JOURNAL OF STATISTICAL MODELING & ANALYTICS (JOSMA)
(ISSN: 2180-3102)
UNIVERSITI MALAYA

Autoregressive Modelling with Seasonal Variations for Malaysian Crude Palm Oil Price Forecasting

Aniza Akaram[1] and Arifah Bahar[1, 2*]

[1]*Department of Mathematical Sciences, Faculty of Science, Universiti Teknologi Malaysia, 81310 Johor Bahru, Johor, MALAYSIA*

[2]*UTM Centre for Industrial and Applied Mathematics, Ibnu Sina Institute of Scientific and Industrial Research, Universiti Teknologi Malaysia Johor Bahru, 81310 UTM Johor Bahru, Johor, MALAYSIA*

**Corresponding author: arifah@utm.my*

RESEARCH ARTICLE

Received 29 Sept 2025
Accepted 10 Dec 2025
Published 26 Dec 2025

Abstract

Crude palm oil (CPO) plays a vital role in Malaysia's economy, yet its price dynamics remain highly volatile due to global market fluctuations, policy changes, and demand uncertainties. Reliable short-term forecasting is therefore essential for industry stakeholders and policymakers. This study employs Autoregressive Integrated Moving Average (ARIMA) and Seasonal ARIMA (SARIMA) modelling to analyze and forecast monthly Malaysian CPO prices from 2015 to 2024. Preliminary seasonality diagnostics using STL decomposition indicated weak and statistically insignificant seasonal patterns, with a seasonal strength of 0.2046. Consequently, seasonal differencing was unnecessary, and first-order non-seasonal differencing was sufficient to achieve stationarity. Model identification and estimation were performed using an in-sample dataset (2015–2022), while model validation was conducted using out-of-sample data (2023–2024). A manual grid search across ARIMA candidates identified ARIMA(3,1,3) as the optimal model based on the Corrected Akaike Information Criterion (AICc). Residual diagnostics confirmed that the model errors behaved as white noise with no remaining autocorrelation. Out-of-sample testing further demonstrated that ARIMA(3,1,3) produced satisfactory predictive accuracy across RMSE, MAE, and MAPE metrics. Overall, the findings indicate that non-seasonal ARIMA models are sufficient for CPO price forecasting and that ARIMA(3,1,3) provides a reliable framework for short-term prediction and decision-making.

Keywords: Crude palm oil, price forecasting, seasonal ARIMA, time series analysis

1. Introduction

Malaysia is one of the world's largest producers and exporters of crude palm oil (CPO), making the commodity a cornerstone of both its agricultural sector and overall economy. Price fluctuations in CPO influence national revenue, trade balances, and the livelihoods of stakeholders across the value chain. Accurate forecasting methods are therefore vital for risk management, investment decisions, and production planning. Beyond its economic role, palm oil cultivation is water-intensive, meaning that price prediction can also serve as an indirect tool for water demand planning in plantations, linking market outcomes with sustainable resource management.

CPO prices are shaped by multiple factors such as supply–demand dynamics, export volumes, climate variability, and international market trends. These drivers often generate clear patterns and seasonal fluctuations that complicate forecasting. In response, numerous studies have investigated forecasting methods. Khalid et al. (2018), for instance, demonstrated that CPO prices are influenced not only by internal autoregressive structures but also by external commodities such as soybean oil and crude oil. Similarly, Amin et al. (2024) used a gravity stochastic frontier model to assess palm oil export performance, highlighting the role of global trade factors in shaping price movements. These studies collectively show that while external determinants are important, autoregressive patterns remain crucial in explaining persistence and lag effects in commodity markets.

A key gap in the literature lies in addressing seasonality explicitly. Palm oil markets exhibit cyclical behavior due to harvesting cycles, demand surges, and climatic variability. Tayib et al. (2021) showed the value of Seasonal Autoregressive Integrated Moving Average (SARIMA) in modeling seasonal production, while Mustapa (2025) demonstrated that combining autoregressive models with Autoregressive Conditional Heteroskedasticity (ARCH) improves volatility forecasts, again emphasizing seasonal and lagged effects. Despite this, few studies have applied SARIMA directly to Malaysian CPO price forecasting. Related evidence from international markets reinforces SARIMA's strengths where Lee et al. (2022) found that autoregressive methods outperformed exponential smoothing in crude oil price forecasting for the United States of America (USA) and Europe, suggesting similar potential for palm oil.

Alternative approaches have also been explored. During the COVID-19 pandemic, exponential smoothing methods achieved relatively low forecast errors for daily CPO prices under heightened uncertainty. More recently, machine learning approaches such as long short-term memory (LSTM) and gated recurrent units (GRU) have been applied to palm oil forecasting, showing excellent predictive capability. For example, an LSTM model achieved an R^2 of about 0.96 in predicting oil palm production (Syarovy et al., 2023), indicating very high accuracy, while GRU and LSTM networks attained low forecasting errors (normalized root mean square error, RMSE < 0.05) for crude palm oil prices (Tardini & Suharjito, 2024). Despite their impressive performance, these models require large datasets, extensive tuning, and provide limited interpretability compared to statistical models. Consequently, SARIMA remains advantageous in contexts where transparency, efficiency, and interpretability are essential for decision-making. Beyond price-focused studies, stochastic modeling has been adopted for sustainability and environmental applications in plantations (Rajakal et al., 2024; Tan et al., 2023), underscoring the broader relevance of statistical approaches in managing uncertainty across the sector.

Taken together, the literature demonstrates that while many techniques exist for CPO forecasting, SARIMA remains underutilized for Malaysian price prediction despite its capacity to handle both autoregressive and seasonal dynamics. Guided by the arguments of Hyndman and Athanasopoulos (2021) on the practical strengths of SARIMA, this study applies the model to Malaysian monthly CPO price data from 2015 to 2024. The goal is to identify an optimal specification that captures long-term trends and seasonal fluctuations, while also linking price dynamics with water demand considerations in oil palm plantations. In doing so, the research aims to provide insights for producers, investors, and policymakers seeking to align financial strategies with sustainable resource management.

2. Materials and Methods

2.1 Data Source

Monthly average crude palm oil prices in Malaysia from January 2015 to December 2024 were collected from publicly available sources of the Malaysian Palm Oil Board (MPOB) (https://mpob.gov.my/?lang=ms) and the Malaysian Palm Oil Council (https://www.mpoc.org.my/). The dataset consists of 120 observations representing price movements over a span of ten years.

2.2 Data Preprocessing

The time series was first visualized to detect trends, seasonality, and potential outliers. The Augmented Dickey-Fuller (ADF) test was applied to assess the stationarity of the series. If the series was found to be non-stationary, differencing was applied accordingly. Seasonal decomposition was also performed to separate the trend, seasonal, and residual components.

2.3 The Seasonal Autoregressive Integrated Moving Average (SARIMA) Model

The SARIMA model is an extension of the ARIMA model designed to handle both trend and seasonal patterns in time series data. While ARIMA focuses on modeling non-stationary data by combining autoregressive (*a*), differencing (*b*), and moving average (*c*) components, it does not explicitly address seasonality. SARIMA fills this gap by incorporating additional seasonal terms (*P*, *D*, *Q*) into the ARIMA framework, making it suitable for time series with regular seasonal fluctuations such as monthly crude palm oil prices, which often reflect annual demand cycles.

The SARIMA model is denoted as SARIMA(*a*,*b*,*c*)(*P*,*D*,*Q*)*s*, which combines both non-seasonal and seasonal components, allowing it to effectively model time series with trend and seasonality. It is particularly well-suited for datasets where certain patterns repeat over fixed intervals, such as commodity price peaks during harvest or export seasons. The general equation of the SARIMA model is:

$$\Phi_P(L^s)\phi_a(L)(1-L)^b(1-L^s)^D y_t = \Theta_Q(L^s)\theta_c(L)\varepsilon_t \qquad (1)$$

where y_t is the observed time series, L is the lag operator such that $L^k y_t = y_{t-k}$, ε_t is the white noise, $\phi_a(L)$ and $\theta_c(L)$ are the non-seasonal autoregressive and moving average components, respectively. $\Phi_P(L^s)$ is the seasonal autoregressive component while $\Theta_Q(L^s)$ is the seasonal moving average component. $(1-L)^b$ and $(1-L^s)^D$ represents non-seasonal and seasonal differencing operator.

Importantly, when all seasonal components P, D, Q are set to zero, the SARIMA model simplifies to a standard ARIMA (*a*, *b*, *c*) model. Then, from equation (1), it will become:

$$\phi_a(L)(1-L)^b y_t = \theta_c(L)\varepsilon_t \qquad (2)$$

This reduction highlights the flexibility of SARIMA where it can accommodate both seasonal and non-seasonal time series data within a unified modelling framework.

2.4 Model Estimation and Validation

Model identification was carried out using the ACF and PACF plots together with the Corrected Akaike Information Criterion (AICc) on the in-sample dataset (2015–2022) to determine the most appropriate ARIMA or SARIMA parameter combinations. After estimating several candidate models, their predictive performance was evaluated using the out-of-sample test data (2023–2024) based on RMSE, MAE, and MAPE. The model with the lowest test-set forecasting errors was then selected as the final model and used to generate the 12-month ahead forecast for 2025.

3. Results and Discussion

3.1 Descriptive Statistics

Table 1 shows that the monthly crude palm oil (CPO) price has an average value of 3,231.77 MYR/ton, representing the typical market level over the study period. The variance is 1,302,032, indicating that prices fluctuate considerably from month to month, reflecting the inherent volatility of the CPO market. Although this numerical value may appear large at first glance, it is a direct consequence of the wide price range, where monthly values fluctuate between 1,795 MYR/ton and 6,877 MYR/ton. Given that variance is expressed in squared units, large deviations from the mean particularly the occasional months with very high prices contribute disproportionately to the overall magnitude of the variance. In practical terms, the variance translates to a standard deviation of approximately 1,141 MYR/ton, meaning that typical monthly price changes amount to roughly one-third of the average price. This level of volatility is consistent with the behaviour of commodity markets, where prices are sensitive to global supply–demand conditions, production uncertainties, and policy interventions.

A skewness of 1.082 shows that the distribution is positively skewed, meaning that most monthly prices are below the average, with occasional high-price months raising the mean. The kurtosis value of 3.937 indicates a leptokurtic distribution, slightly more peaked than a normal distribution. This implies a greater probability of extreme price fluctuations compared to a normal distribution, highlighting the risk of sudden spikes or drops in the CPO market.

Table 1. Descriptive statistics of monthly crude palm oil

Mean	Variance	Minimum	Maximum	Skewness	Kurtosis
3231.77	1302032	1795	6877	1.082029	3.936954

Figure 1 displays the raw monthly price movements of CPO over the period 2015 to 2024. The data reveal sharp fluctuations in CPO prices from 2015 to 2024, reflecting the volatile nature of the market. A gradual increase is seen around 2017, followed by a decline in 2019. Prices then surge sharply between 2020 and 2022, reaching a peak in early 2022 before falling steeply and stabilizing toward late 2023. Another upward movement appears to emerge in 2024. These variations correspond to periods of heightened export demand and production adjustments. This initial visualization is essential for understanding the overall behavior of the dataset and provides valuable context for subsequent statistical modeling. Without such exploratory analysis, it would be difficult to identify whether the series contains a persistent trend, seasonal pattern, or irregular shocks.

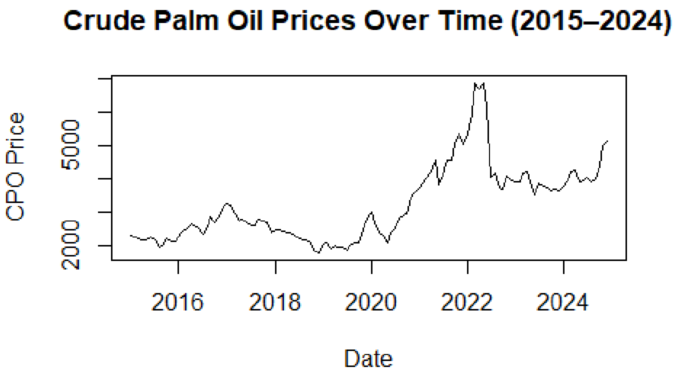

Figure 1. Monthly crude palm oil prices in Malaysia (2015-2024)

The STL (Seasonal-Trend decomposition using Loess) decomposition of monthly CPO prices in Figure 2 reveals three key components: trend, seasonality, and remainder. The seasonal component shows a consistent and repeating annual pattern, reflecting the cyclical nature of palm oil production and demand fluctuations within each year. The seasonal strength derived from STL decomposition was 0.2046, meaning that only about 20% of the overall variation can be attributed to seasonality. This low value indicates that the seasonal component exists but is relatively weak when compared with the trend and irregular fluctuations in the series.

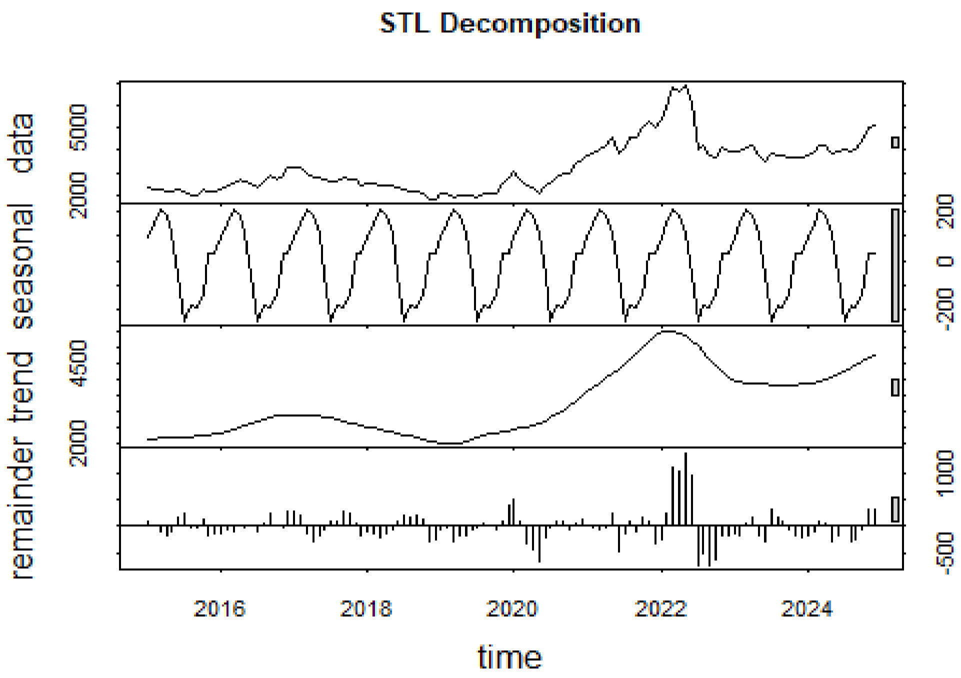

Figure 2. STL decomposition time series

The trend depicted in Figure 2 highlights a relatively stable period between 2015 and 2019, followed by a sharp rise from 2020 to early 2022, largely associated with global supply disruptions and rising demand, before moderating but remaining above pre-2020 levels in 2023–2024. The remainder captures irregular variations, with large spikes observed during 2021–2022, due to pandemic-related shocks and geopolitical uncertainties, while more recent years indicate reduced but still notable volatility. Overall, the decomposition demonstrates that CPO prices are shaped by predictable seasonal cycles, long-term structural changes, and short-term unexpected shocks.

The estimated seasonal amplitude as in Figure 3, calculated as the difference between the maximum and minimum of the seasonal component, was approximately 461.83 MYR. This indicates that, on average, CPO prices vary within a range of nearly 462 MYR due to seasonal effects alone. The presence of this consistent seasonal amplitude highlights the importance of using seasonal models such as SARIMA, as non-seasonal ARIMA would fail to capture these recurring fluctuations. From an applied perspective, the seasonal amplitude not only reflects market dynamics linked to harvest cycles and export demand but also indirectly relates to water demand. Periods of higher prices often coincide with lower yields caused by climatic stress, which increases irrigation needs in plantations. Conversely, periods of lower prices are typically associated with abundant harvests following favorable rainfall. Thus, quantifying seasonal amplitude provides a dual insight into both price forecasting accuracy and resource management challenges, particularly water use planning in oil palm production.

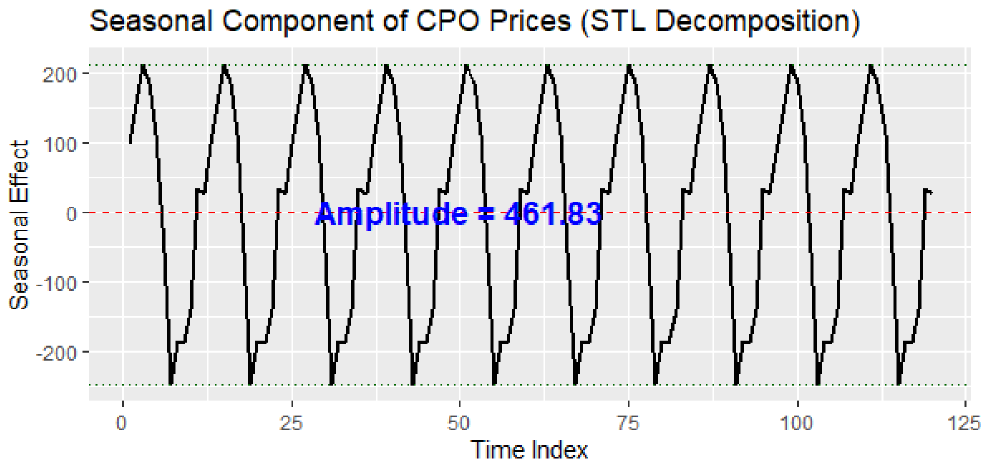

Figure 3. Seasonal amplitude

3.2 Stationary Test

3.2.1 ACF and PACF plot

The combined ACF and PACF plots in Figure 4 indicate that the raw CPO price series behaves in a memoryless manner. The ACF shows rapid decay across successive lags, with correlations dropping close to zero even at short lags. This behaviour is typical of a series dominated by short-term fluctuations around a stable mean, rather than by long-term trends or persistent effects. In a memoryless process, past shocks lose their influence quickly, so current values are largely independent of distant past observations. The PACF in the same figure reinforces this interpretation. Only the very first lag shows a noticeable partial autocorrelation, while subsequent lags are near zero, indicating that once immediate past values are accounted for, earlier observations have little effect on current prices. This pattern is characteristic of a memoryless or weakly dependent series, where price movements are primarily driven by recent shocks and quickly revert to the equilibrium level, rather than retaining long-lasting effects.

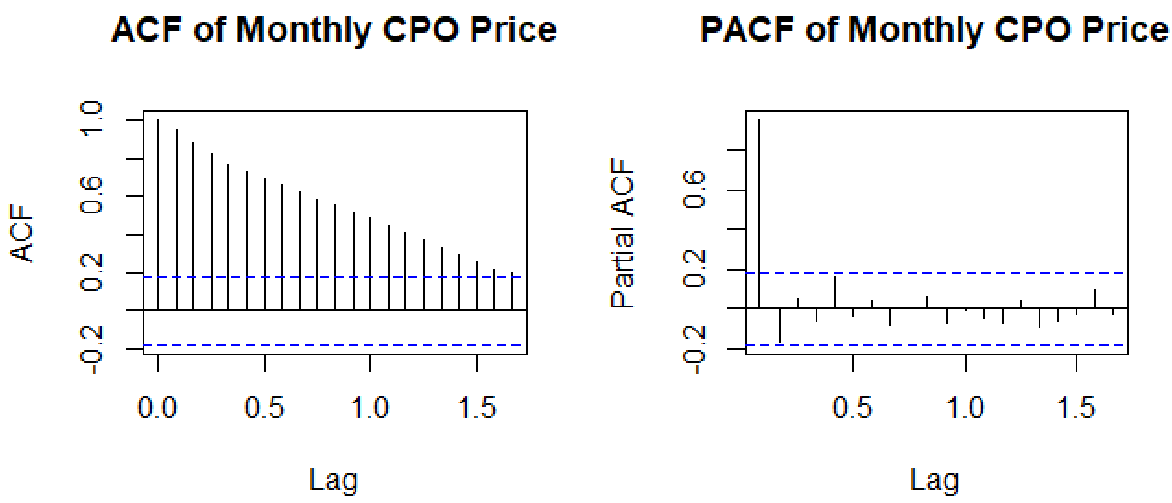

Figure 4. ACF and PACF Plot of Monthly Crude Palm Oil Price

Taken together, the ACF and PACF in Figure 4 demonstrate that the raw CPO price series does not satisfy the conditions of stationarity. The gradual decay in autocorrelation and the prominent partial autocorrelations at initial lags both point to a series dominated by long-term trend components. This diagnostic finding underscores the need for formal stationarity testing and subsequent differencing before fitting ARIMA or SARIMA models. Without correcting for non-stationarity, model parameters would be unreliable and forecasts would be biased due to unresolved trend persistence in the data.

3.2.2 The Augmented Dickey-Fuller Test

To statistically verify non-stationarity, the Augmented Dickey-Fuller (ADF) test was applied. The null hypothesis (H_0) of the ADF test states that the series contains a unit root, which implies it is non-stationary. The alternative hypothesis (H_A) asserts that the series is stationary. The ADF test returned a *p*-value of 0.5384, which is greater than the 5% significance level (α = 0.05). Thus, we fail to reject H_0, confirming that the series is indeed non-stationary in its original form. This result justifies the need for differencing before proceeding to SARIMA modeling.

3.2.3 Differencing

Figure 5 illustrates the ACF and PACF of the CPO price series after applying first-order non-seasonal differencing. The ACF shows a sharp drop in autocorrelation immediately after lag 1, with subsequent correlations remaining within the significance bounds. This pattern indicates that the persistent dependence observed in the original series, driven primarily by a strong upward trend, has been effectively removed. The PACF similarly displays only minor and isolated spikes, further suggesting the absence of sustained lag relationships. Collectively, these plots demonstrate that the differenced series behaves substantially more like a stationary process, with dependence now limited to short-term fluctuations rather than gradual long-term movements.

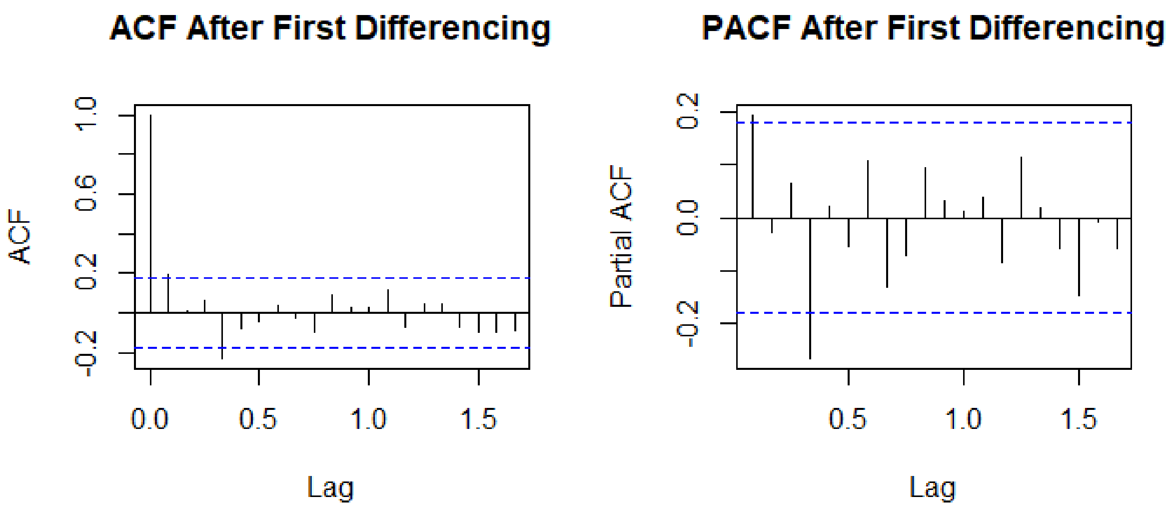

Figure 5. ACF and PACF Plot After First-Order Differencing

The behaviour of the ACF and PACF after differencing also clarifies the rationale for adopting first-order differencing ($b = 1$) rather than seasonal differencing ($D = 1$). The absence of a significant spike at the seasonal lag (lag 12) in the differenced ACF indicates that annual seasonal dependence is not sufficiently strong to warrant seasonal differencing. This finding is fully consistent with the results of the seasonality significance tests. Both the Kruskal–Wallis test (p-value = 1.0000) and the Friedman test (p-value = 0.6545) revealed no statistically meaningful monthly pattern, while the STL-derived seasonal strength value of 0.2046 further confirmed that seasonal effects account for only a small fraction of total variation. As a result, applying seasonal differencing would risk over-differencing the series, potentially introducing excessive noise and diminishing model performance. The evidence therefore supports the choice of $b = 1$ and $D = 0$ for subsequent SARIMA modelling.

The improvement in stationarity after first differencing is further verified through formal statistical tests. The Augmented Dickey–Fuller (ADF) test returned a test statistic of –5.1499 with a p-value of 0.01, allowing the rejection of the unit root hypothesis at the 5% significance level. This confirms that the differenced series does not contain a unit root and is therefore stationary. Complementing this result, the KPSS test yielded a test statistic of 0.06316, which falls well below the critical value for trend stationarity. The KPSS test thus fails to reject the null hypothesis of stationarity.

Taken together, the ACF and PACF behaviour and the outcomes of the ADF and KPSS tests provide consistent evidence that first-order differencing is sufficient to achieve stationarity, and that additional seasonal differencing is unnecessary. The differenced series is therefore appropriate for SARIMA model estimation, ensuring that parameter estimates remain unbiased and forecasting results are reliable.

3.3 In-Sample and Out-of-Sample Framework

Before undertaking model identification, the dataset was divided into two segments to enable a structured and unbiased evaluation of forecasting performance. The first segment, covering the period 2015–2022, was designated as the in-sample (training) dataset, while the subsequent period 2023–2024 served as the out-of-sample (testing) dataset. This division ensures that parameter estimation, diagnostic analysis, and model selection are carried out solely on historical data that would have been available at the start of the forecasting horizon.

The in-sample data are used to identify appropriate SARIMA orders, estimate parameters, examine residual behaviour, and compare competing models using information criteria such as AICc and BIC. These diagnostics reveal how well a model captures the underlying structure of the series. The out-of-sample segment, by contrast, provides an independent basis for evaluating predictive accuracy and detecting overfitting. Forecast errors calculated on the test set (RMSE, MAE, and MAPE) allow the performance of alternative specifications to be compared under realistic forecasting conditions.
By establishing this framework prior to model identification, the analysis ensures that the final model is selected not only for its statistical fit to the training data but also for its demonstrated ability to generalise to unseen observations. This approach increases the robustness and credibility of the forecasting results.

3.4 Model Identification

Model identification was carried out exclusively using the in-sample (training) dataset covering 2015–2022, in accordance with the in-sample/out-of-sample framework established earlier. Prior diagnostics indicated that a single non-seasonal difference was sufficient to stabilise the series. This conclusion was supported by the rapid decay observed in the ACF and PACF of the differenced series, as well as by the results of the ADF and KPSS tests, both of which confirmed stationarity at the 5% significance level. Seasonal differencing was not applied because the STL-based seasonal strength measure showed that the monthly seasonal pattern was weak with seasonal strength equal to 0.2046. The absence of a consistent spike at lag 12 in the differenced ACF further supported the exclusion of seasonal differencing. Accordingly, the differencing orders were fixed at b = 1 and D = 0 for all candidate SARIMA models considered during the in-sample identification stage. Given this structure, candidate models were restricted to the class:

$$SARIMA(a,1,c)(P,0,Q)_{12},$$

where the non-seasonal autoregressive and moving-average orders (a and c) and the seasonal AR and MA orders (P and Q) were allowed to vary. The grid was constrained deliberately to avoid over-parameterisation, which is a common risk in monthly economic series with weak seasonal effects. The parameter ranges were therefore set to:

i. Non-seasonal components: $a, c \in \{0, 1, 2, 3\}$
ii. Seasonal components: $P, Q \in \{0, 1\}$
iii. Differencing orders: $b = 1, D = 0$

This configuration yields 64 possible models a manageable number that covers all realistic autocorrelation structures implied by the ACF and PACF while avoiding unnecessary seasonal complexity.

3.5 Model Estimation

All ARIMA candidates were estimated using the maximum likelihood method. Model comparison was conducted using the Corrected Akaike Information Criterion (AICc), which is preferred for moderate-length monthly datasets because it incorporates a correction term that penalizes unnecessary parameterization. This reduces the risk of selecting an overly complex model while still favoring specifications that provide genuine improvements in fit.

A systematic grid search was performed across a restricted range of autoregressive and moving-average orders. For each candidate specification ARIMA(a,1,c), the AICc value, parameter significance, and residual behaviour were evaluated. The model achieving the smallest AICc value was retained as the preferred specification. Based on this procedure, the ARIMA(3,1,3) model emerged as the best-fitting structure. This model provided the lowest AICc among all evaluated alternatives and demonstrated strong in-sample performance, reflecting its ability to capture short-run dynamics in the differenced CPO price series.

3.6 Residual Diagnostics

Residual diagnostics were carried out to assess whether the ARIMA(3,1,3) model satisfied the fundamental assumptions of independence, stationarity, and approximate normality. The diagnostic plots in Figure 6 show that the residual series fluctuates randomly around zero with no visible trend, although several isolated spikes are present. The residual histogram exhibits slightly heavy tails, suggesting mild departures from normality, which is common for commodity price data. The autocorrelation function of the residuals remains within the 95% confidence bounds, and the Ljung–Box test indicates no significant remaining autocorrelation, confirming that the error structure approximates white noise. Overall, the diagnostic results support the adequacy of the ARIMA(3,1,3) model for forecasting purposes.

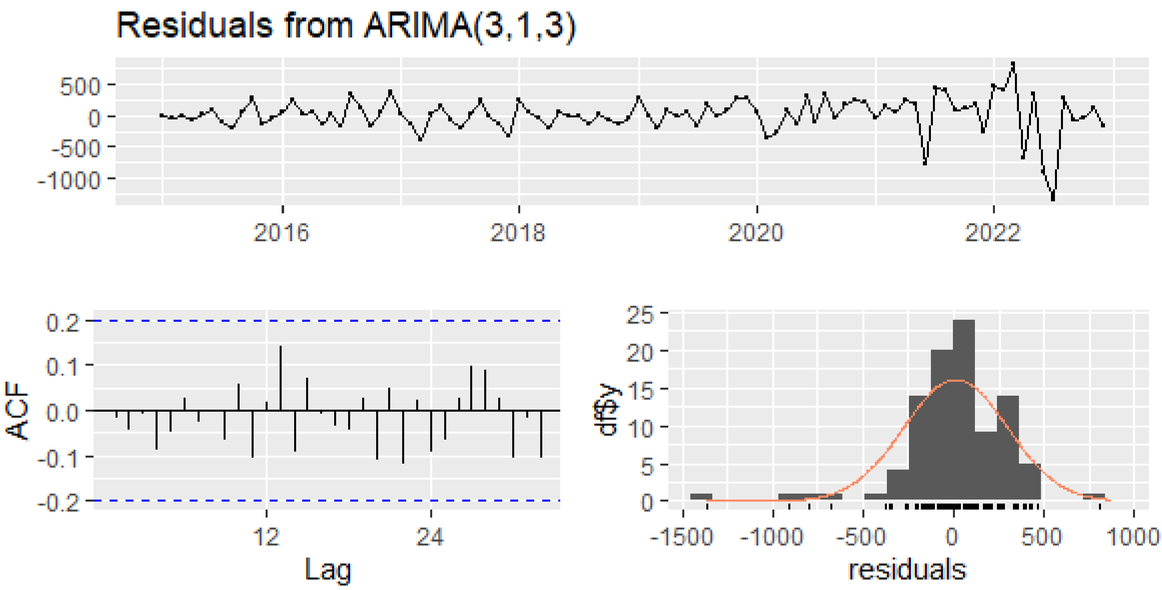

.

Figure 6. Residuals for ARIMA(3, 1, 3)

3.7 Model Validation

To evaluate forecasting performance, the selected ARIMA(3,1,3) model was assessed using the out-of-sample test dataset covering January 2023 to December 2024. Forecast accuracy was measured using the root mean square error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE). Table 2 summarises the results obtained for the validation period.

Table 2. Forecast Performance of Automatic and Manual SARIMA

Model	RMSE	MAE	MAPE
ARIMA (3, 1, 3)	439.96	290.97	6.68%

The forecasting errors indicate that the model is capable of reproducing the short-term evolution of CPO prices but also reflects the high volatility characteristic of the commodity market. The magnitude of the errors is consistent with irregular fluctuations observed in the historical data.

The final selection of the forecasting model was based on three key considerations: overall goodness of fit, residual adequacy, and predictive accuracy. Goodness of fit was evaluated using the AICc and the residual variance, ensuring that the chosen model achieved an optimal balance between explanatory power and parameter efficiency. Residual adequacy was assessed through diagnostic plots and the Ljung–Box test, which confirmed whether the modelled errors exhibited the properties of white noise. Predictive performance was examined using out-of-sample accuracy measures, namely RMSE, MAE, and MAPE, calculated for the 2023–2024 validation period. Among all candidates, the ARIMA(3,1,3) model provided the most favourable combination of these criteria. It produced the lowest AICc value, passed the residual independence checks, and delivered acceptable forecasting accuracy despite minor deviations from normality in the residual distribution. Taken together, these results indicate that ARIMA(3,1,3) offers the best compromise between parsimony, statistical soundness, and predictive capability. Accordingly, this model was selected as the

3.8 Time Series Visualization

Once the ARIMA(3,1,3) model was selected based on its superior predictive accuracy in the validation stage, it was re-estimated using the complete dataset to generate a one-year ahead forecast for 2025. Figure 7 displays the projected monthly CPO prices together with their associated 80% and 95% uncertainty intervals.

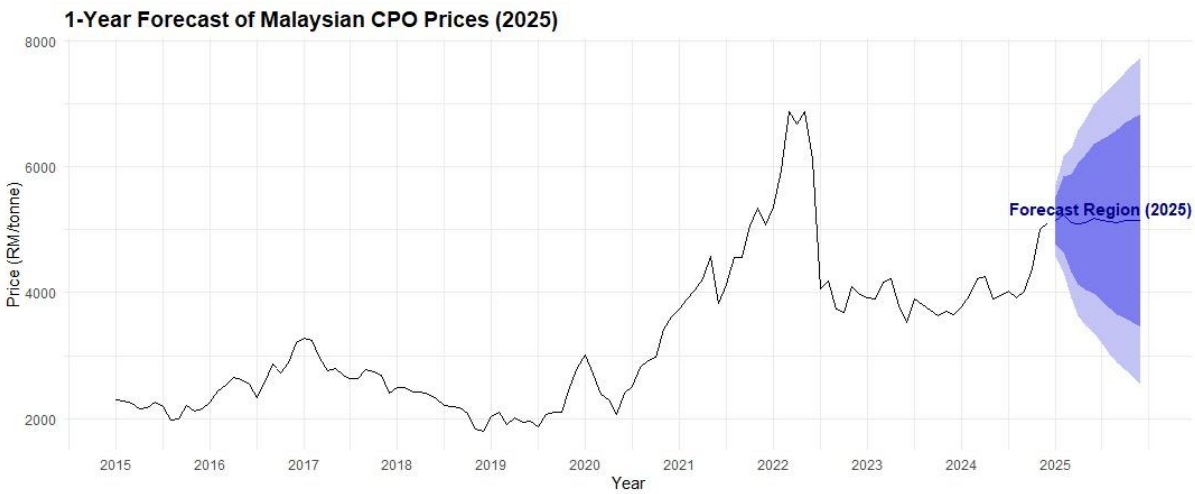

Figure 7. CPO Forecast for 1-year (2025)

The point forecast indicates a relatively stable price trajectory entering 2025, suggesting that the sharp price spikes observed during 2021–2022 are unlikely to recur under current market conditions. The widening fan-shaped confidence region reflects the increasing uncertainty inherent in longer-horizon predictions. This visualisation provides a clear and intuitive depiction of both expected price levels and the plausible range of future fluctuations, which is essential for policymakers and industry stakeholders who must plan under uncertainty.

Figure 8 presents an integrated visualisation of the actual Malaysian CPO prices, the fitted values from the training period (2015–2022), the validation forecasts for 2023–2024, and the final 12-month forecast for 2025. This unified plot allows direct comparison between model behaviour and observed data throughout the entire analysis period.

Figure 8. CPO Visualisation

During the training phase, the ARIMA(3, 1, 3) model reproduces the general direction and turning points of the actual series, although the fitted values appear smoother than the observed data because the model emphasises medium-term dynamics rather than short-term volatility. In the validation period, the model performs similarly, providing forecasts that follow the general level of the series, despite not capturing the small month-to-month fluctuations present in the actual data. This behaviour is consistent with a differenced ARIMA model without seasonal components, which tends to produce relatively stable short-term forecasts.

The 2025 forecast region (red line) shows a nearly flat trajectory. This stability is expected, once differencing removes trend and the model detects no persistent seasonal pattern, the resulting ARIMA(3,1,3) tends to generate level forecasts that revert toward the long-run mean of the differenced process. The shaded confidence region widens over time, indicating increasing uncertainty, but the central forecast remains relatively constant rather than showing any cyclical or trending behaviour.

Overall, this figure illustrates that while the ARIMA(3,1,3) model does not replicate the full volatility of the actual price series, it provides stable and reliable forecasts that capture the underlying level of the series. This reinforces its suitability as a forecasting model for short- to medium-term planning, especially when the objective is to obtain a baseline projection rather than reproduce high-frequency price variability.

4. Conclusion

This study set out to develop an accurate and interpretable forecasting model for Malaysian CPO prices using monthly data from 2015 to 2024. A comprehensive analysis combining descriptive statistics, seasonality testing, stationarity assessment, and model comparison led to several key findings.

First, the assessment of seasonal behaviour revealed that monthly CPO prices exhibit only weak seasonal patterns. The STL decomposition reported a low seasonal strength value of 0.2046. Together with the absence of a pronounced spike at lag 12 in the ACF, these results justified the exclusion of seasonal differencing and supported the use of non-seasonal modelling approaches.

Second, first-order differencing successfully transformed the original non-stationary series into a stationary process, as confirmed by the ACF and PACF behaviour and the outcomes of the ADF and KPSS tests. This enabled the formulation of candidate models of the form SARIMA$(a,1,c)(P,0,Q)[12]$. However, because seasonal effects were minimal, all competitive models effectively reduced to non-seasonal ARIMA specifications. This convergence further reinforces the conclusion that seasonal dynamics do not play a significant role in shaping short-term CPO price movements during the study period.

Third, the model selection and evaluation process was carried out using a systematic in-sample and out-of-sample framework to ensure robustness. All candidate ARIMA models were estimated through manual grid search, and the optimal specification was identified based on its in-sample AICc value and residual adequacy. Among the evaluated models, ARIMA(3,1,3) achieved the lowest AICc and demonstrated well-behaved residuals, with no evidence of autocorrelation as confirmed by diagnostic plots and the Ljung–Box test. Its forecasting capability was further assessed using the 2023–2024 out-of-sample dataset. The model produced satisfactory predictive accuracy in terms of RMSE, MAE, and MAPE, indicating that ARIMA(3,1,3) generalized effectively beyond the training period.

Overall, although seasonal structures were examined at earlier stages, the weak seasonal signal in the data resulted in the modelling framework naturally converging to a non-seasonal ARIMA form. Based on its strong balance of goodness of fit, residual adequacy, and forecast performance, ARIMA(3,1,3) was selected as the final model and used to generate the 12-month price forecast for 2025.

5. Acknowledgements

The authors would like to thank Universiti Teknologi Malaysia, UTM Encouragement Research grant (Q.J130000.3809.31J69) and Ministry of Higher Education, Fundamental Research Grant Scheme (FRGS/1/2024/STG06/UTM/02/14) for their partial financial support.

6. References

Amin, S. I. M., Rahman, A. A., Shukor, M. S. (2024). Determinants of Malaysian palm oil export efficiency: A gravity stochastic frontier model. *Geografia-Malaysian Journal of Society and Space 20*, issue 2: 145-155. e-ISSN 2682-7727. DOI: 10.17576/geo-2024-2002-09.

Mustapa, F. H. (2025). Malaysian palm oil price prediction using ARIMA – ARCH model. *Oil Palm Industry Economic Jurnal*. DOI: 10.21894/opiej.2025.01.

Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and Practice* (3rd ed.). OTexts. https://otexts.com/fpp3/.

Khalid, N., Hamidi, H. N. A, Thinagar, S., Marwan, N. F. (2018). Crude Palm Oil Price Forecasting in Malaysia: An Econometric Approach. *Jurnal Ekonomi Malaysia*, 51(3): 247-259. http://dx.doi.org/10.17576/JEM-2018-5203-19.

Lee, J. Y., Nguyen, T. T., Nguyen, G. G., Lee, J. Y. (2022). Towards Predictive Crude Oil Purchase: A Case Study in the USA and Europe. *Energies*, 15: 4003. https://doi.org/10.3390/en15114003.

Malaysian Palm Oil Council (MPOC) (2024). Monthly palm oil trade statistics. https://www.mpoc.org.my/category/monthly-palm-oil-trade-statistics/.

Rajakal, J. P., Ng, F. Y., Zulkifli, A., How, B. S., Sunarso, J., Ng, D. K. S., Andiappan, V. (2024). Analysis of current state, gaps, and opportunities for technologies in the Malaysian oil palm estates and palm oil mills towards net-zero emissions. *Heliyon*, 10: e30768. https://doi.org/10.1016/j.heliyon.2024.e30768.

Syarovy, M., Nugroho, A. P., Sutiarso, L., Suwardi, Muna, M. S., Wiratmoko, A., Sukarman, Primananda, S. (2023). Prediction of oil palm production using Recurrent Neural Network Long Short-Term Memory (RNN-LSTM). (Eds.): *ICoSIA* 2022, ABSR 29: 55–66. https://doi.org/10.2991/978-94-6463-122-7_6.

Tan, L. H., Lim, M. H., Lee, Y. L., Tan, O. K. (2023). Effective engineering noise control approaches using stochastic noise mapping simulation technique: a case study in palm oil mill. *The Institution of Engineers, Malaysia Journal*, 84(1).

Tardini, G. A, and Suharjito (2024). Selection of Modelling for Forecasting Crude Palm Oil Prices Using Deep Learning (GRU & LSTM). *Emerging Science Journal*, 8(3): 875-898. ISSN: 2610-9182. DOI: http://dx.doi.org/10.28991/ESJ-2024-08-03-05.

Tayib, S. A. M, Nor, S. R. M, Norrulashikin, S. M. (2021). Forecasting on the crude palm oil production in Malaysia using SARIMA Model. *Journal of Physics: Conference Series*, 1988: 012106. DOI:10.1088/1742-6596/1988/1/012106.