

Prediction of Crude Oil and Condensate Resources in Malaysia Using the ARIMA Approach

Rahaini Mohd Said^{1*}, Norhafizah Hussin¹, Siti Haryanti Haitol Anuar¹, Azyan Yusra Kapi², Sinarwati Binti Mohamad Suhaili³

¹ Faculty of Electronic and Computer, Technology and Engineering Universiti Teknikal Malaysia Melaka, 76100 Durian Tunggal Melaka, Malaysia

² Faculty of Computer and Mathematical Sciences, Universiti Teknologi MARA Johor Branch Pasir Gudang Campus, 81750, Pasir Gudang, Johor, Malaysia

³ Faculty of Cognitive Science and Human Development, Universiti Malaysia Sarawak, Jln Datuk Mohammad Musa, 94300 Kota Samarahan, Sarawak

DOI: <https://doi.org/10.47772/IJRISS.2026.100800592>

Received: 26 August 2026; Accepted: 31 August 2026; Published: 12 September 2026

ABSTRACT

Crude oil and condensate resources are hard to plan around without a reasonable forecast, and that matters most in economies like Malaysia, where upstream petroleum production underpins national energy security planning. This study uses the Box–Jenkins ARIMA approach on Malaysia's annual Total crude oil and condensate resources series (1990–2021, $n = 32$), taken from the Malaysia Energy Statistics Handbook 2023. After a literature review covering 20 Q1/Q2-indexed studies on ARIMA-based petroleum forecasting, the series was checked for stationarity, differenced once ($d = 1$), and sixteen candidate ARIMA($p, 1, q$) models ($p, q = 0–3$) were built from the ACF/PACF diagnostics and ranked by AIC and BIC. Rather than stopping at the AIC/BIC ranking, the final model choice came down to which specification actually forecast well against a held-out 2017–2021 test period (RMSE, MAPE). ARIMA(0,1,2) and ARIMA(2,1,0) had the lowest AIC/BIC, but it was ARIMA(0,1,1) that gave the best out-of-sample result (RMSE = 0.400, MAPE = 8.47%), so that is the model carried forward. Refit on the full series, it forecasts resources holding roughly flat at about 4.42 billion barrels through 2026, with a 95% interval that widens from [3.07, 5.77] in 2022 out to [1.15, 7.69] by 2026. The main takeaway is that for a short annual series like this, out-of-sample testing should outrank information criteria when picking a final model. Condensate-specific forecasting at the national level is still an open gap worth more attention.

Keywords: ARIMA; Box–Jenkins methodology; crude oil resources; condensate; time series forecasting

INTRODUCTION

Good forecasts of crude oil and condensate resources (reserves, production volumes, national output) matter for energy security planning, investment decisions, and long-term policy in oil-producing economies. That is especially true for a country like Malaysia, where upstream petroleum resources have moved through distinct phases of growth and decline over the past three decades. The Energy Commission (Suruhanjaya Tenaga), in its Malaysia Energy Statistics Handbook 2023, publishes the annual national crude oil and condensate resource total sourced from PETRONAS, covering 1990 to 2021 (19). Over those 32 years the total has shifted a lot, shaped by new discoveries and periodic reserve revisions. Combined with the usual swings in global oil markets, that makes this a genuinely interesting series to forecast.

Swings in global energy markets, plus the natural decline of both conventional and unconventional reservoirs, keep pushing demand for forecasting tools that hold up under uncertainty. ARIMA, the Autoregressive Integrated Moving Average model built on the Box–Jenkins framework, is still one of the go-to choices here, mainly because it is easy to interpret, does not need much data, and performs well once a series is made stationary

(1). Its linear structure also means it fits naturally into hybrid setups, paired with neural networks or similar learners that pick up whatever nonlinear pattern ARIMA misses on its own (1, 2).

Machine learning and deep learning methods (LSTM networks, Prophet, gradient boosting, various ensembles) have become popular alternatives, yet ARIMA keeps showing up in two roles: as a forecasting method in its own right, and as the baseline everyone compares newer models against (3, 4, 5). That staying power, even with ML adoption accelerating in petroleum forecasting research, is part of why this review looks specifically at how ARIMA and its variants have been used for crude oil and condensate resource prediction, as opposed to oil price forecasting, which most of the existing literature actually focuses on. The review is scoped around three things: first, pulling together recent Q1/Q2 studies that apply ARIMA or ARIMA-family models to oil and condensate resource forecasting; second, working out where the methodological and resource-type gaps sit in that literature; and third, placing this study, an ARIMA-based forecast of Malaysia's total crude oil and condensate resources from 1990 to 2021, within that wider body of work.

LITERATURE REVIEW

Classical ARIMA price studies. A large share of the ARIMA literature applies univariate Box–Jenkins models straight to crude oil prices at national scale, usually following the same identification-estimation-diagnostic cycle. Gasper and Mbwapbo fitted an ARIMA(0,1,1) model to Tanzanian crude oil prices after confirming stationarity through first-order differencing (6); Chamalwa, Abdulsalam, and Abbas took a similar Box–Jenkins ARIMA(1,1,1) approach for Nigerian crude oil prices, using standard diagnostics to settle on the final order (8). Both cases show ARIMA still holding up as a low-data national-level baseline, particularly where more data-hungry ML methods are not really an option. Alam et al. pushed the comparison further, benchmarking ARIMA against Simple Exponential Smoothing and K-Nearest Neighbours across oil, coal, and natural gas prices before and after COVID-19, and found that which model performed best actually changed between the two periods (7). Kim et al. went broader still, folding ARIMA into a supply-chain model alongside VAR and VECM to predict South Korea's crude oil import prices; ARIMA is just one piece of a larger multivariate system (15).

Hybrid ARIMA-ML models for production and resource forecasting. A second strand of the literature, more resource-focused, pairs ARIMA's linear forecasting with nonlinear ML components to pick up whatever pattern a purely linear model misses. Fan et al. built an ARIMA-LSTM setup for well production forecasting that specifically accounts for manual field operations, something purely statistical models tend to skip, and reported better accuracy than standalone ARIMA on RMSE and MAE (1). Dar et al. went the decomposition route instead, applying Ensemble Empirical Mode Decomposition to volatile crude oil price series before modelling the components with ARIMA and a feed-forward neural network, and found decomposition helped stabilise forecasts on highly volatile data (2). In oilfield production specifically, one integrated-learning study combined Singular Spectrum Analysis, ARIMA, Prophet, and neural networks into a single pipeline (14), and short-term single-well production forecasting has also been tackled with multi-agent hybrid models that use ARIMA mainly as a point of comparison rather than the main method (9).

ARIMA as comparative benchmark against machine learning. A number of studies use ARIMA mainly as the baseline that ML methods get measured against, rather than as the thing being studied. Ning, Kazemi, and Tahmasebi compared ARIMA, LSTM, and Prophet directly for oil production forecasting and found that which one wins depends heavily on forecast horizon and how volatile the data is; ARIMA stayed competitive at shorter horizons (3). Alkhamash's genetic-algorithm-tuned gradient boosting model for crude oil production used ARIMA as the comparison point to show what optimised ensemble learning could gain over it (4), and a similar setup, ARIMA against ANN and LSTM for crude oil price forecasting, shows up in related work too (5). Across these comparisons, ML tends to win on longer horizons or more nonlinear series, but ARIMA still holds its ground for shorter-horizon, lower-data situations, which is relevant here given Malaysia's series is only annual.

Reviews and adjacent resource types. Rahmanifard et al. put together a wide-ranging review of data-driven approaches, ARIMA-LSTM variants included, for production forecasting in unconventional reservoirs, useful for spotting methodological gaps across the wider field (10). Other work covers adjacent resource types: shale gas and shale oil production forecasting (11, 13), natural gas production and risk assessment (12), and gas-condensate reservoir forecasting (17), all pointing to active interest in resource-volume forecasting even where

ARIMA is not the main method used. Asafo et al.'s broader trend analysis of energy resources, including crude oil and lease condensate (16), points to growing interest in condensate specifically, though rigorous ARIMA work on condensate resources on their own is still thin.

Research gap. Put together, most of this literature leans heavily toward crude oil price forecasting and single-well or national production series, usually judged by AIC for model choice and RMSE, MAE, or MAPE for accuracy. What is missing is multi-decade forecasting of crude oil and condensates together at the national level, the kind of long-run series Malaysia's PETRONAS-sourced statistics actually provide. That gap, on the resource side specifically (condensates alongside crude oil), is what this study's ARIMA-based forecast of Malaysia's total crude oil and condensate resources over 1990–2021 is meant to address.

METHODOLOGY

Data

This study uses the annual Resources of Crude Oil and Condensates series from the Malaysia Energy Statistics Handbook 2023, compiled by Suruhanjaya Tenaga (Energy Commission Malaysia) from PETRONAS data (18). It reports Malaysia's aggregated national crude oil and condensate resources in billion barrels from 1990 to 2021 (32 annual observations), modelled here as a single national Total series rather than any sub-national breakdown, a scope decision discussed further in Section 4.8. Being annual, the series does not need SARIMA's seasonal terms, so a standard non-seasonal ARIMA(p,d,q) specification is enough, which fits how the annual-frequency ARIMA studies in the price and production literature were set up too (6, 8). With only 32 observations to work with, keeping the model order small is a real constraint carried through into Stage 1.

Stage 1: Model Identification

The series is first checked for stationarity using the Augmented Dickey-Fuller test, since Box–Jenkins requires a stationary series before order selection can happen, the same approach used in the national-level ARIMA price studies at (6) and (8). Given the likely trend over 32 years from cumulative depletion and reserve revisions, a first-order difference is applied where needed:

$$\Delta Y_t = Y_t - Y_{t-1}$$

where Y_t is the resource volume at year t and ΔY_t is the differenced series (a second difference, $\Delta^2 Y_t = \Delta Y_t - \Delta Y_{t-1}$, is applied if one difference is insufficient). Candidate orders (p, q) are then identified via the ACF and PACF of the differenced series, following the identification procedure used for hybrid production-forecasting studies (1, 9) and comparative ARIMA benchmarking (3). Remaining candidates are ranked by Akaike's Information Criterion:

$$AIC = -2 \ln(L) + 2k$$

where L is the maximized likelihood and k the number of estimated parameters. The lowest-AIC model is preferred, balancing fit against parsimony, which matters given the short sample and lines up with the AIC-based order-selection reported in comparative multi-commodity ARIMA forecasting (7) and integrated learning production-forecasting pipelines (14).

Stage 2: Model Estimation and Diagnostic Testing

Parameters of the identified ARIMA(p,d,q) model are estimated via maximum likelihood, as in decomposition-based ARIMA-neural network hybrids for crude oil price volatility (2) and structural ARIMA-VAR-VECM specifications for crude oil import prices (15). The fitted model is then diagnostically checked: residual autocorrelation (Ljung-Box test), coefficient significance, and normality. With only 32 observations, diagnostics also guard against over-parameterization. If diagnostics fail, the specification is revised and the process loops back to Model Identification; this identification-estimation-diagnosis cycle repeats until the model passes, a discipline reflected in reviews of data-driven production-forecasting approaches (10).

Stage 3: Model Application

Once validated, we use the model to forecast Malaysia's crude oil and condensate resources beyond 2021. Accuracy is assessed both in-sample and out-of-sample using:

$$\text{MSE} = \frac{1}{n} \sum_{t=1}^n (Y_t - \hat{Y}_t)^2 \quad \text{RMSE} = \sqrt{\frac{1}{n} \sum_{t=1}^n (Y_t - \hat{Y}_t)^2}$$

$$\text{MAE} = \frac{1}{n} \sum_{t=1}^n |Y_t - \hat{Y}_t| \quad \text{MAPE} = \frac{100}{n} \sum_{t=1}^n \left| \frac{Y_t - \hat{Y}_t}{Y_t} \right|$$

where $\hat{Y}_t$ is the forecast at year t and n is the number of forecast periods evaluated. MAPE's scale-independence makes it useful for comparing forecast accuracy across candidate models on the same series, and its use alongside AIC/BIC follows the accuracy-reporting convention used in multi-intelligent-agent production forecasting (9) and multi-commodity ARIMA comparisons (7).

RESULTS AND DISCUSSION

This section presents the empirical results of applying the Box-Jenkins ARIMA methodology to Malaysia's crude oil and condensate resources data (1990–2021). The analysis begins with a descriptive and exploratory examination of the series, followed by stationarity diagnostics via the autocorrelation (ACF) and partial autocorrelation (PACF) functions on both the original and first-differenced series. Candidate ARIMA($p,1,q$) models are then identified and evaluated using Akaike's Information Criterion (AIC), Bayesian Information Criterion (BIC), and out-of-sample forecast accuracy on a held-out test period (2017–2021), before the final model is selected and used to generate forecasts for 2022–2026.

Descriptive Statistics

Table 4.1 summarizes the descriptive statistics of Malaysia's crude oil and condensate resources series (1990–2021, $n = 32$), measured in billion barrels.

Table 4.1: Descriptive Statistics of Total Crude Oil and Condensate Resources (1990–2021)

Statistic	Total (Billion Barrels)
Mean	4.7058
Maximum	6.1150
Minimum	2.9430
Standard Deviation	0.9281
Skewness	-0.1906
Kurtosis	1.9627
Observations	32

On average, Malaysia held approximately 4.71 billion barrels of crude oil and condensate resources per year over the 1990–2021 period, ranging from a minimum of 2.943 billion barrels in 1990 to a maximum of 6.115

billion barrels in 1993, when a wave of new resource additions was first reported at the national level. The standard deviation of 0.9281 means annual resource levels deviate from the mean by about 0.93 billion barrels on average, a coefficient of variation of roughly 20% that indicates moderate year-to-year variability consistent with a resource series driven by periodic reserve revisions and new discoveries rather than smooth, continuous depletion. Skewness of -0.1906, being close to zero and well within the conventional ± 1 threshold, indicates the distribution of annual resource levels is approximately symmetric around the mean. Kurtosis of 1.9627, being below the normal-distribution benchmark of 3, indicates a platykurtic distribution, flatter and thinner-tailed than normal, consistent with the series' pattern of gradual, sustained shifts rather than sudden, isolated spikes. Finally, the dataset spans 32 annual observations, a relatively short series for ARIMA modeling that constrains the number of parameters ($p + q$) that can be reliably estimated and is the reason parsimony is emphasized during model order selection.

Trend Analysis

Figure 4.1 plots Malaysia's annual crude oil and condensate resources (1990–2021, billion barrels) against a fitted ordinary least squares (OLS) linear trend line, used as a preliminary visual check for trend and stationarity before ARIMA modeling.




Figure 4.1: Resources of Crude Oil and Condensates (Total, 1990–2021) with fitted linear trend

Figure 4.1 shows Malaysia's annual crude oil and condensate resources plotted against a fitted OLS linear trend line, $Y_t = 0.0450 \cdot (\text{Year}) - 85.63$. The positive slope of 0.0450 indicates resources increased by approximately 0.045 billion barrels per year over 1990–2021, and this upward trend is statistically significant ($p = 0.009$). However, the linear trend explains only about 21% of the year-to-year variation ($R^2 = 0.207$): the series moves through distinct phases: a sharp jump in 1992–1993, a decline through the late 1990s, a sustained rise from 2002 to 2012, and a renewed decline from 2016 onward, rather than following a single straight line. This gap between the actual series and the fitted trend is itself evidence that the series is non-stationary in its level but not simply a linear deterministic trend, supporting the use of differencing (rather than de-trending) for Box-Jenkins ARIMA modeling.

Stationarity Diagnostics: Original Data

Figure 4.2 and Figure 4.3 show the ACF and PACF of the original (level, undifferenced) Total series, computed on the training subset spanning 1990–2016 ($n = 27$ observations). The remaining 2017–2021 period is held out as the test set.




Figure 4.2: ACF of the Total series (original data, 1990–2016 training set)




Figure 4.3: PACF of the Total series (original data, 1990–2016 training set)

The lag-1 autocorrelation is 0.663, outside the 95% confidence band and by far the largest value in the correlogram, decaying gradually rather than cutting off sharply, the classic ACF signature of a non-stationary, trending series. A single dominant PACF spike appears at lag 1 (0.663), after which the partial autocorrelations drop sharply back within the band. This sharp PACF cutoff paired with the ACF's slow decay confirms that the original series requires differencing before ARIMA order identification, consistent with the Augmented Dickey-Fuller test on the level series ($t = -2.482$, $p = 0.120$ on the 1990–2016 training subset; $t = -2.721$, $p = 0.070$ on the complete 1990–2021 series), both readings borderline non-stationary at the 5% level

Stationarity Diagnostics: First-Differenced Data

Figure 4.4 and Figure 4.5 show the ACF and PACF of the first-differenced series ($\Delta Y_t = Y_t - Y_{t-1}$), 1990–2016 training set ($n = 26$ after differencing).




Figure 4.4: ACF of the first-differenced Total series (1990–2016 training set)




Figure 4.5: PACF of the first-differenced Total series (1990–2016 training set)

After first differencing, every autocorrelation from lag 1 through lag 12 falls within the 95% confidence band, resembling a white-noise correlogram and confirming that one round of differencing ($d = 1$) is sufficient (ADF $t = -4.572$, $p < 0.001$ on the training subset; $t = -5.062$, $p < 0.001$ on the complete series). Mirroring the ACF, every partial autocorrelation of the differenced series also falls within the confidence band. Because neither points clearly to a specific p or q , order selection was resolved using AIC and out-of-sample forecast accuracy across a grid of candidate $ARIMA(p,1,q)$ specifications.

Model Identification and Selection

Table 4.2 reports AIC, BIC, and out-of-sample forecast accuracy (RMSE, MAPE against the 2017–2021 held-out test set) for all candidate $ARIMA(p,1,q)$ models we fitted on the training set (1990–2016), with p and q each ranging from 0 to 3, plus a pure random-walk baseline $ARIMA(0,1,0)$. $ARIMA(0,1,1)$, the model we selected as final, is highlighted.

Table 4.2: Candidate $ARIMA(p,1,q)$ Models — AIC, BIC, and Out-of-Sample Forecast Accuracy

Model	AIC	BIC	RMSE	MAPE
ARIMA(0,1,2)	60.685	64.459	0.773	16.17%
ARIMA(2,1,0)	60.714	64.489	0.673	13.94%
ARIMA(0,1,0)	60.847	62.105	0.460	9.83%
ARIMA(0,1,3)	62.366	67.399	0.707	14.49%
ARIMA(1,1,2)	62.431	67.463	0.705	14.53%
ARIMA(2,1,1)	62.570	67.602	0.752	15.66%
ARIMA(3,1,0)	62.622	67.655	0.717	14.91%
ARIMA(0,1,1)	62.794	65.310	0.400	8.47%
ARIMA(1,1,0)	62.825	65.341	0.435	9.26%

ARIMA(1,1,1)	63.317	67.091	0.575	12.28%
ARIMA(3,1,1)	64.142	70.432	0.598	12.30%
ARIMA(1,1,3)	64.232	70.522	0.719	14.57%
ARIMA(2,1,2)	64.326	70.617	0.783	16.19%
ARIMA(3,1,2)	66.133	73.682	0.613	12.59%
ARIMA(2,1,3)	66.226	73.774	0.729	14.81%
ARIMA(3,1,3)	66.314	75.120	0.563	11.22%

Justification for Model Selection

ARIMA(0,1,1) was picked as the final model even though it does not have the lowest AIC or BIC. ARIMA(0,1,2) (AIC = 60.685) and ARIMA(2,1,0) (AIC = 60.714) both score better on those two measures. The reasoning comes down to what AIC and BIC actually measure versus what RMSE and MAPE measure: the first pair looks at in-sample fit, while RMSE and MAPE are computed against the 2017–2021 period the model never saw during estimation, so they measure real out-of-sample accuracy. Since the goal here is prediction, out-of-sample accuracy is the number that should decide things, and on that front ARIMA(0,1,1) wins clearly: lowest RMSE (0.400) and lowest MAPE (8.47%) of every model tested, well ahead of the AIC/BIC favourites, whose MAPE sat at 13.94%–16.17%. The reason for that gap: ARIMA(2,1,0) and ARIMA(0,1,2) fit the oscillating pattern in the 1990–2016 training data closely, but that pattern simply did not repeat in 2017–2021, where the series instead settled into a gradual decline. ARIMA(0,1,1)'s single MA(1) term is not statistically significant (coefficient = 0.103, standard error = 0.348, 95% CI [-0.579, 0.784], $p = 0.768$ on the full-series fit), so the model behaves close to a persistence-style, near-random-walk forecast, and that tracked the actual decline much better than the more heavily autoregressive alternatives did.

This point deserves a more critical look rather than a passing mention: with the MA(1) term statistically indistinguishable from zero, ARIMA(0,1,1) is, in practical terms, barely different from ARIMA(0,1,0), a pure random walk with zero estimated parameters. We checked this directly. A naïve forecast that simply repeats the last observed value (4.430 billion barrels from 2021) across the 2017–2021 test period gives RMSE = 0.460 and MAPE = 9.83%, identical to the random-walk model's own out-of-sample result, since ARIMA(0,1,0) is mathematically a random walk and produces the same forecast. ARIMA(0,1,1) improves on this naïve benchmark, but only modestly: RMSE falls from 0.460 to 0.400 (a 13% reduction) and MAPE from 9.83% to 8.47% (a 14% reduction), for the cost of one additional, statistically insignificant parameter. That the AIC/BIC-preferred models, which use two estimated parameters each, forecast markedly worse than even the zero-parameter naïve benchmark (MAPE 13.94% to 16.17% against 9.83%) reinforces the same conclusion from a different angle: added model complexity did not translate into added forecasting power on this series. Taken honestly, the finding is less that ARIMA(0,1,1) is a strong model and more that a simple persistence-style forecast is difficult to beat on this particular 32-observation annual series, and the modest gain ARIMA(0,1,1) does offer over the naïve benchmark should be read with that context attached.

Residual Diagnostics and Parameter Estimates

Beyond the AIC/BIC/RMSE/MAPE comparison, the ARIMA(0,1,1) model refit on the complete 1990–2021 series was checked against the standard Box–Jenkins diagnostic battery. The Ljung–Box test on the model residuals found no significant autocorrelation at either lag examined ($Q = 3.163$, $p = 0.788$ at lag 6; $Q = 5.151$, $p = 0.953$ at lag 12), which supports the adequacy of a single MA(1) term and rules out leftover serial correlation

the model failed to capture. The MA(1) coefficient itself was estimated at 0.103 (standard error = 0.348, $z = 0.295$, $p = 0.768$), with a 95% confidence interval of $[-0.579, 0.784]$ that comfortably spans zero, which is the formal basis for describing the coefficient as statistically insignificant rather than just numerically small.

One diagnostic did flag a genuine concern: a Shapiro-Wilk test on the residuals rejected normality ($W = 0.777$, $p < 0.001$), and the residual distribution shows positive skew (0.95) and heavy tails (kurtosis 8.18 per the Jarque-Bera statistic), driven mainly by the large 1992–1993 jump in the series when a wave of new national resource additions was first reported. Because the ARIMA confidence intervals reported in Table 4.3 rely on a Gaussian assumption for the residuals, this non-normality means those intervals, and by extension the $[3.07, 5.77]$ to $[1.15, 7.69]$ range projected for 2022–2026, are an approximation rather than an exact probabilistic statement, and should be interpreted with that caveat rather than as a precise 95% guarantee. This is discussed further as a study limitation in Section 4.8.

Final Model Forecast: ARIMA(0,1,1), 2022–2026

We refit the selected final model, ARIMA(0,1,1), on the full 1990–2021 series to generate the out-of-sample forecast for 2022–2026, reported in Table 4.3 and plotted in Figure 4.6.

Table 4.3: ARIMA(0,1,1) Point Forecasts and 95% Confidence Intervals, 2022–2026 (Billion Barrels)

Year	Forecast	Lower 95%	Upper 95%
2022	4.422	3.071	5.773
2023	4.422	2.411	6.433
2024	4.422	1.919	6.924
2025	4.422	1.510	7.334
2026	4.422	1.151	7.693




Figure 4.6: Final ARIMA(0,1,1) forecast, 2022–2026, with 95% confidence interval

The final forecast holds essentially flat at 4.422 billion barrels across the full 2022–2026 horizon, consistent with the model's near-random-walk behavior. The 95% confidence interval widens substantially with horizon length, from $[3.071, 5.773]$ in 2022 to $[1.151, 7.693]$ by 2026, reflecting the compounding uncertainty inherent in forecasting a short (32-observation) annual series five years ahead. This wide interval should be reported

alongside the point forecast rather than the point forecast alone, and is itself a limitation worth noting: the short training series constrains how precisely any ARIMA specification can pin down Malaysia's future crude oil and condensate resources beyond the near term.

The identical point forecast across all five years is a direct consequence of the ARIMA(0,1,1) structure rather than an artifact of the computation. An MA(1) model only carries information one step beyond the last observed residual: at $h = 1$, the forecast of the differenced series ΔY uses the model's MA(1) term applied to the final residual from 2021, but since that coefficient is not statistically significant (0.103, $p = 0.768$), its effect on the forecast is negligible. At $h \geq 2$, no further residual information is available, and with no drift term estimated in the model, the forecast of ΔY collapses to zero for every subsequent period. Because the level series is recovered by cumulating the differenced forecasts ($Y_t = Y_{t-1} + \Delta Y_t$), a string of near-zero ΔY forecasts after the first step means the level forecast simply repeats the 2022 value through to 2026, producing the flat line seen in Figure 4.6. This is a known property of low-order, non-drift ARIMA models applied to short annual series: they are well suited to short-horizon point prediction, as confirmed by the out-of-sample accuracy results in Section 4.5, but they carry no mechanism for projecting a continued trend beyond the first forecast step, which is precisely why the interval band, not the flat point line, should be read as the model's actual statement about 2023–2026.

DISCUSSION

ARIMA(0,1,1) beating the AIC/BIC favourites on out-of-sample accuracy lines up with what Ning, Kazemi, and Tahmasebi found for oil production forecasting: which model comes out ahead depends a lot on forecast horizon and how you evaluate it, and a simpler ARIMA specification can still hold its own at shorter horizons even against more flexible methods (3). The near-random-walk behaviour of the chosen model also matches Gasper and Mbwapbo's result for Tanzanian crude oil prices, where an ARIMA(0,1,1) model, the same structure selected here, came out on top after the same stationarity-testing-and-differencing routine. That suggests a low-order, weakly autocorrelated MA process shows up fairly often in short annual petroleum series (6). Chamalwa, Abdulsalam, and Abbas reached a similarly parsimonious Box–Jenkins model for Nigerian crude oil prices using standard AIC/BIC and diagnostic checks, which fits with why parsimony mattered so much here given only 32 observations to work with (8).

That said, leaning on out-of-sample RMSE and MAPE to override the AIC/BIC ranking points to a real limitation, one Alkhamash's genetic-algorithm gradient boosting study for crude oil production makes plain: a linear ARIMA model, however carefully diagnosed, can still lose out to more flexible ML approaches once there is enough data and compute to support them (4). With only 32 annual observations here, that kind of data-hungry approach was not attempted, though it is a reasonable next step for extending this work. Asafo et al.'s broader trend and forecasting work on energy resources, crude oil and lease condensate included, backs up what shows up here too: resource series tend to move through distinct multi-year phases rather than follow one deterministic trend, which is why differencing rather than de-trending was the right call for stationarity in this case (16). None of this sits apart from the wider ARIMA-based petroleum forecasting literature; it mostly confirms it, while still pointing at condensate-specific, national-level series like Malaysia's as an area that has not had much attention.

LIMITATIONS OF THE STUDY

Two limitations bear directly on how the forecast in Section 4.6 should be used, and both are tied to specific evidence produced during the analysis rather than stated as general caveats.

Forecast horizon and interval width. The five-year forecast (2022–2026) rests on a 32-observation training history, and the point forecast itself, as explained in Section 4.6, holds flat because the selected model's MA(1) term is not statistically significant. The 95% interval widens from [3.07, 5.77] billion barrels in 2022 to [1.15, 7.69] by 2026, and Section 4.5.1 adds a further reason to treat that interval as approximate rather than exact: the model residuals fail the Shapiro-Wilk normality test ($p < 0.001$), so the Gaussian-based interval formula understates or overstates the true uncertainty by an amount this study did not separately quantify. Readers should treat the 2022–2026 numbers as indicating a plateau in resources around current levels, bounded by a wide and only approximately Gaussian range, rather than as a precise year-by-year prediction.

Aggregated national scope. This study models Malaysia's crude oil and condensate resources as a single, aggregated national Total series covering 1990–2021. A more granular, sub-national analysis would need a disaggregated dataset that was outside the scope of what was available for this work, and is left as a direction for future research rather than reported here.

CONCLUSION

This study used the Box-Jenkins ARIMA approach to model and forecast Malaysia's crude oil and condensate resources, drawing on annual data from the Malaysia Energy Statistics Handbook 2023 (1990–2021, $n = 32$). The series turned out to be moderately variable but roughly symmetric (mean = 4.71 billion barrels, standard deviation = 0.93), moving through distinct historical phases rather than following one steady trend. A fitted linear trend only explained 21% of the variation, which backs that up. Stationarity testing found the original series non-stationary (ADF $p = 0.070$ on the complete series), but a single first-order difference was enough to fix that (ADF $p < 0.001$), matching the slow-decaying ACF and sharp PACF cutoff seen on the level series and their near-white-noise behaviour once differenced. Residual diagnostics on the fitted model found no leftover autocorrelation (Ljung-Box $p = 0.788$ at lag 6, $p = 0.953$ at lag 12) but did flag non-normal residuals (Shapiro-Wilk $p < 0.001$), a limitation carried through into how the forecast interval should be read.

Sixteen candidate ARIMA($p,1,q$) models were tested, judged both on in-sample criteria (AIC, BIC) and on out-of-sample accuracy (RMSE, MAPE) against a genuine 2017–2021 holdout. One thing stood out: the models AIC and BIC preferred, ARIMA(0,1,2) and ARIMA(2,1,0), actually forecast worse in practice (MAPE of 16.17% and 13.94%) than ARIMA(0,1,1), which ranked lower on AIC/BIC but had the best out-of-sample numbers of the whole set (RMSE = 0.400, MAPE = 8.47%). That gap is worth flagging for anyone doing applied ARIMA work: information criteria tell you how well a model fits what already happened, not how well it predicts what hasn't happened yet, so when the two disagree, out-of-sample accuracy should win for studies where prediction is actually the point.

On that basis, ARIMA(0,1,1) was picked as the final model, and once refit on the full 1990–2021 series it forecasts Malaysia's crude oil and condensate resources holding roughly flat around 4.42 billion barrels through 2026, with a 95% confidence interval that widens quite a bit across the horizon: from about [3.07, 5.77] in 2022 out to [1.15, 7.69] by 2026. That wide interval is a real limitation of the study: with only 32 annual observations to work from, the model simply cannot pin down a precise long-horizon number, so the forecast is better read as a plateau around current levels than an exact point estimate. A longer historical series, if one becomes available, would help; so would trying hybrid ARIMA-machine learning approaches that have improved accuracy elsewhere in the petroleum forecasting literature (1, 2, 4), and re-estimating the model periodically as new annual data comes in to catch any structural shifts in Malaysia's resource trajectory.

Ethical Considerations

This study uses publicly available, aggregated national resource statistics published by Suruhanjaya Tenaga (Energy Commission Malaysia); it does not involve human participants or animal subjects, and ethical approval was therefore not required.

Conflict of Interest: The author(s) declare no conflict of interest.

Data Availability

The annual crude oil and condensate resources data (1990–2021) used in this study are publicly available in the Malaysia Energy Statistics Handbook 2023, published by Suruhanjaya Tenaga (Energy Commission Malaysia). Derived series and model outputs are available from the corresponding author upon reasonable request.

REFERENCES

1. Fan, Y., et al. (2021). Well production forecasting based on ARIMA-LSTM model considering manual operations. *Energy*, 220, 119708.

2. Dar, A. A., et al. (2022). Forecasting crude oil prices volatility by reconstructing EEMD components using ARIMA and FFNN models. *Frontiers in Energy Research*, 10.
3. Ning, C., Kazemi, H., & Tahmasebi, P. (2022). A comparative machine learning study for time series oil production forecasting: ARIMA, LSTM, and Prophet. *Computers & Geosciences*, 164, 105126.
4. Alkhamash, E. H. (2022). An optimized gradient boosting model by genetic algorithm for forecasting crude oil production. *Energies*, 15, 6416.
5. AIMS Press Authors. (2022). Forecasting crude oil price using LSTM neural networks. *Data Science in Finance and Economics*.
6. Gasper, R., & Mbwambo, H. (2023). Forecasting crude oil prices by using ARIMA model: Evidence from Tanzania. *Journal of Accounting, Finance and Auditing Studies*, 9(2), 158–175.
7. Alam, M., et al. (2023). Forecasting oil, coal, and natural gas prices in the pre- and post-COVID scenarios. *Resources Policy*.
8. Chamalwa, H. A., Abdulsalam, D., & Abbas, U. (2024). Modeling and forecasting crude oil prices in Nigeria using ARIMA. *UMYU Scientifica*, 3(3), 313–321.
9. Yan, Z., et al. (2024). Short-term forecasting approach of single well production based on multi-intelligent agent hybrid model. *PLOS ONE*, 19(4).
10. Rahmanifard, H., et al. (2024). A comprehensive review of data-driven approaches for production forecasting in unconventional reservoirs. *Petroleum/reservoir engineering review journal*.
11. Liu, W., et al. (2023). Process of EUR prediction for shale gas wells based on production data. *Frontiers in Energy Research*.
12. Li, X., et al. (2023). Studies of natural gas production prediction and risk assessment. *Frontiers in Earth Science*.
13. Guo, J., Chen, Z., & Yan, W. (2023). Research progress and challenges on production forecast techniques of fractured horizontal wells in shale oil reservoirs. *Frontiers in Energy Research*.
14. ACS Omega Authors. (2023). Research on the application of integrated learning models in oilfield production forecasting. *ACS Omega*.
15. Kim, S., et al. (2023). A supply chain-oriented model to predict crude oil import prices in South Korea based on the hybrid approach. *Sustainability*, 15, 16725.
16. Asafo, S., et al. (2024). Trend analysis and forecasting of energy resources. *Heliyon*.
17. Kumar, A., et al. (2024). Rapid production forecasting for heterogeneous gas-condensate shale reservoir. *Geoenergy Science and Engineering*.
18. Suruhanjaya Tenaga (Energy Commission). (2023). *Malaysia Energy Statistics Handbook 2023*. Putrajaya, Malaysia: Suruhanjaya Tenaga.