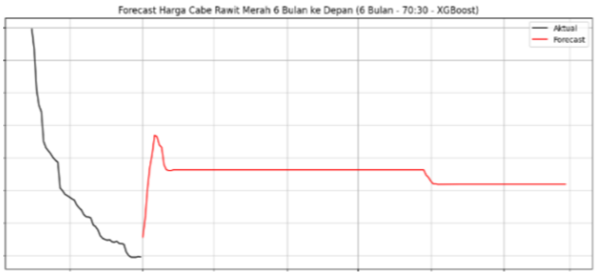

Fig. 9: Red Chili Pepper Price Forecast for the Next 6 Months (6 Months - 70:30 - XGBoost)

- **6-Month Dataset Split 80:20 - XGBoost**

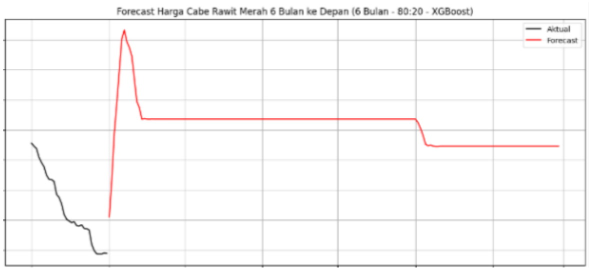

Fig. 10: Red Chili Pepper Price Forecast for the Next 6 Months (6 Months - 80:20 - XGBoost)

- **1-Year Dataset Split 60:40 - XGBoost**

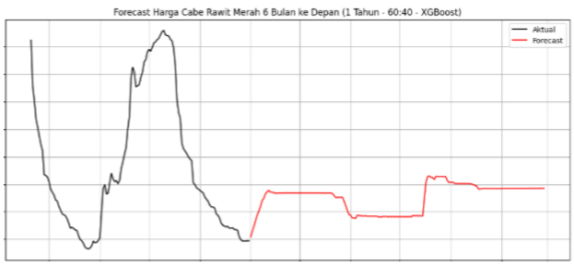

Fig. 11: Red Chili Pepper Price Forecast for the Next 6 Months (1 Year - 60:40 - XGBoost)

- **1-Year Dataset Split 70:30 - XGBoost**

Fig. 12: Red Chili Pepper Price Forecast for the Next 6 Months (1 Year - 70:30 - XGBoost)

- **1-Year Dataset Split 80:20 - XGBoost**

Fig. 13: Red Chili Pepper Price Forecast for the Next 6 Months (1 Year - 80:20 - XGBoost)

The results of this study in figure 5-13 have important implications, both academically and practically. Academically, this study provides empirical evidence that model selection is highly dependent on data characteristics. More complex models are not always the best choice. Practically, the results of this study can be used by the government or market players in selecting a price prediction model for red chili peppers. Using an appropriate model, such as XGBoost, can lead to more accurate price predictions, thus aiding decision-making regarding distribution and price stabilization.

5. Conclusion

The results showed that the XGBoost-based model with parameters n_estimators=500, learning_rate=0.05, max_depth=5, subsample=0.8, colsample_bytree=0.8, and random_state=42 produced the best performance with the lowest error value in all evaluation metrics, namely RMSE, MAE, and MAPE, especially in the 6-month dataset scenario with a data split of 80:20 (RMSE 0.52; MAE 0.40; MAPE 1.42%), which indicates a high ability to capture patterns of fluctuations in the price of red cayenne pepper and has good generalizations in the prediction of time series of food commodities; meanwhile, the Long Short-Term Memory (LSTM) model with epoch=50, batch_size=16, and verbose=1 parameters showed fairly competitive performance but still below XGBoost, with the best results on the 3-month dataset (RMSE 1.74; MAPE 6.38%), but tends to be unstable when the amount of data increases, indicating high sensitivity to parameter tuning, data quality, and pattern complexity; on the other hand, the Neural Prophet model with annual, weekly, and daily seasonality configurations showed the lowest performance with very high error values, especially in the 3-month dataset (MAPE >1000%), which indicates limitations in capturing complex non-linear patterns in fluctuating price data and the possibility of suboptimal configuration and training processes; Overall, model selection has been shown to have a significant impact on prediction accuracy, with XGBoost being the most optimal model, and variations in dataset length and proportion of trained data also contribute significantly to improved model performance, with longer datasets and larger proportions of trained data tend to result in better performance; Therefore, further research is recommended to expand the scope of regions and commodities such as onions, rice, and curly chili peppers to test model generalizations, include external variables such as rainfall, temperature, production volume, and logistics data to improve the representation of market conditions, and develop hybrid approaches such as a combination of LSTM-XGBoost or LSTM-Prophet to integrate the advantages of high accuracy and model interpretability resulting in a more stable and comprehensive prediction system.

Acknowledgement

Thank you to all parties who supported the successful implementation of this research.

References

[1] A. Mandarsari, R. Anindita, and S. Budi, “Price Volatility Analysis of Cayenne Pepper (Capsicum frutescens) in East Java,” Agricultural Socio-Economics Journal, vol. 20, no. 2, Apr. 2020. doi: 10.21776/UB.AGRISE.2020.020.2.5.
[2] M. N. E. Brahmana, Sahara, and N. K. Hidayat, “Price Volatility Analysis of Red and Cayenne Pepper of Java Islands during Covid-19 Pandemic,” Journal of Economics, Finance and Accounting Studies, vol. 4, no. 4, Sep. 2022. doi: 10.32996/jefas.2022.4.4.2.
[3] I. Marina, D. Sukmawati, E. Juliana, and others, “Dinamika Pasar Komoditas Pangan Strategis: Analisis Fluktuasi Harga Dan Produksi,” Paspalum, vol. 12, no. 1, Apr. 2024. doi: 10.35138/paspalum.v12i1.700.
[4] B. W. Sari and D. Prabowo, “Analisis perbandingan prediksi harga rumah dengan Random Forest, Gradient Boosting, dan XGBoost,” Intellect: Indonesian Journal of Learning and Technological Innovation, vol. 4, no. 1, pp. 42–51, 2025.
[5] A. F. Alkayes and T. Sugihartono, “Perbandingan algoritma XGBoost dan LSTM dalam prediksi harga saham Tesla menggunakan data tahun 2025,” J. Pendidik. dan Teknol. Indones., vol. 5, no. 6, pp. 1563–1573, 2025.
[6] A. Primawati, F. A. Mustika, and Y. Wibawanti, “Analisis tren dan prediksi harga emas menggunakan Prophet dan Long Short Term Memory (LSTM),” Simetris: Jurnal Teknik Mesin, Elektro dan Ilmu Komputer, vol. 16, no. 2, 2025.
[7] R. K. K. Sitepu, “Price Transmission in the Indonesian Red Chili Market: Static and Dynamic Models,” Jurnal Ekonomi Kuantitatif Terapan, vol. 15, no. 2, Aug. 2022. doi: 10.24843/jekt.2022.v15.i02.p04.
[8] A. Fauzi and V. Andriani, “Pengaruh meningkatnya harga cabai terhadap permintaan dan penawaran di Indonesia,” Jurnal Akuntansi dan Manajemen Bisnis, vol. 3, no. 1, Apr. 2023. doi: 10.56127/jaman.v3i1.645.
[9] R. J. Hyndman and G. Athanasopoulos, Forecasting: Principles and Practice, 3rd ed. Melbourne: OTexts, 2021.
[10] T. Chen and C. Guestrin, “XGBoost: A Scalable Tree Boosting System,” in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 785–794.
[11] F. Z. Ulya, S. Khomsah, and N. A. F. Tanjung, “Perbandingan Algoritma XGBoost dan LSTM untuk Memprediksi Harga Bitcoin Berdasarkan Harga Harian, Sentimen, dan Google Trends Index,” Jurnal Teknologi Informasi dan Ilmu Komputer, vol. 12, no. 6, 2025.
[12] F. Yuan et al., “An xLSTM–XGBoost Ensemble Model for Forecasting Non-Stationary and Highly Volatile Gasoline Price,” Computers, vol. 14, no. 7, 2025.
[13] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.
[14] S. Mehtab, J. Sen, and A. Dutta, “Stock Price Prediction Using Machine Learning and LSTM-Based Deep Learning Models,” arXiv preprint arXiv:2009.10819, 2020.
[15] A. H. Mahmoud et al., “Enhancing the Exploitation of Natural Resources for Green Energy: An Application of LSTM-Based Meta-Model for Aluminum Prices Forecasting,” Resources Policy, vol. 92, 2024.
[16] O. Triebe et al., “NeuralProphet: Explainable Forecasting at Scale,” arXiv preprint arXiv:2111.15397, 2021.