Sinkron : Jurnal dan Penelitian Teknik Informatika
Volume 10, Number 3, July 2026
DOI : https://doi.org/10.33395/sinkron.v10i3.16410

e-ISSN : 2541-2019
p-ISSN : 2541-044X

Target-Characteristic-Aware Forecasting of Fuel–Equity Rolling Correlations: Evidence That PCA and Persistence Outperform Hybrid Deep Learning Models

Chrisnawan Prastya Atmaja[1)*], Ferian Fauzi Abdulloh[2)]
[1,2)]Universitas AMIKOM Yogyakarta, Indonesia
[1)]Satyachrisna926@students.amikom.ac.id , [2)]ferian@amikom.ac.id

Submitted :Jun 18, 2026 | **Accepted** : Jul 22, 2026 | **Published** : Jul 28, 2026

Abstract: Background: Forecasting dynamic fuel–equity market correlations is important for financial forecasting because fuel price movements can affect market risk, investor sentiment, and cross-market stability. However, most previous studies focus on direct price or volatility prediction, while the forecasting of rolling correlations between multiple fuel types and global equity indices remains less explored. Objective: This study analyzes and predicts time-varying correlations between global fuel prices and major stock market indices by comparing baseline, PCA-based, standalone deep learning, and hybrid deep learning models. Methods: The proposed framework applies data preprocessing, return transformation, 60-day rolling correlation construction across three fuel variables (petrol, diesel, LPG) and five stock indices (S&P 500, NASDAQ, FTSE 100, Nikkei 225, IHSG/JKSE), normalization, Principal Component Analysis (PCA), sequence generation, and comparative evaluation of 17 forecasting models. Model performance was measured using RMSE, MAE, R-squared (R^2), and Directional Accuracy. Results: The PCA Model (PCA + Linear Regression) achieved the lowest RMSE of 0.016909 and the highest R^2 of 0.967045, while the Persistence Model produced the lowest MAE of 0.002293 and the highest Directional Accuracy of 97.767280%. Hybrid deep learning models showed higher errors and negative R^2 values, indicating weaker performance on smooth and persistent rolling correlation targets. Conclusion: The findings show that architectural complexity does not necessarily improve forecasting performance when the target series is smooth and persistent. This study contributes empirical evidence and a replicable comparative framework showing that model selection in financial time-series forecasting should consider target characteristics, particularly smoothness and temporal persistence, rather than relying solely on complex deep learning architectures.

Keywords: Attention-LSTM, CNN-LSTM, Financial Forecasting, Fuel-Equity Correlation, PCA, Persistence Model, Rolling Correlation

INTRODUCTION

The relationship between global fuel prices and stock market performance has become an increasingly critical issue in financial forecasting. Energy price fluctuations can influence production costs, inflation, investor sentiment, and cross-market stability, creating dynamic interdependencies between fuel and equity markets that change over time. Understanding and forecasting these time-varying correlations is important for portfolio risk management, hedging strategies, and cross-market investment decision-making.

Recent forecasting studies have widely adopted machine learning and deep learning methods for financial time-series prediction. Models such as Long Short-Term Memory (LSTM), CNN-LSTM, and Attention-based architectures have demonstrated strong performance in capturing sequential patterns and nonlinear relationships in price and volatility series (Sen & Dutta Choudhury, 2024; Lahmiri, 2026; Chung et al., 2025). Principal Component Analysis (PCA) has also been commonly applied to reduce feature dimensionality while preserving important variance structures in high-dimensional financial datasets (Cao et al., 2016; Wang, 2024).

*name of corresponding author

However, several fundamental gaps remain in the existing literature. First, most studies focus on direct price or volatility prediction rather than forecasting the dynamic correlation structure between fuel and equity markets. Correlation forecasting presents a different challenge because rolling-correlation targets tend to be smoother, bounded, and highly autocorrelated. These characteristics may reduce the advantage of complex deep learning architectures compared with simpler models. Second, although hybrid deep learning models have been widely evaluated on volatile price series, their performance on persistent rolling-correlation targets has not been sufficiently examined. It remains unclear whether architectural complexity provides genuine predictive advantages when the forecasting target exhibits strong temporal persistence. Third, studies that jointly consider multiple fuel types, such as petrol, diesel, and LPG, with multiple global equity indices remain limited, which restricts the generalizability of existing findings across different market conditions.

To address these gaps, this study develops a time-varying fuel–equity correlation forecasting framework using a 60-day rolling window across three fuel variables and five major global stock indices, namely S&P 500, NASDAQ, FTSE 100, Nikkei 225, and IHSG/JKSE. This study provides a systematic empirical comparison of 17 forecasting models across four categories: baseline, PCA-based, standalone deep learning, and hybrid deep learning models. By focusing specifically on rolling-correlation targets, this study examines how target characteristics influence model performance. The findings provide empirical evidence that, in this experimental setting, PCA-based dimensionality reduction and persistence-based forecasting can match or outperform hybrid deep learning architectures when the target series is smooth and persistent. This study therefore contributes a replicable comparative framework for fuel–equity correlation forecasting that can be extended to other cross-market analysis contexts, offering a methodological insight that questions the common assumption that more complex deep learning models always produce superior forecasting performance in financial time-series prediction.

To the best of the authors' knowledge, relatively few studies have systematically compared baseline, PCA-based, standalone deep learning, and hybrid deep learning models specifically for forecasting time-varying rolling correlations between multiple fuel types and global equity indices simultaneously. Whereas most prior studies focus on direct price or volatility prediction, this study treats rolling correlation as the primary forecasting target and provides empirical evidence that target characteristics—particularly smoothness and temporal persistence—can play an important role in determining the relative effectiveness of different model classes. This emphasis constitutes the main contribution of the present study and adds an empirical perspective to the financial forecasting literature, rather than claiming a definitive or exhaustive comparison.

LITERATURE REVIEW

Financial time-series forecasting is important because global markets are affected by volatility, uncertainty, and cross-market spillovers. Fuel and oil price movements are closely related to stock market behavior because energy price changes can influence production costs, inflation, investor sentiment, and market risk (Nguyen et al., 2026; Dang & Liu, 2026; Liu & Wang, 2026). Previous studies also show that the relationship between commodity and stock markets is dynamic and changes over time (Ben Flah et al., 2025; Li & Zhong, 2025). However, most of these studies examine price-level movements, volatility spillovers, or market connectedness rather than forecasting the future values of time-varying fuel–equity correlations as a specific prediction target.

Traditional forecasting models such as ARIMA and GARCH are widely used in financial forecasting, but they often have limitations in capturing nonlinear patterns and complex temporal dependencies (Mishra & Dash, 2024; Chung et al., 2025). Therefore, deep learning models such as LSTM, CNN-LSTM, and attention-based models have been increasingly applied because they can learn sequential patterns and improve forecasting performance in energy and financial markets (Sen & Dutta Choudhury, 2024; Lahmiri, 2026; Ma & Han, 2026). However, these models have been predominantly evaluated on direct price, return, or volatility series, which usually contain stronger fluctuations and nonlinear dynamics. Rolling correlation series, by contrast, are bounded between −1 and 1 and may exhibit strong temporal persistence, especially when constructed using overlapping rolling windows. This difference in target characteristics may reduce the advantage of complex sequential architectures.

PCA is also commonly used to reduce high-dimensional financial features while preserving important information. Dimensionality reduction can improve computational efficiency and reduce feature redundancy in forecasting tasks (Cao et al., 2016; Wang, 2024). Studies such as Chen et al. (2024) and Wang (2024) have shown that combining PCA with deep learning can improve prediction stability on stock index targets. However, the role of PCA as a standalone forecasting approach for persistent rolling correlation targets remains less explored, especially in comparison with hybrid deep learning architectures.

Several studies have explored dynamic correlation forecasting between financial assets. Ni and Xu (2023) forecasted dynamic correlations between stock indices using deep learning methods and showed that sequential models can capture correlation dynamics. Ma and Han (2026) forecasted dynamic correlations between carbon, energy, and stock markets using a graph neural network approach. However, these studies focus mainly on stock-

*name of corresponding author

Sinkron : Jurnal dan Penelitian Teknik Informatika
Volume 10, Number 3, July 2026
DOI : https://doi.org/10.33395/sinkron.v10i3.16410
e-ISSN : 2541-2019
p-ISSN : 2541-044X

to-stock or carbon–energy–stock relationships rather than fuel-to-equity correlations involving multiple fuel types and global stock indices simultaneously. In addition, none of these studies directly examine whether complex sequential models retain their predictive advantage when the forecasting target is a smoother and more persistent rolling correlation series rather than a volatile price or return series.

Beck et al. (2025) highlighted that naive forecasts are frequently underestimated in time series with low predictability or strong persistence, suggesting that persistence-based baselines should be rigorously evaluated before adopting complex models. This argument is particularly relevant for rolling correlation forecasting, where the autocorrelated nature of the target may favor simpler approaches. Despite this, many financial forecasting studies still emphasize complex deep learning models without sufficiently testing whether these models outperform strong baseline and dimensionality-reduction approaches.

Based on this review, three specific gaps are identified. First, forecasting rolling correlations between multiple fuel types and global equity indices remains limited compared with direct price and volatility forecasting. Second, the performance of hybrid deep learning models on persistent rolling-correlation targets, rather than volatile price series, is still not fully understood. Third, the comparative role of PCA-based dimensionality reduction versus hybrid deep learning architectures in fuel–equity correlation forecasting requires further empirical evaluation. This study addresses these gaps by developing a systematic comparative framework that evaluates 17 forecasting models on fuel–equity rolling correlation targets.

METHOD

Research Design

This study employed a quantitative time-series forecasting approach to predict time-varying correlations between global fuel prices and stock market indices. The research compared baseline, PCA-based, standalone deep learning, and hybrid deep learning models to identify the most effective approach for forecasting dynamic fuel–equity market correlations.

Research Workflow

The overall research workflow is shown in Figure 1. The process began with data acquisition, followed by data preprocessing, feature engineering, rolling correlation construction, PCA reduction, sequence preparation, model development, and performance evaluation. This workflow was designed to ensure that the financial time-series data were cleaned, transformed, and structured before being used in the forecasting models.

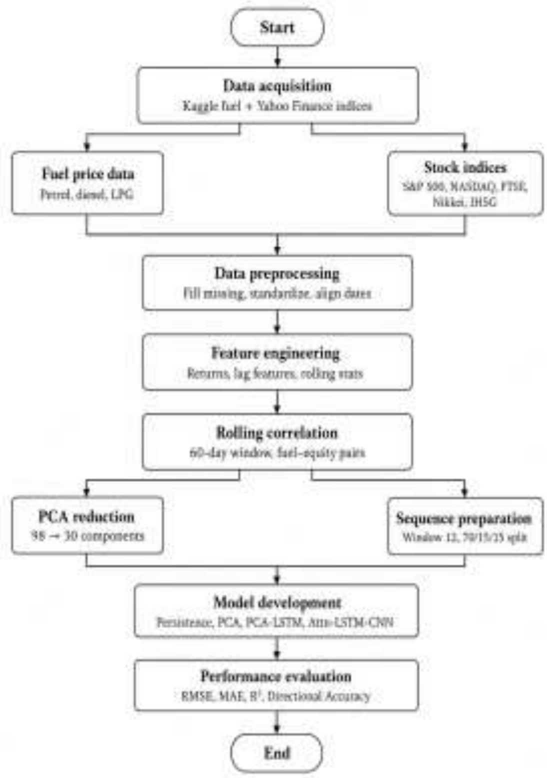

Figure 1. Proposed Research Workflow

*name of corresponding author

Data Collection and Preprocessing

The dataset used in this study was obtained from public data sources. The raw fuel price dataset was collected from Kaggle, while stock market index data were obtained from Yahoo Finance. The fuel price dataset includes global fuel price information, while the stock market data include historical index data from S&P 500, NASDAQ Composite, FTSE 100, Nikkei 225, and IHSG/JKSE. The data sources used in this study are listed below:

Raw fuel price dataset:
Global Fuel Prices 2020–2026:
https://www.kaggle.com/datasets/belbino/global-fuel-prices-20202026

Yahoo Finance index sources:
S&P 500 (^GSPC) Historical Data:
https://finance.yahoo.com/quote/%5EGSPC/history/

NASDAQ Composite (^IXIC) Historical Data:
https://finance.yahoo.com/quote/%5EIXIC/history/

FTSE 100 (^FTSE) Historical Data:
https://finance.yahoo.com/quote/%5EFTSE/history/

Nikkei 225 (^N225) Historical Data:
https://finance.yahoo.com/quote/%5EN225/history/

IHSG / IDX Composite (^JKSE) Historical Data:
https://finance.yahoo.com/quote/%5EJKSE/history/

The datasets were then synchronized based on date to maintain chronological consistency. Missing values were handled using forward fill and backward fill, while numerical variables were standardized before model training. Chronological ordering was maintained because the forecasting task depends on the correct temporal structure of the data.

Feature Engineering and Rolling Correlation

Feature engineering was conducted by transforming raw price data into returns, constructing lag features, and generating rolling statistical features. Return transformation was used to represent relative market movements rather than absolute price values. The return was calculated as follows:

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}}$$

where R_t represents the return at time t, P_t denotes the asset price at time t, and P_{t-1} denotes the asset price at the previous time step.

The forecasting target was constructed using rolling correlation between fuel returns and stock market returns. A 60-day rolling window was applied to capture the dynamic relationship between fuel prices and equity markets. The rolling correlation was calculated as follows:

$$\rho_{xy,t} = \mathrm{Corr}(x_{t-w:t}, y_{t-w:t})$$

where $\rho_{xy,t}$ represents the rolling correlation between fuel return x and equity return y at time t, while w represents the rolling window size, which was set to 60 days in this study.

PCA Reduction and Sequence Preparation

After the rolling correlation target was generated, Principal Component Analysis (PCA) was applied to reduce the dimensionality of the feature space. The original feature set, consisting of 98 features, was reduced into 30 principal components to preserve important information while reducing redundancy and computational

*name of corresponding author

Sinkron : Jurnal dan Penelitian Teknik Informatika
Volume 10, Number 3, July 2026 e-ISSN : 2541-2019
DOI : https://doi.org/10.33395/sinkron.v10i3.16410 p-ISSN : 2541-044X

complexity. The PCA-transformed data were then converted into sequential input data using a sliding window approach with a window size of 12.

The dataset was divided chronologically into training, validation, and testing subsets using a 70:15:15 ratio. Chronological splitting was used instead of random splitting to prevent future information from leaking into the training process.

Model Development and Evaluation

This study evaluated 17 forecasting models grouped into four categories. The first category was the baseline category, represented by the Persistence Model, which predicts the next rolling correlation value using the most recently observed correlation value without any training. The second category was the PCA-based category, represented by the PCA Model. In this study, the PCA Model refers to a two-stage procedure in which Principal Component Analysis is first applied to reduce the high-dimensional feature set into 30 principal components, and a linear regression model is then trained to map these principal components to the rolling correlation target. The PCA Model therefore combines PCA dimensionality reduction with linear regression (PCA + Linear Regression) and does not use any sequential or recurrent component. The third category consisted of standalone deep learning models, namely LSTM, CNN, and Attention-based models, which learn temporal patterns directly from the sequential input. The fourth category consisted of hybrid deep learning models, namely PCA-LSTM, CNN-LSTM, Attention-LSTM-CNN, PCA-Attention, PCA-CNN, LSTM-Attention, CNN-Attention, PCA-Attention-LSTM, PCA-CNN-LSTM, and PCA-Attention-LSTM-CNN, which combine dimensionality reduction and multiple sequential mechanisms. In addition, two classical machine learning baselines were included for comparison, namely a Linear Regression model trained directly on the full processed feature set without dimensionality reduction, and an XGBoost model, a gradient boosting tree ensemble, also trained on the full processed feature set. These classical baselines were added to provide a stronger and more conventional reference against which the PCA-based, persistence, and deep learning models could be evaluated. All models were compared to evaluate whether more complex architectures improved forecasting performance compared with simpler baseline and PCA-based approaches.

Model performance was measured using Root Mean Square Error (RMSE), Mean Absolute Error (MAE), R-squared (R^2), and Directional Accuracy. RMSE and MAE were used to measure prediction error, R^2 was used to evaluate the ability of the model to explain target variation, and Directional Accuracy was used to measure whether the model correctly predicted the direction of correlation movement. To assess whether the difference in predictive accuracy between the two best-performing models was statistically meaningful rather than incidental, a Diebold-Mariano test was applied to their forecast errors, and a bootstrap resampling procedure with 1,000 resamples was used to construct a confidence interval for the difference in RMSE. The model evaluation in the program also used RMSE, MAE, R^2, and Directional Accuracy as the main comparison metrics.

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}$$

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|$$

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}$$

where y_i represents the actual value, $\hat{y}_i$ represents the predicted value, $\bar{y}$ represents the mean of actual values, and n represents the number of observations.

RESULT

This section presents the experimental results of the proposed forecasting framework for predicting dynamic correlations between global fuel prices and stock market indices. The evaluation compares baseline, PCA-based, standalone deep learning, and hybrid deep learning models using RMSE, MAE, R^2, and Directional Accuracy. These metrics were selected to assess numerical prediction error, explanatory ability, and the capacity of each model to predict the direction of correlation movement.

Before model evaluation, the forecasting target was generated using rolling correlation between fuel returns and stock market returns. The fuel variables consisted of petrol, diesel, and LPG, while the equity variables

*name of corresponding author

Sinkron : Jurnal dan Penelitian Teknik Informatika
Volume 10, Number 3, July 2026 e-ISSN : 2541-2019
DOI : https://doi.org/10.33395/sinkron.v10i3.16410 p-ISSN : 2541-044X

consisted of S&P 500, NASDAQ Composite, FTSE 100, Nikkei 225, and IHSG/JKSE. A 60-day rolling window and a one-day forecast horizon were applied to produce time-varying fuel–equity correlation targets.

After the rolling correlation target was constructed, PCA was applied to reduce the feature space. The original feature set was reduced into 30 principal components to preserve important information while lowering feature complexity. The PCA-transformed data were then converted into sequential input using a sliding window size of 12 and divided chronologically into training, validation, and testing subsets using a 70:15:15 ratio.
A total of 17 forecasting models were evaluated, consisting of baseline models, PCA-based models, standalone deep learning models, and hybrid deep learning models.

The performance comparison in Table 1 shows a clear separation between simpler models and hybrid deep learning models. The PCA Model and Persistence Model produced substantially lower errors than PCA-LSTM, Attention-LSTM-CNN, and PCA-Attention. This pattern suggests that the rolling correlation target is highly persistent and can be represented effectively using lower-dimensional features or recent historical information.

Table 1 presents the complete forecasting performance of all 17 evaluated models, ranked by RMSE. The evaluated models represent different forecasting strategies, including dimensionality reduction, naive persistence, standalone deep learning, and hybrid sequence-learning architectures. Therefore, the table not only presents model ranking but also shows how model complexity affects forecasting performance in this dataset. The RMSE and Directional Accuracy of all evaluated models are further compared in Figure 2 and Figure 4.

Table 1. Forecasting Performance Comparison of All 17 Evaluated Models (Ranked by RMSE)

Model	RMSE	MAE	R^2	Directional Accuracy (%)
PCA Model	0.016909	0.002769	0.967045	97.240179
Linear Regression	0.016975	0.002870	0.966787	97.230234
Persistence	0.017002	0.002293	0.966682	97.767280
XGBoost	0.020410	0.004702	0.951986	96.757832
PCA-LSTM	0.093303	0.047417	-0.003783	49.303413
Attention-LSTM-CNN	0.093317	0.047930	-0.004089	53.060006
PCA-Attention	0.093318	0.047222	-0.004110	48.890437
CNN	0.093354	0.048232	-0.004886	50.019902
PCA-Attention-LSTM-CNN	0.093395	0.047725	-0.005774	55.040303
PCA-CNN	0.093414	0.047309	-0.006193	50.661757
LSTM-Attention	0.093460	0.047640	-0.007182	52.910737
CNN-Attention	0.093515	0.048051	-0.008367	52.522639
LSTM	0.093527	0.047287	-0.008610	53.328689
CNN-LSTM	0.093601	0.047579	-0.010215	53.189372
PCA-Attention-LSTM	0.093610	0.047458	-0.010409	49.995024
Attention	0.093665	0.046957	-0.011602	47.676386

*name of corresponding author

Sinkron : Jurnal dan Penelitian Teknik Informatika
Volume 10, Number 3, July 2026
DOI : https://doi.org/10.33395/sinkron.v10i3.16410
e-ISSN : 2541-2019
p-ISSN : 2541-044X

Model	RMSE	MAE	R^2	Directional Accuracy (%)
PCA-CNN-LSTM	0.093751	0.047766	-0.013453	52.826152

As shown in Table 1, the PCA Model (PCA + Linear Regression) obtained the lowest RMSE value of 0.016909. This indicates that the PCA Model generated the smallest overall prediction error when larger deviations were penalized more strongly. The PCA Model also achieved the highest R^2 value of 0.967045, showing that the reduced feature representation explained most of the variation in the rolling correlation target.

The Persistence Model produced an RMSE value of 0.017002, which was only slightly higher than the PCA Model. However, it achieved the lowest MAE value of 0.002293 and the highest Directional Accuracy value of 97.767280%. This result indicates that the most recent correlation value was highly informative for predicting the next value, especially when the evaluation focused on average absolute error and movement direction.

The hybrid deep learning models showed considerably weaker results. The PCA-LSTM model produced an RMSE value of 0.093303, an MAE value of 0.047417, an R^2 value of -0.003783, and a Directional Accuracy value of 49.303413%. The Attention-LSTM-CNN model produced an RMSE value of 0.093317, an MAE value of 0.047930, an R^2 value of -0.004089, and a Directional Accuracy value of 53.060006%. Meanwhile, the PCA-Attention model produced an RMSE value of 0.093318, an MAE value of 0.047222, an R^2 value of -0.004110, and a Directional Accuracy value of 48.890437%.

The two classical machine learning baselines occupied an intermediate position that further supports the main finding. The Linear Regression model, trained on the full processed feature set, achieved an RMSE of 0.016975, an MAE of 0.002870, an R^2 of 0.966787, and a Directional Accuracy of 97.230234%, placing it within the same top-performing group as the PCA Model and the Persistence Model. The XGBoost model achieved an RMSE of 0.020410, an MAE of 0.004702, an R^2 of 0.951986, and a Directional Accuracy of 96.757832%. Although XGBoost is a strong non-linear ensemble method, it produced a higher error than the simple linear models and the persistence baseline, while still clearly outperforming all hybrid deep learning models, which all produced negative R^2 values. This pattern reinforces the central argument of this study: on a smooth and persistent rolling correlation target, neither deep sequential architectures nor non-linear tree ensembles provided an advantage over simple linear and persistence-based approaches.

These values indicate that the hybrid models not only produced higher errors but also failed to explain the target variance effectively. The negative R^2 values suggest that their predictions were less reliable than a simple mean-based reference in this testing setting. Directional Accuracy values close to 50% further show that the hybrid models had limited ability to detect whether the rolling correlation would increase or decrease.

Statistical Validation of Target Characteristics

To verify that the forecasting target is indeed smooth and persistent, its statistical properties were examined directly. Because the evaluation dataset concatenates the test samples of many fuel–equity pairs, the persistence and smoothness of the target were measured within each individual pair series (765 series in total) and then summarized, rather than on the concatenated vector, where transitions between different pairs would distort the autocorrelation structure. As shown in Table 2, the rolling correlation target is strongly persistent: the mean lag-1 autocorrelation across all pair series is 0.936, and 85.2% of the series have a lag-1 autocorrelation above 0.9. The target is also smooth, with a mean smoothness ratio of 0.282, indicating that one-step changes are small relative to the overall variation of the series. Consistent with this high persistence, an Augmented Dickey-Fuller test classified only about 9.9% of the pair series as stationary at the 5% level, reflecting near-unit-root behavior. These results provide direct statistical evidence that the most recent correlation value carries substantial predictive information, which explains why the Persistence Model and the simple PCA-based and linear models perform strongly, and why the additional complexity of hybrid deep learning architectures does not translate into improved accuracy on this target.

Table 2. Statistical Characteristics of the Rolling Correlation Target (computed per fuel–equity pair series, then averaged)

Statistic	Value (per fuel–equity pair series)
Mean lag-1 autocorrelation	0.936
Median lag-1 autocorrelation	0.951

*name of corresponding author

Sinkron : Jurnal dan Penelitian Teknik Informatika
Volume 10, Number 3, July 2026
DOI : https://doi.org/10.33395/sinkron.v10i3.16410

e-ISSN : 2541-2019
p-ISSN : 2541-044X

Statistic	Value (per fuel–equity pair series)
Series with lag-1 autocorrelation > 0.9	85.2%
Series with lag-1 autocorrelation > 0.8	96.3%
Mean smoothness ratio (std of first difference / std of level)	0.282
Series stationary at 5% (ADF)	9.9%

Because the difference between the PCA Model and the Persistence Model is very small (RMSE of 0.016909 versus 0.017002, and R^2 of 0.967045 versus 0.966682), the statistical significance of this difference was examined rather than relying on the raw metric values alone. A Diebold-Mariano test was applied to compare the predictive accuracy of the two models, and a bootstrap resampling procedure with 1,000 resamples was used to construct a confidence interval for the difference in RMSE. Using a squared-error loss, the Diebold-Mariano test produced a statistic of −3.10 with a p-value of 0.0019, and the bootstrap 95% confidence interval for the RMSE difference, $[-1.355\times10^{-4}, -3.80\times10^{-5}]$, did not contain zero. This indicates that the difference between the PCA Model and the Persistence Model is statistically significant at the 5% level. However, the magnitude of this difference is very small in practical terms, since the RMSE differs by only about 9×10^{-5} (approximately 0.5%), and the Persistence Model in fact achieved a lower MAE and a higher Directional Accuracy than the PCA Model. The two simpler models should therefore be interpreted as a single top-performing group: although the PCA Model holds a small but statistically detectable advantage in RMSE, the practical difference between the two is negligible. In contrast, the difference between this top-performing group and the hybrid deep learning models is large and consistent across all metrics, confirming that the lower-error group is clearly separated from the higher-error group. Therefore, the central finding of this study is not that the PCA Model strictly outperforms the Persistence Model, but that both simpler models clearly outperform the more complex hybrid deep learning models on this smooth and persistent rolling correlation target.

The RMSE comparison in Figure 2 visually reinforces the numerical results in Table 1. The PCA Model, Linear Regression, and Persistence Model form a low-error group, joined closely by XGBoost, while all standalone and hybrid deep learning models form a high-error group with nearly identical RMSE values. This pattern suggests that adding LSTM, CNN, or Attention components did not provide additional predictive advantage for this smooth and persistent correlation series.

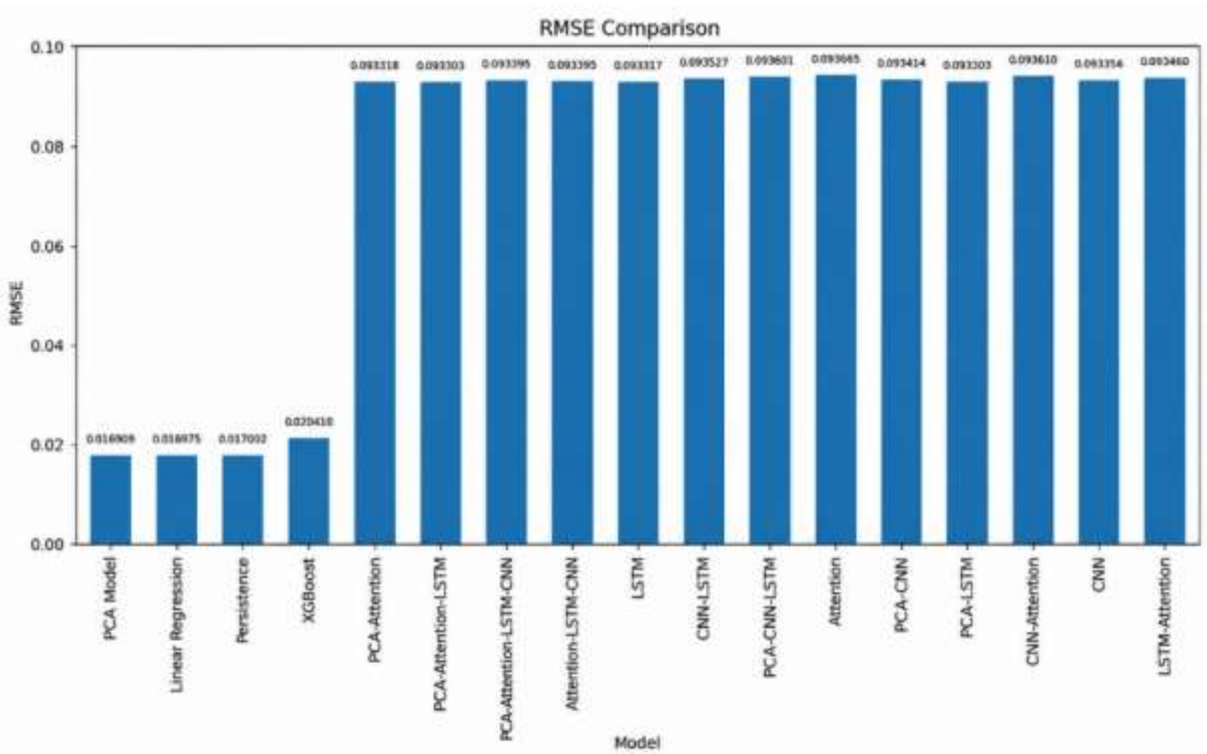

*name of corresponding author

Figure 2. RMSE Comparison of All Evaluated Forecasting Models

The actual and predicted rolling correlation values of the best-performing model are shown in **Figure 3**. This visualization helps interpret the numerical metrics by showing whether the prediction curve follows the level and movement of the actual target. A close alignment between the actual and predicted curves supports the high **R^2** score of the PCA Model, while visible deviations indicate periods where the model had difficulty adjusting to changes in correlation behavior. It should be noted that the horizontal axis in Figure 3 does not represent a single continuous calendar timeline. Because the forecasting target was constructed for every fuel-equity pair, the test samples of all pairs (petrol, diesel, and LPG against the S&P 500, NASDAQ, FTSE 100, Nikkei 225, and IHSG/JKSE) were concatenated into a single evaluation sequence. The time-step index therefore extends to approximately 20,000 because it counts the combined test observations across all fuel-equity pairs rather than the number of trading days for a single pair.

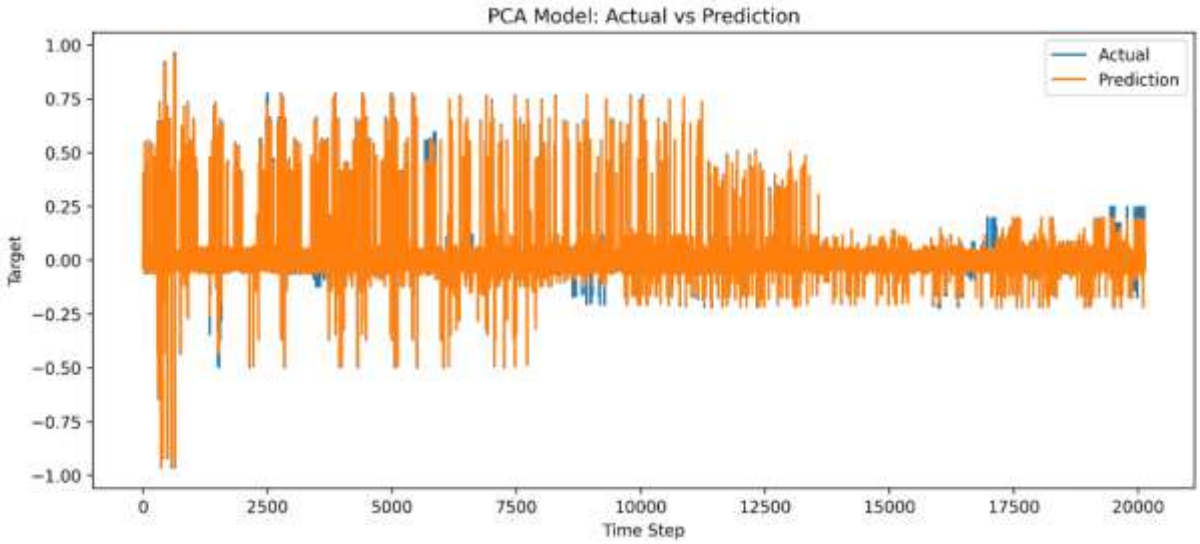

Figure 3. Actual and predicted rolling correlation values of the best-performing model.

The Directional Accuracy comparison in **Figure 4** provides another interpretation of model performance. The Persistence Model achieved the strongest directional performance, indicating that short-term correlation movements were often continuation-based. In contrast, the hybrid models produced directional scores around the random-guessing range, suggesting that their more complex architectures did not improve short-term movement prediction.

*name of corresponding author

Sinkron : Jurnal dan Penelitian Teknik Informatika
Volume 10, Number 3, July 2026
DOI : https://doi.org/10.33395/sinkron.v10i3.16410
e-ISSN : 2541-2019
p-ISSN : 2541-044X

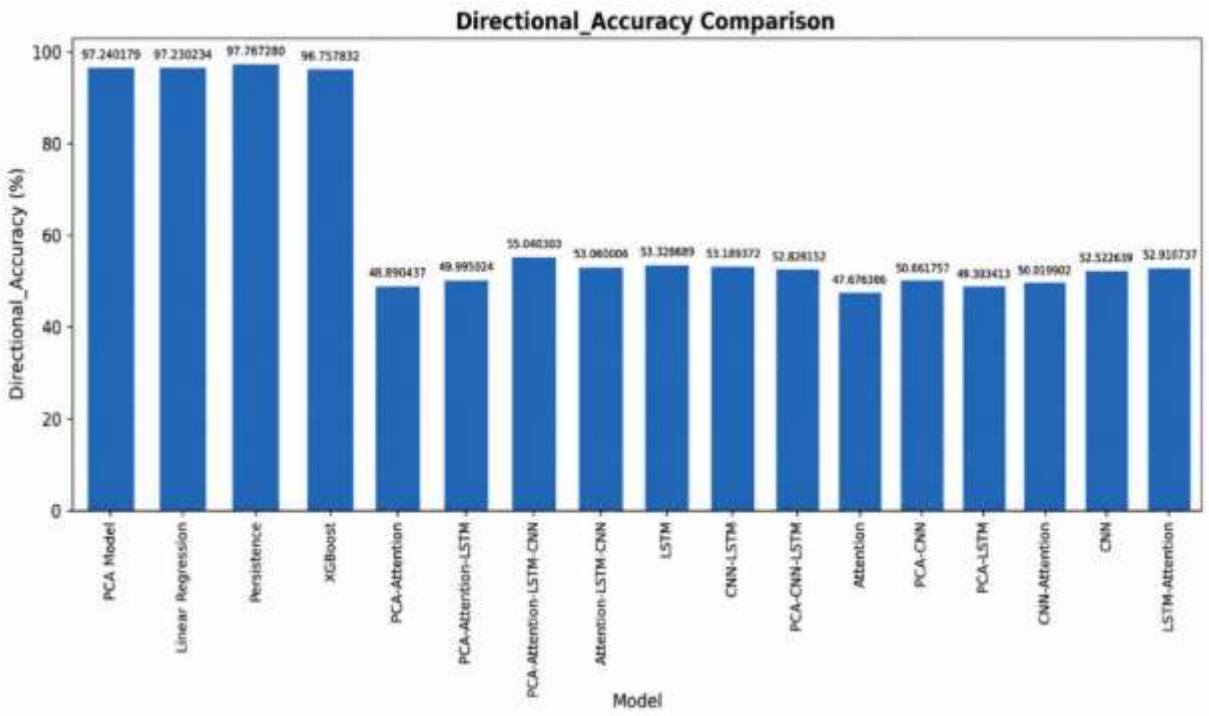

Figure 4. Directional Accuracy Comparison of All Evaluated Forecasting Models

Overall, the result section indicates that simpler models were more suitable for the forecasting target used in this study. The PCA Model was the strongest model in terms of **RMSE and R^2**, while the Persistence Model was the strongest model in terms of **MAE and Directional Accuracy**. These findings indicate that model selection for dynamic correlation forecasting should consider target smoothness, temporal persistence, and the risk of unnecessary model complexity.

DISCUSSION

The experimental results reveal a theoretically meaningful finding: simpler models outperformed hybrid deep learning architectures in forecasting dynamic fuel–equity correlations. This outcome is not coincidental but can be explained through the statistical characteristics of the forecasting target. Rolling correlations computed over a 60-day window tend to produce smoother, bounded, and highly autocorrelated series. Under these conditions, the most recent observation carries substantial predictive information, which explains the strong performance of the Persistence Model in terms of MAE and Directional Accuracy. This finding is consistent with Beck et al. (2025), who argued that naive forecasts are frequently underestimated in time series with strong persistence, and extends their argument to the specific context of fuel–equity correlation forecasting.

The superior performance of the PCA Model in terms of RMSE and R^2 can be attributed to the effectiveness of dimensionality reduction in reducing redundant and noisy features while preserving the principal variance structure of the data. By compressing the original feature space into 30 principal components, PCA retained the dominant information needed to explain rolling correlation dynamics while lowering feature complexity. This finding is consistent with Cao et al. (2016) and Wang (2024), who demonstrated that feature compression can improve model stability in financial forecasting. Importantly, this study extends their findings by showing that PCA can remain highly competitive even without a complex sequential modeling component when applied to persistent rolling correlation targets.

The underperformance of hybrid deep learning models requires deeper explanation beyond simply reporting higher error values. LSTM, CNN, and attention mechanisms are designed to capture long-range temporal dependencies and complex nonlinear patterns in volatile series. When applied to a rolling correlation series that is smoother, bounded, and strongly autocorrelated, these mechanisms may introduce unnecessary model complexity without providing additional predictive information. The negative R^2 values observed across the hybrid models indicate that their predictions were less reliable than a simple mean-based reference in this testing setting. Directional Accuracy values near 50% further show that the hybrid models had limited ability to distinguish whether the rolling correlation would increase or decrease from one time step to the next. This pattern suggests

*name of corresponding author

e-ISSN : 2541-2019
p-ISSN : 2541-044X

that the additional parameters introduced by CNN, LSTM, and attention layers may have reduced generalization performance on this smooth and persistent target.

These findings differ from several studies that report strong hybrid deep learning performance on different forecasting targets. Sen and Dutta Choudhury (2024) demonstrated that deep learning models can effectively forecast crude oil prices over four decades, achieving strong predictive accuracy on highly volatile and non-stationary price series. Similarly, Chung et al. (2025) showed that deep learning outperforms GARCH models in forecasting financial volatility under structural breaks, where nonlinear dynamics are prominent. Lahmiri (2026) reported that a hybrid CNN-LSTM-GRU system achieved strong performance in gold and Brent price prediction, where the target series exhibited complex temporal patterns. The key distinction between these studies and the present research lies in the nature of the forecasting target. All of the above studies evaluate model performance on direct price, return, or volatility series, which are typically non-stationary, highly volatile, and may contain complex nonlinear dynamics that benefit from the pattern recognition capability of deep learning architectures. Rolling correlation series, by contrast, are bounded and persistence-dominated, especially when constructed using overlapping rolling windows. This fundamental difference in target characteristics conditions the relative effectiveness of different model classes. Therefore, the results of this study highlight that the superiority of deep learning models is target-dependent rather than universal, which represents an empirical contribution to the fuel–equity forecasting literature by showing that hybrid deep learning models may lose their predictive advantage when the forecasting target is smoother, bounded, and persistence-dominated.

From a theoretical perspective, these results question the assumption that increasing architectural complexity consistently improves forecasting accuracy in financial time-series prediction. The findings suggest that model selection should be guided primarily by the statistical properties of the forecasting target, particularly its smoothness, persistence, and autocorrelation structure, rather than by architectural sophistication alone. This aligns with the principle of parsimony in statistical modeling, which favors simpler models when they provide equivalent or superior predictive performance.

From a practical perspective, practitioners in financial risk management and portfolio optimization can benefit from these findings by recognizing that PCA-based approaches offer computationally efficient and accurate alternatives to deep learning when monitoring dynamic fuel–equity correlations. The lower computational cost of PCA-based forecasting may also make it more suitable for real-time risk monitoring applications compared with hybrid deep learning architectures that require more training time, more parameter tuning, and greater computational resources.

This study has several limitations that should be acknowledged. The forecasting target was based on a fixed 60-day rolling window and a one-day forecast horizon, so different window sizes or forecast horizons may produce different results regarding the relative performance of model categories. The dataset was limited to three fuel price variables and five stock market indices, which may not fully represent the diversity of global fuel and equity markets. In addition, the hyperparameters of the deep learning models were tuned during the validation phase but were not subjected to exhaustive optimization. Therefore, further hyperparameter tuning could potentially improve hybrid model performance.

Future research should test alternative rolling window sizes, incorporate additional macroeconomic variables such as inflation rates and exchange rates, explore longer forecast horizons, and apply more systematic hyperparameter optimization techniques such as Bayesian optimization. Future studies may also examine whether hybrid models recover their predictive advantage under different target conditions, especially during periods of structural market shocks such as financial crises or geopolitical disruptions, where correlation dynamics may become more volatile and nonlinear.

CONCLUSION

This study developed a systematic comparative framework for forecasting time-varying correlations between global fuel prices and stock market indices. The research evaluated 17 forecasting models across four categories, namely baseline, PCA-based, standalone deep learning, and hybrid deep learning models. The forecasting target was constructed using a 60-day rolling correlation based on three fuel variables, namely petrol, diesel, and LPG, and five major global stock indices, namely S&P 500, NASDAQ, FTSE 100, Nikkei 225, and IHSG/JKSE. The research process included data preprocessing, feature engineering, rolling correlation construction, PCA reduction, sequence preparation, model development, and performance evaluation using RMSE, MAE, R^2, and Directional Accuracy.

The results showed that the PCA Model achieved the best performance in terms of RMSE and R^2, with an RMSE value of 0.016909 and an R^2 value of 0.967045. The Persistence Model achieved the lowest MAE value of 0.002293 and the highest Directional Accuracy of 97.767280%. In contrast, the hybrid deep learning models produced negative R^2 values and Directional Accuracy values near 50%, indicating that architectural complexity did not improve forecasting performance on this smooth and persistent rolling correlation target. These findings

*name of corresponding author

e-ISSN : 2541-2019
p-ISSN : 2541-044X

provide empirical evidence that simpler models can be more effective than hybrid deep learning models when the forecasting target exhibits strong temporal persistence and smoothness.

This study contributes to the financial forecasting literature in three ways. First, it shows that the statistical properties of the forecasting target, particularly smoothness and temporal persistence, should be a key consideration in model selection rather than relying on architectural complexity alone. This finding questions the assumption that more complex deep learning models always produce superior forecasting performance. Second, it demonstrates that PCA-based dimensionality reduction can serve as a competitive standalone forecasting approach for persistent rolling correlation targets, outperforming hybrid architectures such as CNN-LSTM, Attention-LSTM-CNN, and PCA-Attention in this experimental setting. Third, it provides a replicable comparative framework for fuel–equity correlation forecasting that combines rolling correlation construction, PCA reduction, sequence preparation, and systematic model evaluation across multiple fuel types and global stock indices, which can be extended to other cross-market correlation forecasting problems beyond the fuel–equity context.

However, this study has several limitations. The forecasting framework was based on a fixed 60-day rolling window and a one-day forecast horizon, which may not fully capture longer-term correlation dynamics or perform optimally under different window configurations. The dataset was limited to three fuel price variables and five stock market indices, so the findings may not represent all global energy and equity market conditions. In addition, the hyperparameters of the deep learning models were tuned during the validation phase but were not subjected to exhaustive optimization. Therefore, more systematic tuning may improve hybrid model performance under different experimental settings.

Future research should test alternative rolling window sizes and longer forecast horizons to examine whether model rankings remain consistent under different temporal configurations. Additional macroeconomic variables, such as inflation rates, exchange rates, interest rates, and geopolitical risk indices, should also be incorporated to enrich the feature space and potentially improve forecasting accuracy. More systematic hyperparameter optimization techniques, such as Bayesian optimization, can be applied to evaluate hybrid models more comprehensively. Future studies may also examine whether the findings generalize to other cross-market correlation forecasting problems, such as cryptocurrency–equity, commodity–bond, or currency–equity correlations.

ACKNOWLEDGMENT

The author would like to express sincere gratitude to the academic supervisor for providing guidance, feedback, and constructive suggestions throughout the research process. The author also acknowledges the availability of public datasets from Kaggle and Yahoo Finance, which supported the data collection process in this study. Appreciation is also extended to all individuals who provided technical assistance, comments, and support during the preparation of this research. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

REFERENCES

Akusta, A., Yıldırım, H. H., Gün, M., & Sakarya, Ş. (2026). Deep learning enhanced energy market prediction: A robust ARIMAX–LSTM fusion for crude oil pricing. *Journal of Computational and Applied Mathematics*, *474*, 117006. https://doi.org/10.1016/j.cam.2025.117006

Amagbo, R., & Geman, H. (2026). Dynamic and asymmetric spillovers between crude oil, biofuels and agricultural commodities: Evidence from periods of geopolitical tensions and energy policy uncertainty. *International Review of Economics & Finance*, *108*, 105276. https://doi.org/10.1016/j.iref.2026.105276

Beck, N., Schemm, J., Ehrig, C., Sonnleitner, B., Neumann, U., & Zimmermann, H. G. (2025). Mind the naive forecast! A rigorous evaluation of forecasting models for time series with low predictability. *Applied Intelligence*. https://doi.org/10.1007/s10489-025-06268-w

Ben Flah, I., Samet, K., El Ammari, A., & Terzi, C. (2025). Shock and Volatility Transmissions Across Global Commodity and Stock Markets Spillovers: Empirical Evidence from Africa. *Journal of Risk and Financial Management*, *18*(6), 332. https://doi.org/10.3390/jrfm18060332

Ben Romdhane, W., & Boubaker, H. (2026). A Hybrid HAR-LSTM-GARCH Model for Forecasting Volatility in Energy Markets. *Journal of Risk and Financial Management*,*19*(1),77.https://doi.org/10.3390/jrfm19010077

Cao, L., et al. (2016). Forecasting daily stock market return using dimensionality reduction. *Expert Systems with Applications*, *67*, 126–139. https://doi.org/10.1016/j.eswa.2016.09.027

Chen, Z., et al. (2024). PCA-ICA-LSTM: A hybrid deep learning model based on dimension reduction methods to predict S&P 500 index price. *Computational Economics*. https://doi.org/10.1007/s10614-024-10629-x

Choudhary, J., Sharma, H. K., Malik, P., & Majumder, S. (2025). Price Forecasting of Crude Oil Using Hybrid Machine Learning Models. *Journal of Risk and Financial Management*, *18*(7),346.https://doi.org/10.3390/jrfm18070346

*name of corresponding author

Chung, V., Espinoza, J., & Quispe, R. (2025). Forecasting Financial Volatility Under Structural Breaks: A Comparative Study of GARCH Models and Deep Learning Techniques. *Journal of Risk and Financial Management*, *18*(9), 494. https://doi.org/10.3390/jrfm18090494

Dang, T. P. T., & Liu, B. (2026). The spillover effects of gold and oil prices: a study of the Vietnamese and global markets. *Journal of Financial Economic Policy*, 1–18. https://doi.org/10.1108/JFEP-07-2025-0291

Egbe, A. A., Eyupoglu, S. Z., & Seraj, M. (2026). Sustainability Investment in Distress: Volatility Spillovers and ESG Markets Portfolio Implications. *Sustainability*, *18*(9), 4403. https://doi.org/10.3390/su18094403

Firouzjaee, M. A., et al. (2024). The interpretability of LSTM models for predicting oil company stocks: Impact of correlated features. *International Journal of Energy Research*. https://doi.org/10.1155/2024/5526692

Kakade, K., Jain, I., & Mishra, A. K. (2022). Value-at-Risk forecasting: A hybrid ensemble learning GARCH-LSTM based approach. *Resources Policy*,*78*,102903.https://doi.org/10.1016/j.resourpol.2022.102903

Lahmiri, S. (2026). A hybrid CNN-LSTM-GRU system tuned by whale optimization algorithm for enhanced gold and brent price prediction. *Finance Research Open*, *2*(1), 100100.https://doi.org/10.1016/j.finr.2026.100100

Li, A., & Zhong, B. (2025). Asymmetric spillover connectedness between clean energy markets and industrial stock markets: How uncertainties affect it. *PLoS ONE*, *20*(3), e0316171. https://doi.org/10.1371/journal.pone.0316171

Li, S., & Xu, S. (2025). Enhancing stock price prediction using GANs and transformer-based attention mechanisms. *Empirical Economics*, *68*, 373–403. https://doi.org/10.1007/s00181-024-02644-6

Liu, Z., & Wang, J. (2026). Stock markets, Sino–US monetary policy, and international crude oil prices. *Finance Research Letters*, *94*, 109629. https://doi.org/10.1016/j.frl.2026.109629

Ma, Q., & Han, M. (2026). Forecasting Dynamic Correlations Between Carbon, Energy, and Stock Markets Using a BOHB-Optimized Multivariable Graph Neural Network. *Mathematics*,*14*(1),171.https://doi.org/10.3390/math14010171

Mishra, A., & Dash, A. K. (2024). Return volatility of Asian stock exchanges; a GARCH DCC analysis with reference of Bitcoin and global crude oil price movement. *Journal of Chinese Economic and Foreign Trade Studies*, *17*(1), 29–48. https://doi.org/10.1108/JCEFTS-01-2024-0009

Nguyen, A. T. H., Nguyen, P. T. T., Pham, U. T. T., Le, A. T., Le, A. T. N., Tran, T. P., & Nguyen, O. V. H. (2026). Dynamic volatility spillovers during market turbulence: examining the interconnectedness among gold, oil, cryptocurrencies and stock markets. *Journal of International Economics and Management*, 1–16. https://doi.org/10.1108/JIEM-04-2025-0003

Ni, J., & Xu, Y. (2023). Forecasting the Dynamic Correlation of Stock Indices Based on Deep Learning Method. *Computational Economics*, *61*(1), 35–55. https://doi.org/10.1007/s10614-021-10198-3

Oyewola, D. O., Akinwunmi, S. A., & Omotehinwa, T. O. (2024). Deep LSTM and LSTM-Attention Q-learning based reinforcement learning in oil and gas sector prediction. *Knowledge-Based Systems*, *284*, 111290. https://doi.org/10.1016/j.knosys.2023.111290

Rahmi, A., & Haryanto, T. (2025). Examining the link between world oil price volatility and sectoral stock performance in Indonesia. *Cogent Economics & Finance*, *13*(1). https://doi.org/10.1080/23322039.2025.2535484

Sen, A., & Dutta Choudhury, K. (2024). Forecasting the Crude Oil prices for last four decades using deep learning approach. *Resources Policy*, *88*, 104438. https://doi.org/10.1016/j.resourpol.2023.104438

Siami-Namini, S., Tavakoli, N., & Namin, A. S. (2019). The performance of LSTM and BiLSTM in forecasting time series. In *Proceedings of the 2019 IEEE International Conference on Big Data* (pp. 3285–3292). IEEE. https://doi.org/10.1109/BigData47090.2019.9005997

Su, Y., & Feng, X. (2026). Machine learning approaches to predicting energy price correlation: From a responsible AI perspective. *Technological Forecasting and Social Change*,*225*,124515.https://doi.org/10.1016/j.techfore.2025.124515

Wang, X. (2024). Index prediction model based on LASSO-PCA and deep learning. *International Journal of Crowd Science*. https://doi.org/10.26599/IJCS.2024.9100013

Yang, A. (2025). Big data-driven corporate financial forecasting and decision support: A study of CNN-LSTM machine learning models. *Frontiers in Applied Mathematics and Statistics*, *11*. https://doi.org/10.3389/fams.2025.1566078

Zhang, X., et al. (2024). CNN-CBAM-LSTM: Enhancing stock return prediction through long and short information mining. *Mathematics*, *12*(23), 3738. https://doi.org/10.3390/math12233738

Zheng, L., Sun, Y., & Wang, S. (2024). A novel interval-based hybrid framework for crude oil price forecasting and trading. *Energy Economics*, *130*, 107266. https://doi.org/10.1016/j.eneco.2023.107266

*name of corresponding author