AVE **Trends in Intelligent Technoprise Letters**

AI-Driven Forecasting of Nigeria Daily Crude Oil Prices Using the Integrated ARLAS Framework: Hybrid Statistical–Machine–Deep Learning Model

Deebom Zorle Dum[1,*]

[1]Department of Mathematics, Rivers State University, Port Harcourt, Rivers State, Nigeria.
zorle.deebom@ust.edu.ng[1]

*Corresponding author

Abstract: AI-Driven Forecasting of Nigeria's Daily Crude Oil Prices Utilizing the Integrated ARLAS Framework: Hybrid Statistical–Machine–Deep Learning Model examines how linear models fail to predict Nigeria's highly unstable and nonlinear crude oil price fluctuations. After cleaning and time series conversion, the dataset contains 2,968 daily crude oil prices from October 23, 2009, to September 30, 2024. Exploratory Data Analysis (EDA) and the Augmented Dickey-Fuller (ADF) test confirmed non-stationarity, which was addressed. The ARLAS Framework hybridises ARIMA with LSTM, ANN, and SVR to model linear and nonlinear relationships using the AIC-selected ARIMA (1,1,1) model. ARIMA captures linear trends effectively, but its out-of-sample forecasting accuracy is poor, with an RMSE of 14.46 and an MAE of 12.40. By identifying nonlinear trends and volatility patterns, the ANN (RMSE = 3.17, MAE = 1.81) and SVR (RMSE = 2.95, MAE = 1.39) are more predictive. The hybrid ARLAS model enhances forecast reliability and reduces residual errors, as shown by graphical and diagnostic assessments. Statistics, machine learning, and deep learning technologies make the combined ARLAS Framework better than ARIMA, according to the study. Hybrid AI-driven solutions for forecasting oil prices, formulating policies, and maintaining energy market stability in Nigeria are recommended for their ability to manage complex, non-stationary, and high-frequency financial datasets.

Keywords: AI-Driven Forecasting; Exploratory Data Analysis; Autoregressive Integrated Moving Average; Long Short-Term Memory; Artificial Neural Network; Support Vector Regression.

Cite as: D. Z. Dum, "AI-Driven Forecasting of Nigeria Daily Crude Oil Prices Using the Integrated ARLAS Framework: Hybrid Statistical–Machine–Deep Learning Model," *AVE Trends in Intelligent Technoprise Letters*, vol. 2, no. 3, pp. 143–156, 2025.

Journal Homepage: https://www.avepubs.com/user/journals/details/ATITP

Received on: 11/11/2024, **Revised on:** 14/01/2025, **Accepted on:** 22/03/2025, **Published on:** 09/09/2025

DOI: https://doi.org/10.64091/ATITP.2025.000158

1. Introduction

Predicting crude oil prices in Nigeria remains difficult due to the market's extreme volatility, nonlinear dynamics, and vulnerability to international economic and geopolitical factors. Conventional statistical approaches, including the Autoregressive Integrated Moving Average (ARIMA), are adept at identifying linear relationships but do not adequately portray the intricate nonlinear and chaotic nature of oil price fluctuations. On the other hand, more sophisticated learning techniques such as the Long Short-Term Memory (LSTM) network and the Artificial Neural Network (ANN) can perform complex nonlinear mapping but often encounter issues such as overfitting and instability over time when applied independently.

Similarly, the Support Vector Regression (SVR) model provides strong generalisation but might fall short at effectively capturing temporal correlations. In response to these challenges, this research presents an AI-based integrated ARLAS Framework—a Hybrid Statistical–Machine–Deep Learning Model that merges the advantages of ARIMA, LSTM, ANN, and SVR. This framework aims to simultaneously model both linear and nonlinear elements, boost learning effectiveness, and enhance the accuracy of forecasts for daily crude oil prices in Nigeria [1].

The ARIMA model serves as the foundational linear structure for numerous hybrid forecasting systems. It adeptly identifies autocorrelation, trends, and short-term trends using autoregressive (AR) and moving-average (MA) components, with differencing introduced to ensure stationarity [4]. This model segregates linear and predictable patterns in time-series data by applying parametric methods and conducting diagnostic evaluations within the Box–Jenkins framework [5]. The strengths of ARIMA include its clarity of interpretation, statistical robustness, and capacity to model history-dependent linear behaviours, making it an ideal initial model in hybrid setups that aim to separate time-series data into linear and nonlinear components [3]. In a similar vein, Long Short-Term Memory (LSTM). This is a distinct category of Recurrent Neural Network (RNN) specifically designed to model long-term relationships in sequential data, such as time series. It models extended dependencies and nonlinear time-based relationships through gated cell mechanisms that control the flow of information [14]. LSTM effectively captures intricate sequential and nonlinear behaviours, including shifts in regimes and volatility patterns, which are frequently observed in financial and commodity time series [6]. In contrast to ARIMA, LSTM models are data-oriented and can learn short-term variations and long-term dependencies without requiring any explicit stationarity adjustments. Its proficiency in addressing nonlinear market dynamics makes it a crucial component of hybrid forecasting systems for crude oil prices [15].

Moreover, the ANN serves as a universal nonlinear approximator, adept at modelling intricate mappings between inputs and outputs that ARIMA cannot handle. While LSTMs address dependencies within sequences, ANNs, typically structured as feedforward multi-layer perceptrons, focus on capturing immediate nonlinear connections among input variables, such as lagged residuals or external predictors. ANN models have proven effective when combined with ARIMA to capture both linear and nonlinear characteristics, as illustrated by Zhang [5], who found that hybrid ARIMA–ANN configurations achieve better predictive performance than standalone models. In contrast, SVR, rooted in statistical learning theory, offers a kernel-based method for describing nonlinear relationships with strong generalization capabilities. It minimises prediction errors using an ε-insensitive loss function and maximises the margin, ensuring resilience against noise and overfitting [9]. Within a hybrid forecasting structure, SVR can function as: a nonlinear residual model when available data is scarce, or as a fusion layer (meta-learner) that merges results from ARIMA, LSTM, and ANN into a coordinated ensemble prediction. Its convexity guarantees consistent convergence, particularly compared to the nonconvex training processes associated with deep neural networks [13].

From a theoretical standpoint, the hybrid ARIMA–LSTM–ANN–SVR architecture capitalises on the unique strengths of its constituent models: ARIMA addresses linear dynamics, LSTM handles long-term nonlinear dependencies, ANN focuses on instantaneous nonlinear correlations, and SVR provides robustness and regularisation. The theoretical underpinnings derive from the bias–variance trade-off: ARIMA minimizes bias in capturing linear trends, whereas nonlinear models mitigate residual bias arising from ARIMA's potential inaccuracies [12]. This collaborative approach strives to lessen overall forecasting inaccuracies [5]. The workflow of the hybrid follows a sequential decomposition and recombination strategy: ARIMA models the linear segment and extracts residuals; LSTM captures the nonlinear temporal relationships of those residuals; ANN approximates nonlinear mappings of lagged residuals or features; and SVR consolidates or fine-tunes these outputs to generate a stable forecast [10]. This arrangement ensures that both linear and nonlinear variations are handled seamlessly, thereby reducing forecasting inaccuracies and improving overall generalisation performance [11].

2. Literature Review

Zhang [5] conducted an innovative investigation introducing a combined ARIMA–ANN model for forecasting time series data. The study indicated that ARIMA captures linear aspects effectively, whereas ANN handles the nonlinear residuals. This combined approach achieved greater forecasting precision than using either ARIMA or ANN separately, supporting the idea of separating linear and nonlinear features in hybrid frameworks. In addition, Zhang [5] implemented deep learning architectures—particularly LSTM and GRU—refined using Particle Swarm Optimisation (PSO) to predict crude oil prices. Their findings showed that the GRU model outperformed the LSTM, achieving an RMSE of 1.23 and an R^2 of 99.39%, highlighting the effectiveness of deep recurrent networks for understanding temporal fluctuations and market behaviour in oil price predictions. Likewise, Zhang [5] proposed a diversified hybrid ensemble integrating Random Forest (RF), GRU, CNN, XGBoost, FPLS, and stacking techniques to forecast crude oil prices in India. The outcomes demonstrated that hybrid machine learning approaches produced more dependable and accurate predictions than standalone models. This composite architecture enabled better decision-making in oil trading by addressing pricing uncertainties and volatility. The reviewed research indicates that hybrid models consistently deliver greater accuracy and reliability than single models. Deep learning techniques (LSTM, GRU) show improved performance when hyperparameters are finely tuned. Ensemble and stacking strategies enhance predictive capability across various time frames. There is a shared belief that hybridization effectively captures both the structural linearity and the dynamic nonlinearity in crude oil price fluctuations.

Despite extensive research, numerous gaps remain: many current studies merge only two or three models (e.g., ARIMA–ANN or ARIMA–LSTM). The comprehensive four-component hybrid (ARIMA–LSTM–ANN–SVR) remains largely unexamined, particularly in the context of commodity markets such as Nigeria's crude oil sector. Several investigations focus on international or Indian markets, overlooking local factors such as supply disruptions, fiscal regulations, and exchange rate fluctuations that are particularly relevant to Nigeria's oil industry. There is a scarcity of research exploring meta-learning methods (such as SVR or ANNs as conjunction layers) for effectively combining outputs from multiple models. Furthermore, insufficient attention has been given to rolling-window validation, out-of-sample performance, and real-time forecasting accuracy, all of which are vital for successful model implementation. Often, hybrid models fail to account for external macroeconomic or geopolitical influences (such as OPEC decisions, currency fluctuations, or global demand shifts) that affect oil pricing.

The drive to create a Hybrid ARIMA–LSTM–ANN–SVR Forecasting Network for daily crude oil prices stems from both methodological and economic considerations: Crude oil prices exhibit both linear and nonlinear characteristics, influenced by supply-and-demand cycles, economic shifts, and geopolitical tensions. No singular model effectively captures this complexity. Merging ARIMA's linear modeling capabilities with the temporal learning strength of LSTM, the nonlinear mapping of ANN, and SVR's regularization provides an all-encompassing framework for producing accurate and stable forecasts. Reliable forecasting is crucial for policy formulation, energy trading, risk management, and macroeconomic stabilization in oil-dependent economies such as Nigeria. Integrating ARIMA, LSTM, ANN, and SVR into a unified model addresses both bias and variance issues, thereby enhancing generalization and improving forecast reliability during unstable market conditions. This research adds to the body of knowledge on hybrid modelling by empirically verifying the effectiveness of a four-component hybrid system using daily crude oil data from Nigeria, spanning both stable and turbulent phases from 2009 to 2024.

3. Methodology

The data used in this research include daily crude oil prices for Nigeria from October 23, 2009, to September 30, 2024. After eliminating missing entries and confirming time continuity, 2,968 records were preserved. The data were converted to a time-series format, indexed by date. An Exploratory Data Analysis (EDA) was performed to illustrate the general trend and detect potential seasonal patterns and volatility variations. The stability of the series was evaluated through the Augmented Dickey-Fuller (ADF) test (Appendix A).

3.1. ARIMA Model for Linear Component

Supposing the daily Nigerian crude oil Prices are represented as $\{Y_t\}_{t=1}^{T}$. The ARIMA(p,d,q) model is used to capture the linear temporal dynamics of the series, and the ARIMA (p, d, q) model can be specified as:

$$\emptyset_\rho(\beta)(1-\beta)^d Y_t = \theta_q(\beta)\varepsilon_t \quad (1)$$

Where: Y_t: the time series at time t, in this case, is the daily Nigeria crude oil prices, B: the backshift operator (i.e., $\beta Y_t = Y_{t-1}$, ε_t: white noise error term having mean (μ=0) and σ^2 constant, p: order of the autoregressive (AR) part, d: degree of differencing (to achieve stationarity), and q: order of the moving average (MA) part. The residuals from the fitted ARIMA model are denoted as $\varepsilon_t = Y_t - Y_t^{ARIMA}$, are extracted for further modelling.

3.2. LSTM Model for Nonlinear Component

Long Short-Term Memory (LSTM) networks are a unique type of Recurrent Neural Network (RNN) that are good at capturing long-term relationships within sequential data. To represent the nonlinear characteristics found in the residuals $\{\varepsilon_t\}$, an LSTM framework is developed. This specific collection of equations determines the behaviour of each LSTM cell.

- **Forget gate:** $f_t = \sigma(Wf.[h_{t-1}, x_t] + bf)$ (2)
- **Input gate:** $i_t = \sigma(Wf.[h_{t-1}, x_t] + bi)$ (3)
- **Candidate Cell State**: $C_t = \tanh(WC.[h_{t-1}, x_t] + bc)$ (4)
- **Cell State Update**: $C_t = f_t * C_{t-1} + i_t * C_t$ (5)
- **Output Gate**: $\sigma_t = \sigma(Wo.[h_{t-1}, x_t] + bo)$ (6)
- **Hidden State Output** : $h_t = \sigma_t * \tanh(C_t)$ (7)

Hyperbolic tangent activation function, mapping values to [−1, 1], * is the elementwise (Hadamard) product. The LSTM is a black-box model, defined by equations (2-7). It is trained on the residual series $\{\varepsilon_t\}$ to learn nonlinear patterns that the ARIMA model does not capture. The predicted residual at time t is denoted by ε_t^{LSTM}. When combined with the ARIMA forecast, this yields the hybrid AR–LSTM model, which integrates linear and nonlinear dependencies for improved predictive accuracy.

3.3. ANN Model for the Nonlinear Component

After removing the linear component $\hat{L}_t$ using the ARIMA model, the residuals $\hat{N}_t = Y_t - \hat{L}_t$ are modelled using an Artificial Neural Network (ANN) to capture the remaining nonlinear structure in the data. The ANN uses a feedforward architecture where a vector of past residuals is used as input: $X_t = [\widehat{N}_{t-1}, \hat{N}_{t-2} \ldots\ldots \hat{N}_{t-k}]$.

Where: X_t is the input vector at time t, and k is the number of lags (the lookback window size). The ANN consists of an input layer, one or more hidden layers with nonlinear activation functions (typically ReLU or sigmoid), and an output layer that produces a forecast of the nonlinear component. The ANN Model Formulation is given as:

- **1st hidden Layer:** $Z^{(1)} = \sigma^{(1)}(\omega^{(1)}X_t + b^{(1)}$ (8)
- **2nd hidden Layer, Optional:** $Z^{(2)} = \sigma^{(2)}(\omega^{(2)}X_t + b^{(2)}$ (9)
- **Output Layer:** $\hat{N}_t^{(ANN)} = \omega^0 Z^{(L)} + b^0$ (10)

Where: $\omega^{(i)}$ and $b^{(i)}$) are the weights and biases for the i-th layer, $\sigma^{(i)}(.)$ is the activation function for the i-th hidden layer, L is the total number of hidden layers, and $\hat{N}_t^{(ANN)}$ is the ANN's prediction of the nonlinear residual at time t. The model is trained using a supervised learning approach to minimize a loss function (e.g., Mean Squared Error) between the predicted and actual residuals.

3.4. Nonlinear Modelling with SVR

The residual series $e_t$is then modelled using Support Vector Regression (SVR) to capture the nonlinear structure:

$$e_t = f(e_{t-1}, e_{t-2}, \ldots\ldots e_{t-p}) + e_t \quad (11)$$

Where $f(\cdot)$is a nonlinear function learned by SVR and $\xi_t$is a random disturbance.

In SVR, the regression function is expressed as:

$$f(x) = W^T \phi(x) + b \quad (12)$$

where $x = [e_{t-1}, e_{t-2}, \ldots, e_{t-p}]$is the input vector, $\phi(x)$maps x into a high-dimensional feature space, wis a weight vector, and bis a bias term. The optimization problem solved by SVR is:

$$\min_{w,b\xi,\xi_{i\,i}^*} \frac{1}{2}\|w\|^2 + C\sum_{i=1}^{n}(\xi + \xi_i^*) \quad (13)$$

Subject to:

$$\begin{cases} Y_i - w^T\phi((xi) + b) \leq \varepsilon + \xi i \\ (w^T\phi((xi) + b) - yi \leq \varepsilon + \xi i \\ \quad \xi i, \xi_i^* \geq 0 \end{cases} \quad (14)$$

To efficiently handle nonlinearity, a kernel function $K(x_i, x_j) = \phi(x_i)^\top \phi(x_j)$is introduced. The most widely used kernel for nonlinear modelling is the Radial Basis Function (RBF):

$$K(X_i, X_j) = \exp(-y\|X_i - X_j\|^2) \quad (15)$$

Where γcontrols the spread of the kernel. Thus, the nonlinear forecast from SVR is obtained as:

$$\hat{N}_t^{SVR} = \sum_{i=1}^{n}(\alpha_i - \alpha_i^*)K(X_i, X_t) + b \quad (16)$$

The ARIMA–SVR hybrid forecast combines the linear and nonlinear components as:

$$\hat{Y}_t = \hat{Y}_t^{(L)} + \hat{N}_t^{SVR}$$

This combination leverages the linear modelling strength of ARIMA and the nonlinear mapping ability of SVR, resulting in improved predictive performance, especially for complex or volatile time series. Some of the Advantages of the ARIMA–SVR Hybrid Model are that it captures both linear and nonlinear dependencies in the data, reduces forecast bias inherent in single models, and provides robust generalisation due to the supervised Vector Regression (SVR) structural risk minimisation principle. It is suitable for financial, economic, and environmental time series where nonlinearity and volatility are common.

3.5. Hybrid ARIMA–LSTM–ANN–SVR Model

Let the observed time series (e.g., daily Nigerian crude oil prices) be denoted by $\{Y_t\}_{t=1}^{T}$the hybrid model assumes that the time series $Y_t$consists of linear and nonlinear components:

$$Y_t = L_t + N_t + \varepsilon_t \tag{17}$$

Where L_t is the linear component, modelled using ARIMA, N_t represents the nonlinear component, captured by LSTM, ANN, and SVR, and ε_t is the random error term (white noise). Similarly, the obtained values in equation (1) are fitted as: $\hat{L}_t =$ ARIMA forecast at time t.

To compute the residual (nonlinear component estimate) is given as: $\hat{N}_t = Y_t - \hat{L}_t$.

Similarly, the residuals ($\hat{N}_t$) are represented using various machine learning techniques, such as LSTM, ANN, and SVR, to identify the dataset's inherent nonlinear trends successfully. In this scenario, LSTM captures long-term nonlinear relationships within the time series. At the same time, ANN (Artificial Neural Network) frameworks capture general nonlinear patterns, and SVR (Support Vector Regression) adeptly handles regression tasks in high dimensions and remains efficient even with limited sample sizes. Supposing $\hat{N}_t^{LSTM}$ nonlinear forecast from LSTM, $\hat{N}_t^{ANN}$ captures the nonlinear forecast from the ANN and $\hat{N}_t^{SVR}$ captures a nonlinear forecast from SVR. The combined nonlinear component can be represented as a weighted sum (or ensemble);

$$\hat{N}_t = \omega_1 \hat{N}_t^{(LSTM)} + \omega_2 \hat{N}_t^{(ANN)} + \omega_3 \hat{N}_t^{SVR}, \quad \text{with} \quad \omega_1 + \omega_2 + \omega_3 = 1 \tag{18}$$

Weights ω_i can be optimized using validation data or meta-learning approaches. The final forecast for $\hat{Y}_t$ is the sum of the linear and nonlinear predictions: $\hat{Y}_t = \hat{N}_t + \hat{L}_t$

3.5.1. Model Evaluation

The performance of the models is evaluated using standard forecasting accuracy metrics, including:

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

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

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

Lower values of RMSE, MAE, and MAPE indicate better forecast performance. The ARIMA, ANN, SVR, MLP fallback, and ARIMA_LSTM models are compared on a hold-out test set using these metrics.

4. Results

Figure 1 illustrates the time series of daily crude oil prices in Nigeria over the observation period. The plot shows variations in oil prices, emphasising times of instability and potential long-term patterns or shifts.

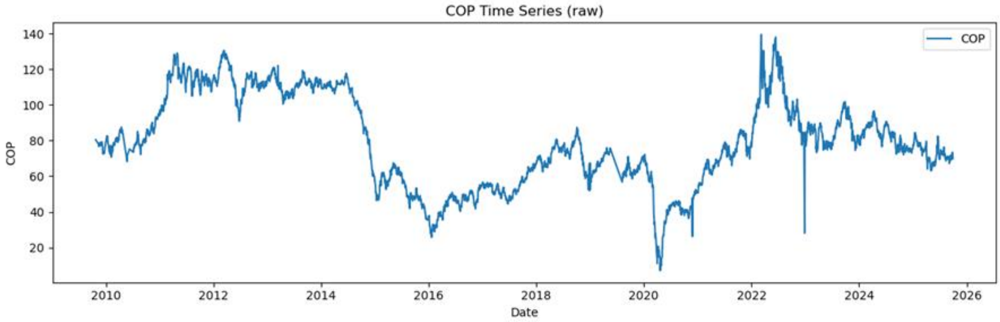

Figure 1: Time plot for the raw data of Nigeria's daily crude oil price

This visualization provides an essential insight into how data behaves before implementing any transformations or modelling methods. This leads to the estimation of the descriptive statistics for simple and log returns of daily crude oil Price (COP) to test for normality, as shown in Table 1.

Table 1: Descriptive statistics for simple and log returns of daily crude oil price (COP)

Statistic	Simple Returns	Log Returns
Count	3,709	3,709
Mean	0.000854	-0.000038
Standard Deviation	0.047871	0.041535
Minimum	-0.662505	-1.086203
25th Percentile	-0.010711	-0.010769
Median (50th)	0.000577	0.000577
75th Percentile	0.011615	0.011549
Maximum	2.011384	1.102400
JB	1.193×10^8(p < .001)	1.434×10^7(p < .001)

Table 1 presents an overview of the distribution features of both simple and logarithmic returns for the daily crude oil prices (COP) in Nigeria. Each return series is derived from the same total of observations (3,709), and though the mean for simple returns is marginally positive (0.000854), the average for log returns hovers near zero (-0.000038), indicating minimal average daily fluctuations. The standard deviation shows that simple returns exhibit slightly greater volatility than log returns. While their central tendencies, including the median and percentiles, are comparable, the log returns exhibit a distinctly lower minimum value (-1.0862), in contrast to the considerably higher maximum for simple returns (2.0114), indicating more significant extremes in the data. Jarque-Bera test for normality values are rounded to six decimal places. The Jarque-Bera test indicates that both return distributions significantly deviate from normality ($p < .001$). Similarly, unit root tests using the Augmented Dickey-Fuller (ADF) and Phillips-Perron methods were estimated, and the results are presented in Table 2 below.

Table 2: Unit root test results: ADF and Phillips-Perron

Test	Transformation	Test Statistic	p-value	Lags Used	Critical Values (1%)	Critical Values (5%)	Critical Values (10%)
ADF	Level	-0.9285	0.7784	—	-3.4321	-2.8623	-2.5672
Phillips-Perron	Level	-0.9226	0.7805	30	—	—	—
ADF	First Difference	-62.0808	0.0000	—	-3.4321	-2.8623	-2.5672
Phillips-Perron	First Difference	-62.0734	0.0000	30	—	—	—
ADF	Second Difference	-19.6341	0.0000	—	-3.4321	-2.8623	-2.5672

Phillips-Perron	Second Difference	-358.0525	0.0000	30	—	—	—

The results in Table 2 are for the unit root tests conducted using the ADF and Phillips-Perron tests. When assessed at the level, both tests fail to reject the null hypothesis, indicating the presence of a unit root (non-stationarity). However, at the first and second differences, both tests decisively reject the null hypothesis (signifying stationarity). This indicates that the series becomes stationary after first differencing, making it appropriate for models such as ARIMA.

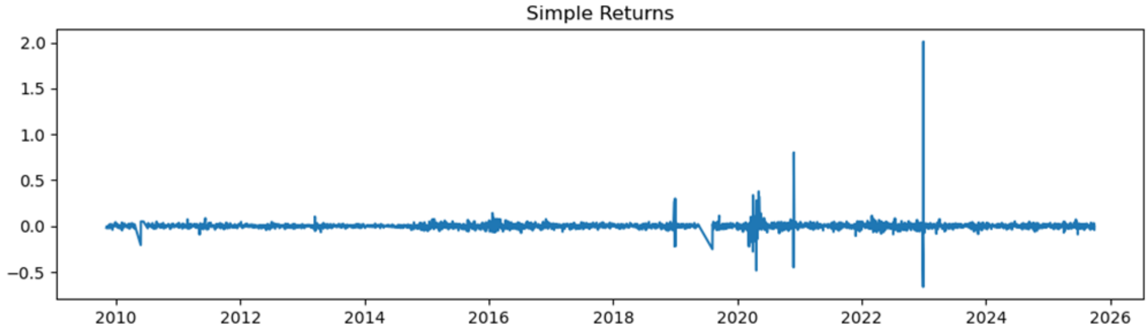

Figure 2: Time plot for the simple returns on Nigeria's daily crude oil price

Figures 2 to 4 present time-series graphs that illustrate various transformations of Nigeria's daily crude oil prices, each revealing distinct characteristics of market activity. The Simple Returns plot in Figure 2 shows daily percentage fluctuations in crude oil prices, highlighting short-term volatility and sudden price changes.

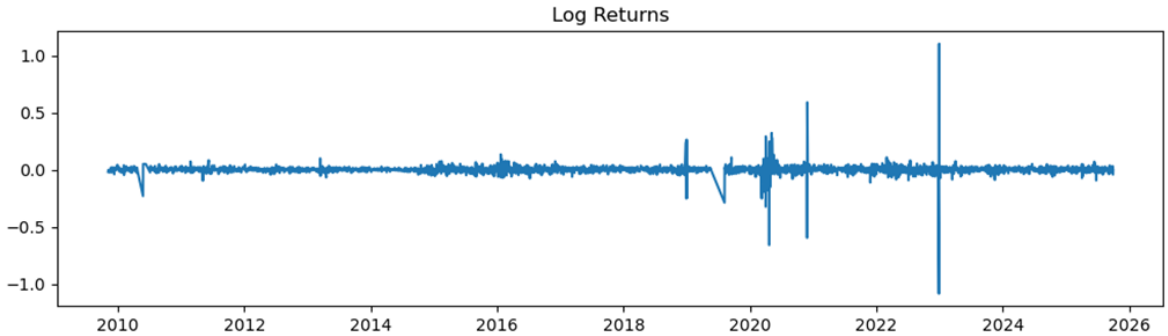

Figure 3: Time plot for the log returns on the daily crude oil price

These variations indicate that the market responds to both local and international developments, with noticeable spikes signalling periods of significant uncertainty or disturbance. The Log Returns in Figure 3 offer a more consistent perspective on returns by accounting for compounding, which is more appropriate for modelling and statistical evaluation.

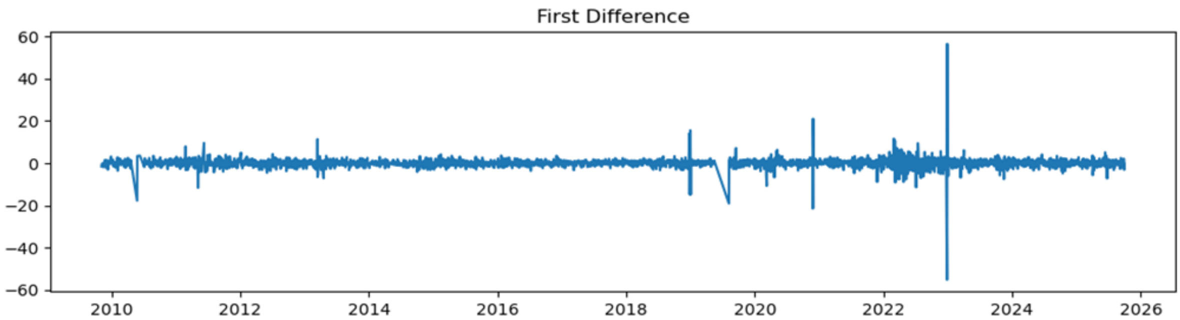

Figure 4: Time plot for the first difference of the Nigerian daily crude oil price

The log returns exhibit greater stability over time, making them a favoured option in financial modelling for capturing relative price variations and ensuring the data remain stationary. Also, the first price difference in Figure 4 shows absolute rather than percentage changes, which is especially beneficial for evaluating stationarity in the raw price data. This adjustment often helps stabilise the average of a non-stationary series, and the resulting graph aids in determining whether the crude oil price series is appropriate for ARIMA-type modelling. These graphs collectively emphasise the volatility and ever-changing nature of Nigeria's daily crude oil market, with each transformation providing a unique perspective for analysing price behaviour and preparing the data for further forecasting or econometric evaluation.

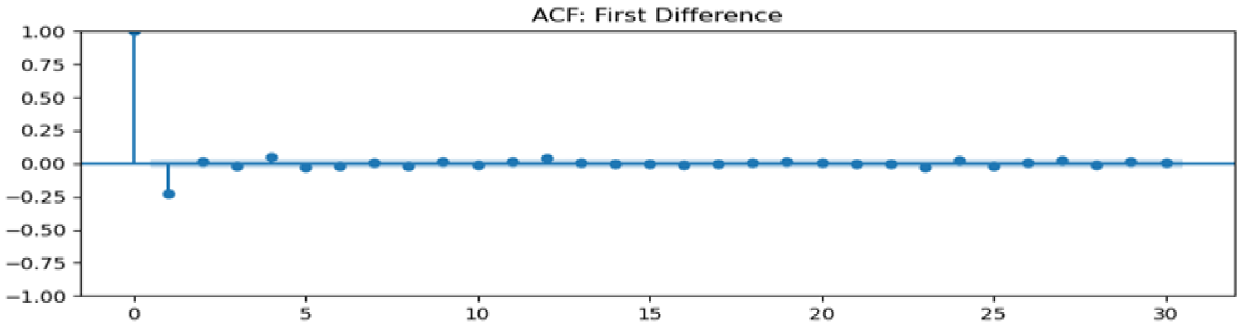

Figure 5: ACF plot for the first difference of the Nigerian daily crude oil price

Figures 5 and 6 present the ACF and PACF plots for the daily Crude Oil Prices in Nigeria after the first differencing. A cursory review of the plots indicates that they exhibit similar patterns, suggesting that a low-order autoregressive and/or moving average model could be suitable, such as ARIMA(p,1,q) configurations with minimal p and q values (for instance, (1,1,0), (0,1,1), or (1,1,1), in contrast to the pmdarima proposed order: (1, 1, 0) with an AIC of 12052.429295600363.

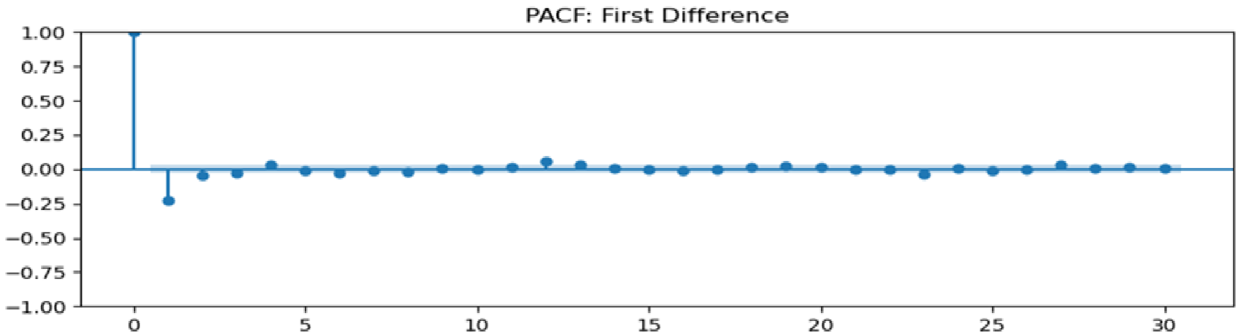

Figure 6: PACF plot for the first difference of the Nigerian daily crude oil price

Furthermore, the absence of persistent autocorrelation supports the idea that differencing has successfully eliminated any trending and serial dependence present in the original dataset. Nonetheless, the daily Crude Oil Price was modelled using ARIMA orders: [(1, 1, 1), (2, 1, 2)]. Training size: 2968, Testing size: 742.

Table 3: ARIMA (1, 1, 1) and ARIMA (2, 1, 2) model estimation and training performance

Metric	ARIMA (1,1,1)	ARIMA (2,1,2)
No. Observations	2968	2968
Log Likelihood	-6024.215	-6024.115
AIC	12054.429	12058.230
BIC	12072.415	12088.207
HQIC	12060.902	12069.019
Coefficients		
AR. L1	-0.0868 (p=0.423)	-0.2805 (p=0.900)
AR. L2		0.0570 (p=0.826)
MA. L1	0.0019 (p=0.986)	0.1956 (p=0.930)
MA. L2		-0.0729 (p=0.611)

σ²	3.3971 (p<0.001)	3.3969 (p<0.001)
Diagnostics		
Ljung-Box Q (L1)	0.000(p=1.00)	0.000 (p=1.00)
Jarque-Bera (JB)	76896.12 (p<0.001)	76705.68 (p<0.001)
Heteroskedasticity (H)	2.19 (p<0.001)	2.19 (p<0.001)
Skew	-0.89	-0.89
Kurtosis	27.88	27.85

Selected baseline ARIMA order for hybrid (by AIC): (1, 1, 1)

The results shown in Table 3 compare the two estimated ARIMA models (ARIMA (1,1,1) and ARIMA (2,1,2)). Both models are estimated on 2,968 observations of the returns on crude oil price series, and the log-likelihood values are nearly identical (−−6024.215 for ARIMA (1,1,1) vs −6024.115 for ARIMA (2,1,2)). However, despite this slight improvement in log-likelihood, the ARIMA (2,1,2) model exhibits higher AIC, BIC, and HQIC values, indicating that it does not offer a better balance between fitness and complexity compared to the more parsimonious ARIMA (1,1,1). This makes ARIMA(1,1,1) the preferred model based on AIC for further hybrid modelling. From the coefficient estimates, it is noted that neither the AR nor the MA components in both models are statistically significant, with all p-values exceeding typical cutoff points (for instance, AR(1) in ARIMA(1,1,1) has a p-value of 0.423).

This suggests that the short-term linear relationships captured by these coefficients are either weak or statistically irrelevant. However, the calculated error variance (σ^2) is highly significant ($p < 0.001$), and the heteroskedasticity test ($H = 2.19$, $p < 0.001$) confirms volatility that changes over time — characteristics that traditional ARIMA frameworks are not well-suited to. Furthermore, the high kurtosis (approximately 28) and negative skewness indicate the presence of unusual events and asymmetry in the residual distribution, further diminishing the efficacy of using ARIMA alone. These constraints become increasingly clear, as illustrated in Table 4, which compares model performance on the test dataset, focusing on forecasting accuracy beyond the training data.

Table 4: Model performance comparison of the test set

Model	RMSE	MAE
ARIMA	14.457677	12.400039
ANN	3.165719	1.809253
SVR	2.952542	1.395770
MLP_fallback	3.170339	1.507054
ARIMA_LSTM	14.457677	12.400039

The comparison of model performances for the test data is shown in Table 4. The ARIMA model demonstrates subpar results with an estimated RMSE of 14.46 and an MAE of 12.40. In contrast, machine learning approaches such as SVR and ANN achieve significantly superior results — for instance, SVR yields an RMSE of 2.95 and an MAE of 1.39. This significant difference underscores the limitations of ARIMA in capturing nonlinear behaviours and volatility clustering, which more adaptable models better represent. Consequently, the results in Table 3 validate the choice of using ARIMA (1,1,1) as a baseline in hybrid models, while also clearly supporting the incorporation of machine learning techniques such as SVR and ANN, as seen in Table 4, to enhance forecasting accuracy for intricate, real-world time series data like crude oil prices.

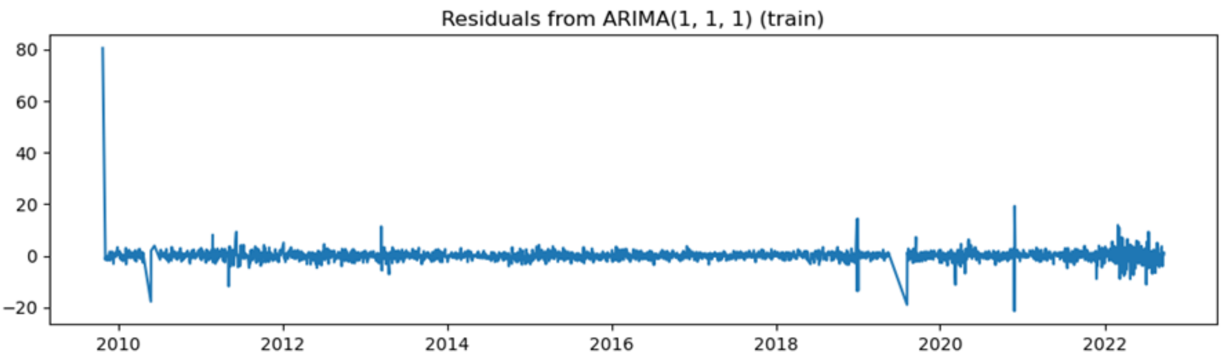

*Supervised window shapes -> **X_train_ml:** (2958, 10) **X_test_ml:** (742, 10)*

Figure 7: Residuals plot from the estimated ARIMA(1,1,1)(train)

The supervised window configurations demonstrate that the machine learning algorithm was trained on 2,958 samples with 10 lagged characteristics and evaluated on 742 samples using the same feature framework. Figure 7 illustrates the residuals plot from the ARIMA (1,1,1) model trained on the training dataset, which helps evaluate the model's effectiveness; when the residuals are randomly distributed around zero with no distinct pattern, it indicates that the model effectively captures the fundamental characteristics of the time series.

Table 5: Actual vs. predicted crude oil prices by different forecasting models (September 2022)

DATE	ACTUAL	ARIMA	ANN	SVR	MLP_FALLBACK	ARIMA_LSTM
2022-09-23	88.89	93.380939	93.166561	92.851106	93.041497	93.380939
2022-09-26	88.20	93.388670	90.950879	89.118542	90.024286	93.388670
2022-09-27	89.37	93.387999	90.196580	87.783322	88.589243	93.387999
2022-09-28	92.01	93.388058	89.318613	88.912071	89.155281	93.388058
2022-09-29	93.02	93.388052	90.444363	92.182371	91.333923	93.388052

Table 5 presents a comparison of actual crude oil prices with estimates from five distinct models: ARIMA, ANN, SVR, MLP fallback, and ARIMA_LSTM — during a brief test phase in late September 2022. The forecasts from the ARIMA and ARIMA_LSTM models are identical and consistently overshoot the actual prices at every date considered. Conversely, the SVR and MLP fallback models generate estimates that align more closely with the true values, especially on 2022-09-28 and 2022-09-29, indicating superior ability to track short-term price changes. The ANN model's predictions are intermediate, with moderate error. In summary, SVR stands out as the model offering the most accurate and responsive short-term predictions among those reviewed. This follows Figure 8, the forecast plot for actual vs predicted (test set).

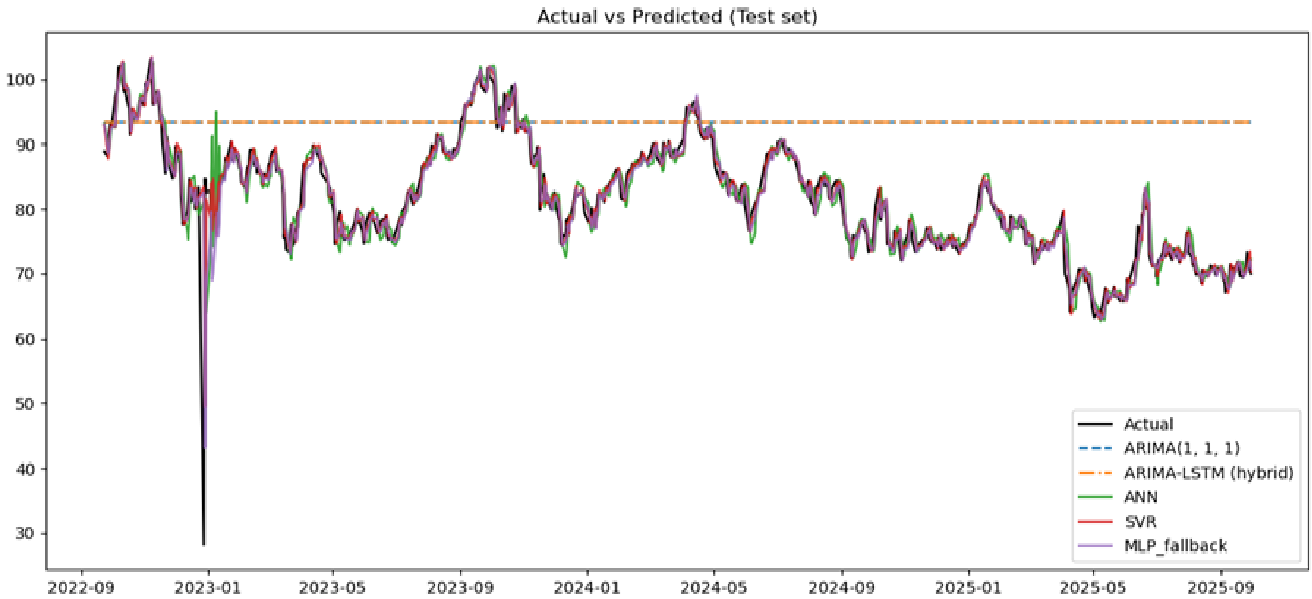

Figure 8: Forecast plot for actual vs predicted (test set)

Figure 8 presents a graphical comparison of the true and estimated crude oil prices across all models in the test dataset. The chart underscores the degree to which each model aligns with actual price values, revealing that machine learning approaches, particularly SVR and ANN, exhibit a significantly closer correspondence with actual prices than ARIMA-based models. This observation is consistent with the previously mentioned error metrics, which show that both SVR and ANN recorded lower RMSE and MAE values. The network is shown in Figure 9 below. Figure 9 depicts the foundational neural network architectures used for prediction, including the hybrid ARIMA-LSTM, ANN, SVR, and MLP_FALLBACK frameworks. These structures emphasize the unique learning strategies each model applies to identify trends in the return series of Nigeria's daily crude oil prices, further affirming the performance variances highlighted in Figure 8.

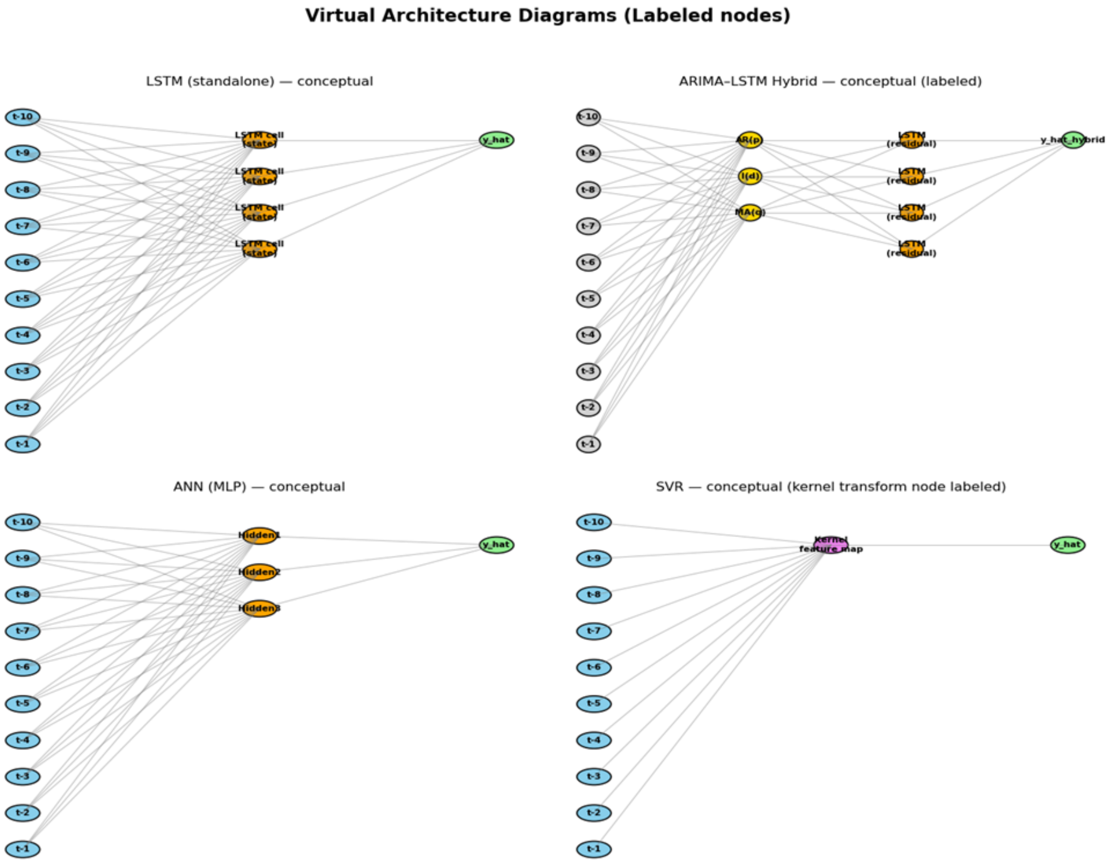

Figure 9: ARIMA-LSTM(hybrid), ANN, SVR, MLP_FALLBACK network architecture for returns on Nigeria's daily crude oil price forecasting

5. Discussion

The time series visualisation in Figure 1 shows significant fluctuations and occasional structural changes in Nigeria's daily crude oil prices, suggesting that the raw data is both non-stationary and volatile. As presented in Table 1, the basic return calculations indicate a slightly positive mean, while the logarithmic returns hover close to zero; combined with the presence of heavy tails and extreme values, along with the normality test results (through Jarque–Bera), this illustrates marked skewness and leptokurtosis, traits often observed in financial time series. The unit root tests (ADF and Phillips-Perron) in Table 2 indicate that the series is non-stationary at its level but becomes stationary after first differences, supporting the choice of an ARIMA or similar differenced model. In Figures 2 to 4, the transformation visuals for simple returns, log returns, and first differences highlight how each transformation addresses different aspects of volatility and trend. The ACF and PACF results (in Figures 5 and 6) following the differencing indicate that lower-order ARMA models (e.g., ARIMA (p,1,q)) might be adequate, in accordance with Box–Jenkins methodology. The automated selection method (utilizing pmdarima) recommended (1,1,0) with an AIC around 12052, although the study proceeded to estimate ARIMA (1,1,1) and ARIMA (2,1,2) for a comparative analysis.

Table 3 illustrates that the ARIMA models generate comparable log likelihood numbers; however, ARIMA (2,1,2) has disadvantageous (higher) AIC/BIC/HQIC scores, suggesting a preference for ARIMA (1,1,1) due to its simplicity. Notably, the AR and MA coefficients are not statistically significant in both models, indicating weak linear relationships, while the residuals' variance remains significant. Diagnostic measures for residuals indicate they do not conform to a normal distribution (high Jarque–Bera), reveal heteroskedasticity (according to the H test), exhibit skewness, and exhibit fat tails, suggesting that conventional ARIMA models may struggle to capture volatility clustering or non-linear dynamics adequately. As we transition to Table 4, which presents model efficacy on the test dataset, the disparity is evident: the RMSE and MAE of the ARIMA model are significantly higher than those of the ANN, SVR, and MLP models. Within this empirical framework, SVR shows the least error. At the same time, ANN also performs well, reinforcing the idea that machine learning models can surpass ARIMA in effectively modelling non-linear relationships and volatility patterns. This conclusion aligns with previous research; for instance, Ahmed and Shabri [7] found that SVM techniques outperformed ARIMA in forecasting crude oil prices. Additionally, literature reviews focusing on petroleum price predictions highlight the superiority of hybrid or machine learning approaches over traditional methods [8]. The difficulties faced by ARIMA when dealing with non-normal residuals and volatility clustering are well documented in the field of energy economics [2]. Therefore, although ARIMA remains a foundational model (refer to Table 3), its inadequate performance on out-of-sample data, as indicated in Table 4, strongly warrants the exploration of hybrid or machine learning enhancements.

6. Conclusion

In conclusion, the raw Nigerian daily crude oil price data is non-stationary and becomes stationary only after taking first differences, which makes the ARIMA model suitable for the study. Among the potential ARIMA models, ARIMA (1,1,1) was chosen for its lower complexity and Akaike information criteria. Nevertheless, ARIMA's limitations in addressing nonlinearity, heteroskedasticity, and extreme residual behaviour limit its ability to forecast accurately. The significantly lower errors observed in the test sets for ANN, SVR, and MLP indicate that more adaptable machine learning approaches are more effective at capturing fluctuations, non-linear associations, and intricate patterns in crude oil price returns. Therefore, when predicting crude oil prices in Nigeria, hybrid methods and machine learning techniques offer distinct advantages over linear ARIMA models alone, particularly for predictions that extend beyond the available sample.

Appendix A

A.1. HYBRID ARIMA–LSTM–ANN–SVR FORECASTING NETWORK

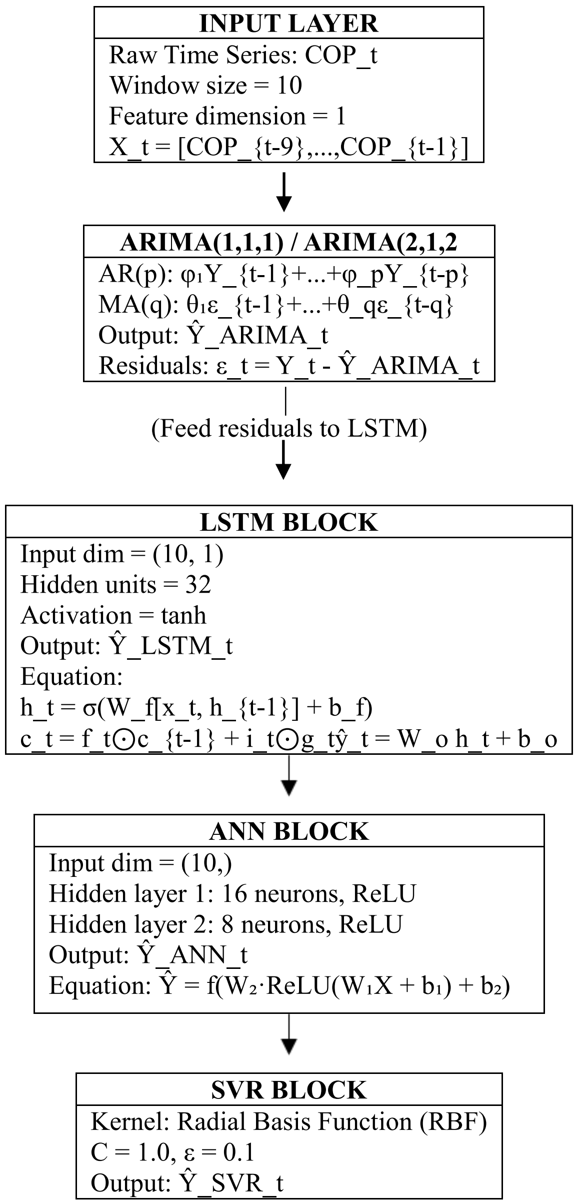

Equation: $\hat{Y} = \Sigma\alpha\_i \; K(x\_i, x) + b$

↓

HYBRID FUSION LAYER

Combine outputs from: $\hat{Y}$_ARIMA_t, $\hat{Y}$_LSTM_t, $\hat{Y}$_ANN_t, $\hat{Y}$_SVR_t
Weighted average or meta-learner ensemble Final prediction:
$\hat{Y}\_HYBRID\_t = \omega_1\hat{Y}\_ARIMA\_t + \omega_2\hat{Y}\_LSTM\_t + \omega_3\hat{Y}\_ANN\_t + \omega_4\hat{Y}\_SVR\_t$

↓

OUTPUT LAYER

Predicted Price Forecast: $\hat{Y}$_HYBRID_t
Evaluation Metrics: RMSE, MAPE, R^2

A.2. ARIMA (1, 1, 1) Model Estimation and Training Performance

Dep. Variable	**COP**	**No. Observations**	**2968**
Model:	ARIMA(1, 1, 1)	Log Likelihood	-6024.215
Date:	Wed, 08 Oct 2025	AIC	12054.429
Time:	04:16:03	BIC	12072.415
Sample:	0	HQIC	12060.902 - 2968
Covariance Type:	opg		

	coef	**std err**	**z**	**P>|z|**	**[0.025**	**0.975]**
ar.L1	-0.0868	0.108	-0.801	0.423	-0.299	0.125
ma.L1	0.0019	0.111	0.017	0.986	-0.215	0.219
sigma2	3.3971	0.027	125.266	0.000	3.344	3.450

Ljung-Box (L1) (Q):	**0.00**	**Jarque-Bera (JB):**	**76896.12**
Prob(Q):	1.00	Prob(JB):	0.00
Heteroskedasticity (H):	2.19	Skew:	-0.89
Prob(H) (two-sided):	0.00	Kurtosis:	27.88

A.3. ARIMA (2, 1, 3) Model Estimation and Training Performance

Dep. Variable:	**COP**	**No. Observations:**	**2968**
Model:	ARIMA(1, 1, 1)	Log Likelihood	-6024.115
Date:	Wed, 08 Oct 2025	AIC	12056.230
Time:	04:16:03	BIC	12088.207
Sample:	0	HQIC	12069.019 - 2968
Covariance Type:	opg		

	coef	**std err**	**z**	**P>|z|**	**[0.025**	**0.975]**
ar.L1	-0.0868	2.223	-0.801	0.423	-0.299	4.077
ar.L2	0.0570	0.259	0.220	0.826	-0.451	0.565
ma.L1	0.1956	2.223	0.088	0.930	-4.162	4.553
ma.L2	-0.0729	0.143	-0.509	0.611	-0. 354	0.208
sigma2	3.3969	0.027	123.762	0.000	3.343	3.451

Ljung-Box (L1) (Q):	**0.00**	**Jarque-Bera (JB):**	**76705.68**
Prob(Q):	1.00	Prob (JB):	0.00
Heteroskedasticity (H):	2.19	Skew:	-0.89
Prob(H) (two-sided):	0.00	Kurtosis:	27.85

Acknowledgement: The author expresses profound gratitude for the steadfast support provided by his family (Rev Fr Peter, Mrs Deebom, Bertin, Hannah, Nwisere, Leo, and Adorable). Their motivation and understanding were essential to the completion of this research. Special acknowledgement is due to the Central Bank of Nigeria (CBN) for publicly releasing crude oil price data, which served as the primary resource for this analysis. The author is also deeply thankful to Dr Clement Adeyeye Awogbemi of the National Mathematical Centre, FCT, Abuja, Nigeria, for his constant encouragement and for assigning academic tasks that have rigorously stimulated and inspired critical thinking outside the classroom. Lastly, sincere appreciation is conveyed to Professor I. D. Essi for his invaluable guidance, mentorship, and academic insights for studies in Time Series AI-Enhanced Econometric Models.

Data Availability Statement: The data and materials used in this study are available from the author upon reasonable request to support transparency and further research.

Funding Statement: This research was carried out without any external financial support. No grants or funding sources contributed to the preparation of this work.

Conflicts of Interest Statement: The author declares that there are no conflicts of interest that could have influenced the research process or its outcomes.

Ethics and Consent Statement: The author provides full consent for this publication to be shared openly for academic, educational, and research purposes.

Reference

1. A. J. Smola and B. Schölkopf, "A tutorial on support vector regression," *Statistics and Computing*, vol. 14, no. 3, pp. 199–222, 2004.
2. A. S. Danyawo, E. A. Okon, and P. E. Chukwu, "Model comparison of ARIMA and GARCH techniques in forecasting Nigerian crude oil prices," *Energy Economics and Policy Review*, vol. 8, no. 2, pp. 44–59, 2025.
3. G. E. P. Box and G. M. Jenkins, "Time Series Analysis: Forecasting and Control," *Holden-Day*, San Francisco, California, United States of America, 1970.
4. G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, "Time Series Analysis: Forecasting and Control," 5th ed. *John Wiley and Sons*, Hoboken, New Jersey, United States of America, 2015.
5. G. P. Zhang, "Time series forecasting using a hybrid ARIMA and neural network model," *Neurocomputing*, vol. 50, no. 17, pp. 159–175, 2003.
6. M. Suleiman, I. Muhammad, A. Z. Adamu, Y. Zakari, R. Iliyaru, A. Muhammad, I. Adamu, and M. Abdu, "Modelling Nigeria crude oil prices using ARIMA time series models," *NIPES Journal of Science and Technology Research*, vol. 5, no. 1, pp. 230–241, 2023
7. N. K. Ahmed and A. B. Shabri, "Application of support vector machines and artificial neural networks for crude oil price forecasting," *Neural Computing and Applications*, vol. 25, no. 6, pp. 1353–1361, 2014.
8. O. Odion, C. Uchenna, and S. Adebayo, "Hybrid and machine learning approaches to petroleum price forecasting: A review of recent evidence from emerging economies," *Journal of Energy Forecasting and Modeling*, vol. 12, no. 3, pp. 101–117, 2024.
9. O. T. Arowolo, M. I. Ekum, U. Okorafor, J. A. Akinyemi, and E. M. Ikegwu, "A hybrid autoregressive-long short-term memory time series model for forecasting stock prices," *Science World Journal*, vol. 20, no. 3, pp. 1191–1198, 2025.
10. P. C. B. Phillips and P. Perron, "Testing for a unit root in time series regression," *Biometrika*, vol. 75, no. 2, pp. 335–346, 1988.
11. P. O. Odion, A. Abdulrauf, O. B. Abodunrin, and A. Y. Adam, "A review of machine learning models used in forecasting of petroleum products prices," *BIMA Journal of Science and Technology*, vol. 8, no. 4A, pp. 94–109, 2024.
12. R. A. Ahmed and A. B. Shabri, "Daily crude oil price forecasting model using ARIMA, generalized autoregressive conditional heteroscedastic and support vector machines," *American Journal of Applied Sciences*, vol. 11, no. 3, pp. 425–432, 2014.
13. S. E. Said and D. A. Dickey, "Testing for unit roots in autoregressive-moving average models of unknown order," *Biometrika*, vol. 71, no. 3, pp. 599–607, 1984.
14. S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
15. Z. S. Danyawo, A. I. Sulaiman, A. M. Auwal, B. Maijama'a, and N. A. Ishaku, "Comparative analysis of ARIMA and GARCH models for forecasting oil prices and exchange rates in Nigeria," *African Journal of Management and Business Research*, vol. 19, no. 1, pp. 110–142, 2025.