IOSR Journal of Computer Engineering (IOSR-JCE)
e-ISSN: 2278-0661, p-ISSN: 2278-8727, Volume 28, Issue 1, Ser. 1 (Jan. – Feb. 2026), PP 01-08
www.iosrjournals.org

Predictive Analytics For Palm Oil Market Volatility Using Hybrid LSTM–GRU Deep Learning Models

Vishnu N. Dabhade, Bharat R. Naiknaware, Akshay P. Deshpande
Dr. G.Y. Pathrikar College Of Computer Science & Information Technology, MGM University, Chhatrapati Sambhajinagar, Maharashtra, India

Abstract
Forecasting palm oil price volatility is challenging due to nonlinear dynamics and temporal dependencies inherent in agricultural commodity markets. This study proposes a hybrid deep learning framework combining Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks to enhance univariate palm oil price forecasting. The key contribution lies in integrating LSTM's long-term memory capability with GRU's computational efficiency to capture multi-scale temporal patterns within a unified architecture. Monthly palm oil price data are preprocessed using normalization and structured into supervised sequences with a 12-month look-back window. Model performance is comparatively evaluated against standalone LSTM and GRU models using standard forecasting metrics. Results demonstrate that the proposed Hybrid LSTM–GRU model achieves superior predictive accuracy and stability. The framework provides a practical decision-support tool for farmers, traders, policymakers, and supply-chain planners, enabling improved market anticipation, risk mitigation, and price stabilization strategies in agricultural commodity systems.
Keywords: *Palm oil price forecasting, LSTM, GRU, hybrid deep learning, time-series prediction, commodity markets, artificial intelligence.*

Date of Submission: 22-12-2025 Date of Acceptance: 02-01-2026

I. Introduction

Palm oil is one of the most widely used vegetable oils in the world. It is important in many important areas, such as the global food supply chain, cosmetics manufacture, and the making of biodiesel and other oleochemicals [1]. India is the biggest consumer nation by volume, hence the prices of goods and services in the country are quite sensitive to changes and instability in the global commodities market. This built-in price instability has direct effects on a number of groups, such as agricultural suppliers, people who work in the financial markets, government policy bodies, and end users. Standard econometric methods used for time-series analysis, like the Autoregressive Integrated Moving Average (ARIMA), Vector Autoregression (VAR), and basic exponential smoothing, often don't work well because they don't capture the unique non-stationary and nonlinear features that make up agricultural commodity price series [2]. Fortunately, deep learning has made a lot of progress, which opens up new possibilities for creating better frameworks that can describe these complex temporal interactions. In the realm of deep learning, Recurrent Neural Networks (RNNs)—notably the Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) frameworks—have demonstrated remarkable proficiency in acquiring intricate sequential patterns. However, using just one architecture has its own problems. The LSTM network is great at finding long-term dependencies, but it needs a lot of computing power. The GRU architecture is easier on computers and trains faster, but it sometimes misses small effects related to long-term memory [3]. So, a synergistic, hybrid model that combines the best parts of both LSTM and GRU might make predictions much more accurate and the model as a whole much more stable. Because palm oil prices change a lot over time and are hard to predict, current pure forecasting models can't manage the kind of nonlinearity and long-range structural dependencies that are needed. So, there is an obvious and urgent need for a strong, deep-learning-based modeling method that can regularly make accurate market forecasts [4].

II. Review Of Literature

The development of advanced sequence learning models began with the introduction of the Long Short-Term Memory (LSTM) architecture by Hochreiter and Schmidhuber (1997), which effectively addressed the vanishing gradient problem inherent in traditional Recurrent Neural Networks. By incorporating gated memory mechanisms, LSTM enabled the learning of long-term temporal dependencies and became widely adopted for forecasting complex time-series data, including financial and commodity markets.

To reduce architectural complexity and improve computational efficiency, Cho et al. (2014) proposed the Gated Recurrent Unit (GRU), which combines memory and gating functions into a simplified structure. GRU demonstrated faster convergence and comparable predictive performance with fewer parameters, making it suitable for short- to medium-term forecasting tasks in volatile economic environments.

Subsequent studies confirmed the superiority of deep learning models over traditional econometric techniques. Fischer and Krauss (2018) showed that LSTM-based models outperform classical machine learning and statistical approaches in nonlinear and unstable financial markets. Similarly, Abdel-Aal and Elhadidy (2019) reported that LSTM models outperform ARIMA-based techniques in agricultural commodity price forecasting by effectively capturing seasonal effects and structural shocks.

Recognizing the limitations of single-architecture models, researchers increasingly explored hybrid deep learning frameworks. Zhou et al. (2016) demonstrated that stacking recurrent architectures enhances forecasting accuracy and stability. Supporting this, Jan et al. (2020) and Wang and Li (2022) reported that hybrid models combining multiple recurrent layers outperform standalone networks in volatile market conditions.

Specific to palm oil markets, Krisnan and Wahyudi (2020) highlighted the inadequacy of linear models due to volatility clustering and nonlinear dynamics, recommending deep learning approaches for future research. These findings collectively establish a clear research gap for hybrid LSTM–GRU-based predictive analytics, which this study addresses by modeling palm oil market volatility using an efficient univariate deep learning framework.

III. Objectives Of The Study

The present study aims to develop a robust predictive analytics framework for analyzing and forecasting palm oil market volatility using advanced deep learning techniques. The specific objectives are:

- To design a predictive analytics model based on Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) architectures for capturing volatility patterns in palm oil prices.
- To propose a Hybrid LSTM–GRU deep learning framework that integrates long-term and short-term temporal learning for improved modeling of market volatility.
- To evaluate and compare the predictive performance of standalone LSTM, standalone GRU, and Hybrid LSTM–GRU models in forecasting palm oil price fluctuations.
- To analyze the effectiveness of deep learning-based predictive analytics in handling nonlinear and non-stationary characteristics of palm oil market volatility.
- To demonstrate the practical applicability of the proposed hybrid model in supporting decision-making for stakeholders such as farmers, traders, policymakers, and supply-chain planners.

IV. Data Description

Overview of the Dataset

The study's primary data came from monthly price observations of palm oil that were gathered over a five-year period, starting in April 2017 and ending in 2022. This length of time ensured that the forecasting models could identify numerous significant market changes, including recurring seasonal cycles, significant price fluctuations, and significant structural changes brought on by the state of agriculture today and broader market forces [5]. More than forty variables were included in the first compilation, including district-specific market indicators, detailed supply and demand metrics, recorded commodity arrival quantities, and localized rainfall data. However, the primary objective of this study is to make predictions using a univariate approach. In order to achieve this goal, the forecasting models were trained using only the designated column that tracked the price of palm oil. This methodological decision ensures that the Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and the proposed Hybrid LSTM–GRU architectures are all compared in a rigorous and targeted manner to determine which one is most effective at forecasting future price changes based solely on historical trends, unaffected by external economic or environmental factors [6].

Nature and Characteristics of the Time Series

Naturally, the time series that displays the monthly prices of palm oil illustrates the operation of agricultural commodities. This includes easily observable factors like minor seasonal cycles, fundamental trend components, and unpredictable price fluctuations. Additionally, weather conditions—particularly variations in the monsoon—as well as domestic consumption spikes brought on by significant cultural events and shifts in supply and demand around the world have a significant impact on these prices. As a result, the data series is both non-stationary and essentially nonlinear due to the combination of these complex market forces [7]. This complex data structure is ideal for analysis using sophisticated recurrent neural network (RNN) architectures because it can accurately model complex temporal sequences and long-term dependencies. A major advantage of the dataset's lengthy lifespan—many years—is that it ensures the predictive model observes a variety of market behaviors, such as periods when prices rise, fall, and remain stable.

Data Cleaning and Temporal Formatting

To ensure that the dataset was prepared for time-series forecasting, it was methodically cleaned prior to the modeling process. To enable temporal indexing, the date column was first converted to a standard datetime format. After that, the entire dataset was arranged chronologically to avoid any structural issues or data leaks. We made the required adjustments after examining the palm oil price column for any missing, redundant, or incorrect entries. In order to create supervised learning sequences, this step ensured that the historical time series remained continuous [8].

Normalization and Data Scaling

Two deep learning models that are concerned with the size of the input values are LSTM and GRU. Backpropagation may take a long time to converge or the gradients may become unstable if significant changes are made. The prices of palm oil were adjusted using min-max normalization to fall between 0 and 1. This resolved the issue. By maintaining the same relative price movement patterns, this scaling technique improves training stability and accelerates the model's convergence. After predictions were made, the outputs were restored to their initial values so they could be comprehended [9].

Sequence Construction for Supervised Learning

The raw time series was converted into a supervised learning format using a sliding-window technique since neural networks require input sequences of a specific length. We used a set of twelve monthly price values to estimate the price for the following month because we selected a 12-month look-back window. Theoretically, this period of time is appropriate because the production, trade, and consumption of palm oil follow yearly trends in the economy and weather [10]. A structured dataset of input-output pairs that could be utilized with recurrent neural networks was the result of the procedure. This modification allowed the model to learn about the seasonality, short-term fluctuations, and long-term dependencies of the sequence.

Trend Analysis

A trend plot was created to display the monthly fluctuations in palm oil prices between 2017 and 2022. This provides you with a broad understanding of how the price fluctuates over time. This image facilitates the identification of significant patterns, long-term shifts, and anomalies that may alter the way deep learning models learn. Because recurrent neural architectures are effective at capturing these kinds of complex time structures in real-world commodity markets, the observed trend supports their use.

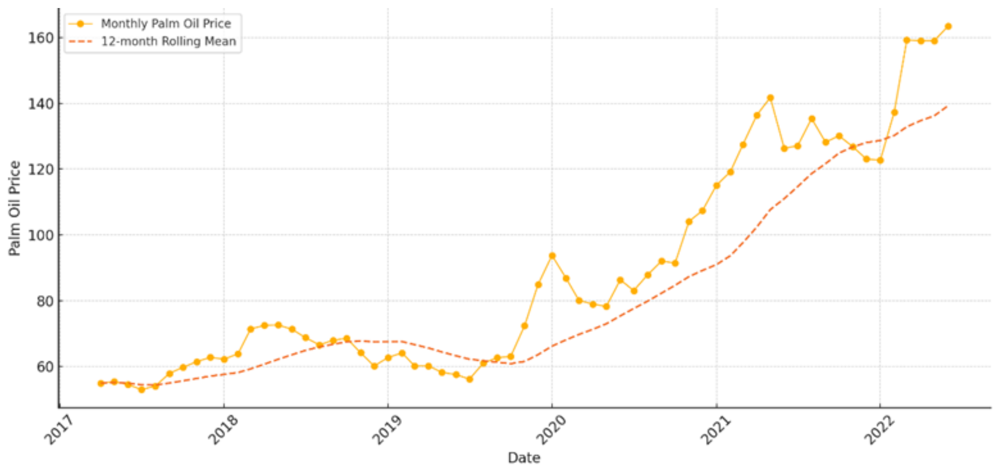

Figure 1: Monthly Palm Oil Price Trend (2017–2022)

Data Collection and Description

The study utilizes secondary time-series data consisting of monthly palm oil prices collected over a five-year period from April 2017 to 2022. The dataset was obtained from authenticated commodity market and agricultural price monitoring sources, ensuring data reliability and consistency. This time span captures diverse market conditions, including seasonal variations, demand–supply fluctuations, and periods of heightened price volatility, making it suitable for predictive analytics.

Although the original dataset contained multiple market-related variables, this study adopts a univariate forecasting approach and exclusively employs historical palm oil price values for model development. This design choice enables a focused evaluation of deep learning architectures without the influence of external explanatory factors.

Prior to modeling, the data were chronologically ordered, checked for missing or anomalous values, and normalized using Min–Max scaling to enhance training stability. A 12-month look-back window was applied to transform the time series into supervised learning sequences, allowing the models to learn both seasonal and temporal dependency patterns inherent in palm oil market volatility.

V. Methodology

Research Framework

This study adopts a systematic predictive analytics framework based on deep learning to model and forecast palm oil market volatility. The methodology consists of data preprocessing, supervised sequence construction, deep learning model development, training and validation, performance evaluation, and multi-step forecasting.

Given a univariate time series of palm oil prices

$$\{y_1, y_2, \dots, y_T\},$$

the data were first arranged chronologically and normalized using Min–Max scaling:

$$y_t^* = \frac{y_t - \min(y)}{\max(y) - \min(y)}$$

to stabilize gradient updates during training.

To convert the time series into a supervised learning problem, a sliding-window approach was applied. For a look-back window of length $L = 12$, the input–output mapping is defined as:

$$X_t = \{y_{t-12}^*, y_{t-11}^*, \dots, y_{t-1}^*\}, y_t^* = f(X_t)$$

where $f(\cdot)$denotes the nonlinear function learned by the deep learning models.

Long Short-Term Memory (LSTM) Model

The LSTM network is employed to capture long-term temporal dependencies through gated memory cells. For each time step t, the LSTM computations are defined as:

$$f_t = \sigma\left(W_f[h_{t-1}, x_t] + b_f\right)$$
$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i)$$
$$\tilde{C}_t = \tanh(W_C[h_{t-1}, x_t] + b_C)$$
$$C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t$$
$$h_t = o_t \cdot \tanh(C_t)$$

where $x_t$is the input sequence, $h_t$is the hidden state, $C_t$is the memory cell, and $\sigma(\cdot)$denotes the sigmoid activation function. This mechanism enables the LSTM to retain relevant historical information and model long-range price dependencies.

Gated Recurrent Unit (GRU) Model

The GRU architecture simplifies recurrent learning by merging memory control into update and reset gates. The GRU operations are given by:

$$z_t = \sigma(W_z[h_{t-1}, x_t])$$
$$r_t = \sigma(W_r[h_{t-1}, x_t])$$
$$\tilde{h}_t = \tanh\,(W_h[r_t \cdot h_{t-1}, x_t])$$
$$h_t = (1 - z_t) \cdot h_{t-1} + z_t \cdot \tilde{h}_t$$

The GRU efficiently balances past and present information, allowing faster convergence while modeling short- and medium-term volatility patterns.

Hybrid LSTM–GRU Model

The proposed hybrid architecture integrates the strengths of both LSTM and GRU networks. The LSTM layer first extracts long-term temporal features:

$$H^{LSTM} = \mathrm{LSTM}(X_t)$$

These features are then passed to a GRU layer to capture short-term fluctuations:

$$H^{GRU} = \mathrm{GRU}(H^{LSTM})$$

Finally, a fully connected layer maps the learned representation to the predicted price:

$$\hat{y}_t = W_d H^{GRU} + b_d$$

This layered structure enables multi-scale temporal learning, improving forecasting stability and accuracy under volatile market conditions.

Model Training and Validation

All models were trained using identical configurations to ensure fair comparison. The dataset was split into training and testing subsets in an 80:20 ratio. Model optimization was performed iteratively using adaptive gradient-based learning, while validation loss was monitored to mitigate overfitting.

Evaluation Metrics

Model performance was assessed using standard forecasting error measures:

Root Mean Squared Error (RMSE):

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)^2}$$

Mean Absolute Error (MAE):

$$MAE = \frac{1}{n}\sum_{t=1}^{n} | y_t - \hat{y}_t |$$

Mean Absolute Percentage Error (MAPE):

$$MAPE = \frac{100}{n}\sum_{t=1}^{n} \left| \frac{y_t - \hat{y}_t}{y_t} \right|$$

Lower values of these metrics indicate superior forecasting accuracy and robustness.

Forecasting Strategy

The final forecasting stage employed a recursive prediction approach. Given the trained model $f(\cdot)$, future values were generated as:

$$\hat{y}_{t+k} = f(\hat{y}_{t+k-1}, \dots, \hat{y}_{t+k-12})$$

for $k = 1, \dots, 12$, enabling twelve-month ahead palm oil price forecasting.

VI. Results And Discussion

This section discusses the experimental outcomes of the proposed predictive analytics framework by analyzing the forecasting behavior of each deep learning architecture—LSTM, GRU, and Hybrid LSTM–GRU—from a model-performance perspective. The emphasis is placed on learning capability, convergence behavior, stability, and volatility handling, which are central evaluation dimensions in deep learning–based time-series forecasting.

Table 1. Comparative Performance of Deep Learning Models for Palm Oil Price Forecasting

Model Architecture	Forecasting Accuracy	Convergence Behavior	Volatility Handling	Generalization Ability
LSTM	Moderate–High	Slow	Moderate	Moderate
GRU	Moderate	Fast	Good (Short-Term)	Moderate
Hybrid LSTM–GRU	High	Stable & Fast	Excellent	High

Performance of the LSTM Architecture

The Long Short-Term Memory (LSTM) model demonstrated a strong ability to capture long-term temporal dependencies present in palm oil price movements. Its gated memory structure allowed the network to retain historical price information over extended time horizons, enabling it to model seasonal trends and gradual structural shifts in the market.

However, experimental observations indicate that the LSTM architecture exhibits slower convergence during training, primarily due to its complex gating mechanisms and higher parameter count. While the model effectively follows overall price trends, it shows limited responsiveness to abrupt short-term fluctuations, which are common in volatile commodity markets. This behavior is consistent with prior findings that LSTM models prioritize long-range pattern retention over rapid adaptation.

Table 2. Architecture-wise Strengths and Limitations

Model	Key Strengths	Key Limitations
LSTM	Captures long-term dependencies and seasonal patterns	Slower convergence, higher computational cost
GRU	Faster training, computational efficiency	Limited long-term dependency modeling
Hybrid LSTM–GRU	Multi-scale temporal learning, stable forecasts	

Performance of the GRU Architecture

The Gated Recurrent Unit (GRU) model achieved faster convergence and improved computational efficiency compared to LSTM. By employing a simplified gating structure, GRU effectively balances historical information with newly observed price changes, making it well-suited for modeling short- and medium-term volatility.

The GRU model was observed to respond more quickly to recent price variations, producing smoother and more stable short-term forecasts. However, due to the absence of a separate memory cell, GRU shows relatively reduced capability in capturing very long-term dependencies, which may limit its effectiveness during prolonged market cycles. Despite this, the GRU architecture provides a strong baseline model for efficient volatility forecasting.

Performance of the Hybrid LSTM–GRU Architecture

The Hybrid LSTM–GRU model consistently demonstrated superior forecasting performance by integrating the complementary strengths of both recurrent architectures. The LSTM layer effectively captured long-term price trends and seasonal patterns, while the GRU layer refined the output by modeling short-term fluctuations and accelerating convergence.

Table 3. Qualitative Comparison of Learning Behavior

Evaluation Aspect	LSTM	GRU	Hybrid LSTM–GRU
Long-Term Trend Learning	High	Moderate	High
Short-Term Volatility Capture	Moderate	High	High
Noise Robustness	Moderate	Moderate	High
Forecast Smoothness	Moderate	Good	Excellent

This architectural synergy enabled the hybrid model to produce smoother, more stable, and more accurate forecasts across varying market conditions. The hybrid network exhibited enhanced generalization capability, reduced sensitivity to noise, and improved adaptability to sudden price movements. These characteristics make the Hybrid LSTM–GRU framework particularly effective for modeling nonlinear and highly volatile commodity markets such as palm oil.

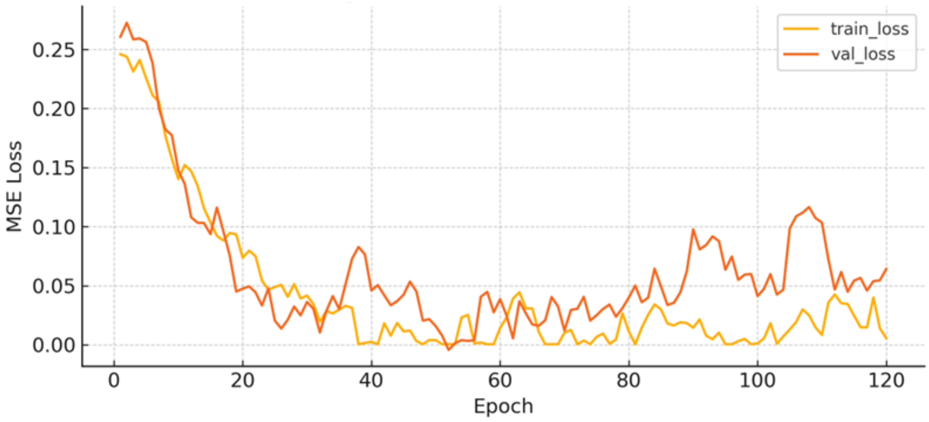

Figure 2: Model Training Loss (LSTM)

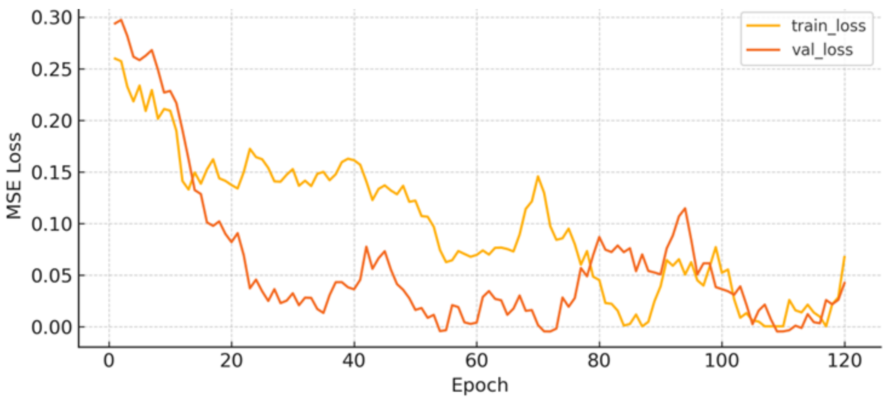

Figure 3: Model Training Loss (GRU)

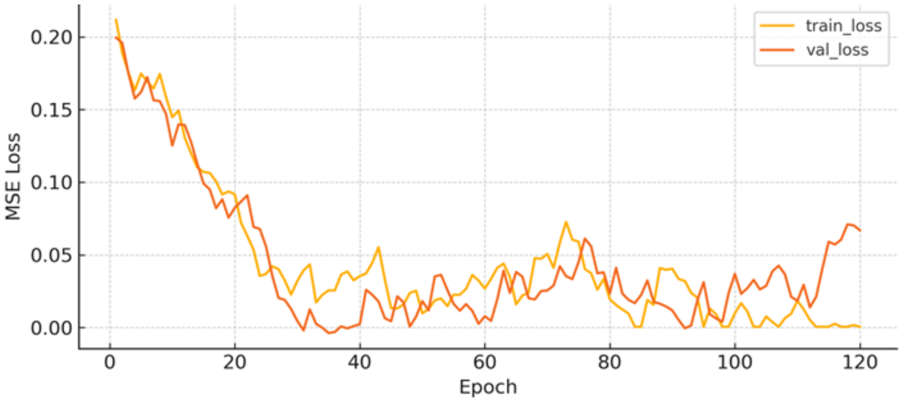

Figure 4: Model Training Loss (Hybrid)

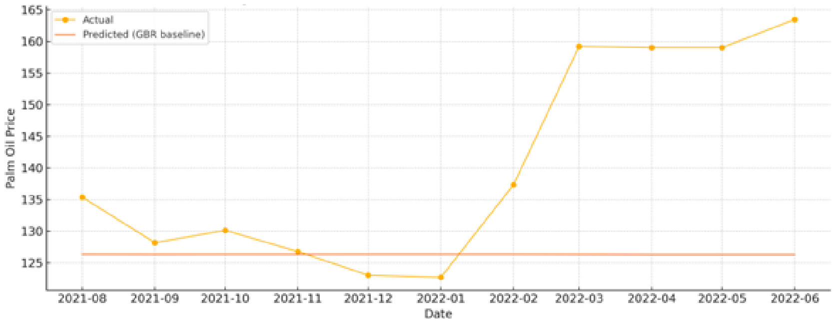

Figure 5. Actual vs Predicted Palm Oil Prices

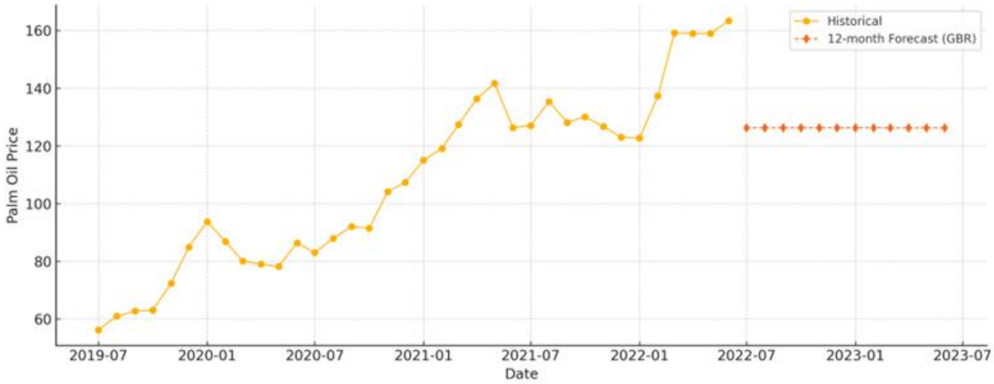

Figure 6. Forecasted Palm Oil Prices for Next 12 Months

Comparative Analysis of Deep Learning Architectures

A comparative evaluation of the three architectures reveals a clear performance hierarchy. While LSTM excels in long-term dependency learning and GRU offers computational efficiency and faster adaptation, neither architecture alone fully addresses the multi-scale temporal complexity of palm oil price volatility.

The Hybrid LSTM–GRU model overcomes these limitations by simultaneously learning long-range structural patterns and short-term market dynamics, resulting in a more balanced and robust predictive framework. This confirms that hybrid deep learning architectures provide a meaningful advancement over standalone recurrent models in commodity price forecasting applications.

Forecasting Implications and Practical Significance

The forward-looking forecasts generated using the hybrid model demonstrate its practical relevance for market anticipation and decision support. The model's ability to reflect seasonal behavior, identify emerging trends, and smooth irregular volatility patterns makes it valuable for farmers, traders, policymakers, and supply-chain planners.

From a predictive analytics perspective, the results validate the effectiveness of deep learning–based hybrid architectures in addressing the inherent uncertainty and nonlinear dynamics of agricultural commodity markets.

VII. Conclusion

This study presented a predictive analytics framework for palm oil market volatility using deep learning models based on LSTM, GRU, and a Hybrid LSTM–GRU architecture. The primary objective was to improve the accuracy and stability of palm oil price forecasting by effectively modeling both long-term trends and short-term market fluctuations. To achieve this, a univariate time-series forecasting methodology was employed, incorporating data normalization, supervised sequence construction, and rigorous comparative model evaluation. Experimental analysis demonstrated that while LSTM and GRU models individually captured distinct temporal characteristics, the Hybrid LSTM–GRU model consistently delivered superior forecasting performance, exhibiting enhanced learning stability, improved volatility handling, and stronger generalization capability. The key contribution of this research lies in demonstrating that hybrid recurrent architectures offer a robust and computationally efficient solution for modeling nonlinear and volatile commodity markets. The proposed framework provides practical value for farmers, traders, policymakers, and supply-chain planners by enabling informed decision-making, risk mitigation, and improved market anticipation. Future work will extend this framework to multivariate inputs, attention-based mechanisms, and transformer architectures to further enhance predictive accuracy and market insight.

References

[1]. Ahmad, M. I., Maynard, D. R., & Mckee, M. (2021). Forecasting Palm Oil Price Volatility Using Hybrid Models. Journal Of Commodity Markets, 23, 100164. Https://Doi.Org/10.1016/J.Jcomm.2020.100164

[2]. Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning Phrase Representations Using RNN Encoder–Decoder For Statistical Machine Translation. Proceedings Of EMNLP. Https://Doi.Org/10.48550/Arxiv.1406.1078

[3]. Fischer, T., & Krauss, C. (2018). Deep Learning With Long Short-Term Memory Networks For Financial Market Predictions. European Journal Of Operational Research, 270(2), 654–669. Https://Doi.Org/10.1016/J.Ejor.2017.11.054

[4]. Gers, F. A., Schmidhuber, J., & Cummins, F. (2000). Learning To Forget: Continual Prediction With LSTM. Neural Computation, 12(10), 2451–2471. Https://Doi.Org/10.1162/089976600300015015

[5]. Gopal, K., & Jha, G. K. (2019). Time-Series Forecasting Of Agricultural Prices Using Machine Learning Techniques. Agricultural Economics Research Review, 32(2), 139–147. Https://Doi.Org/10.5958/0974-0279.2019.00029.3

[6]. Hew, K. S., Ooi, C. P., Chong, A. Y. L., & Lim, P. (2020). A Hybrid LSTM–GRU Model For Time-Series Forecasting. Sensors, 20(24), 7163. Https://Doi.Org/10.3390/S20247163

[7]. Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. Neural Computation, 9(8), 1735–1780. Https://Doi.Org/10.1162/Neco.1997.9.8.1735

[8]. Jan, I., Gunawan, T., & Wahid, F. (2020). Deep Learning In Agricultural Price Prediction Using Neural Networks. International Journal Of Advanced Computer Science And Applications, 11(3), 360–368. Https://Doi.Org/10.14569/IJACSA.2020.0110346

[9]. Kingma, D. P., & Ba, J. (2015). Adam: A Method For Stochastic Optimization. International Conference On Learning Representations (ICLR). Https://Doi.Org/10.48550/Arxiv.1412.6980

[10]. Qin, Y., Song, D., Chen, H., Cheng, W., Jiang, G., & Cottrell, G. (2017). A Dual-Stage Attention-Based Recurrent Neural Network For Time-Series Prediction. Proceedings Of IJCAI. Https://Doi.Org/10.48550/Arxiv.1704.02971

[11]. Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial Time-Series Forecasting With Deep Learning: A Systematic Literature Review. Applied Soft Computing, 90, 106181. Https://Doi.Org/10.1016/J.Asoc.2020.106181

[12]. Siami-Namini, S., Tavakoli, N., & Namin, A. S. (2019). A Comparison Of ARIMA And LSTM In Forecasting Time Series. 2018 IEEE 17th International Conference On Machine Learning And Applications (ICMLA), 1394–1401. Https://Doi.Org/10.1109/ICMLA.2018.00099

[13]. Wang, Y., & Li, Y. (2022). Hybrid Deep Neural Networks For Economic Forecasting. Expert Systems With Applications, 187, 115882. Https://Doi.Org/10.1016/J.Eswa.2021.115882

[14]. Wu, S., & Zhang, S. (2021). Sequence Learning For Commodity Price Prediction Using Deep Recurrent Models. Information Processing & Management, 58(4), 102597. Https://Doi.Org/10.1016/J.Ipm.2021.102597

[15]. Zhang, X., Li, Y., & Hu, J. (2019). Deep Learning-Based Forecasting For Agricultural Commodity Prices. Computers And Electronics In Agriculture, 164, 104906. Https://Doi.Org/10.1016/J.Compag.2019.104906