



Publish open access at [Caravel Press](#)

## Artificial Intelligence for Sustainable Cities

Journal homepage: [aiforsustainablecities.com](http://aiforsustainablecities.com)



Technical article

# An AI-Based Hybrid Prophet–LSTM Model for Forecasting and Financial Optimisation in Sustainable Energy Grids

Tajul Rosli Razak <sup>2,3,4\*</sup>, Ahmed El Alaoui <sup>1</sup>, Hasila Jarimi <sup>3,4</sup>, Ahmad Zia Ul-Saufie <sup>2</sup>, Mohammad Hafiz Ismail <sup>5</sup> and Mohd Faris Mohd Fuzi <sup>5</sup>

<sup>1</sup> *Emines – School of Industrial Management, Mohammed VI Polytechnic University (UM6P), Ben Guerir, Morocco.*

<sup>2</sup> *Faculty of Computer and Mathematical Sciences, Universiti Teknologi MARA (UiTM), Shah Alam, Malaysia.*

<sup>3</sup> *Solar Energy Research Institute, The National University of Malaysia, Bangi, Malaysia.*

<sup>4</sup> *Department of Architecture and Built Environment, University of Nottingham, Nottingham NG7 2RD, United Kingdom.*

<sup>5</sup> *Faculty of Computer and Mathematical Sciences, Universiti Teknologi MARA (UiTM), Arau, Perlis, Malaysia.*

## ABSTRACT

Accurate and reliable forecasting is critical for sustainable energy grid planning, as fluctuating demand, market volatility, and policy uncertainty pose significant challenges to both operational and financial stability. This paper proposes a novel AI-based hybrid forecasting framework that integrates Facebook Prophet, for interpretable long-term trend and seasonality decomposition, with Long Short-Term Memory (LSTM) networks, for modelling nonlinear short-term residual dynamics. This design explicitly addresses the dual requirement of transparency and high predictive accuracy in complex, non-stationary energy systems. The framework is evaluated using 18 years of historical financial data from Tenaga Nasional Berhad (TNB), Malaysia's largest electricity utility, as a representative large-scale grid operation case study. Performance is benchmarked against ARIMA, standalone Prophet, and standalone LSTM models using RMSE, MAE, MAPE, and SMAPE, with statistical significance assessed via the Diebold–Mariano test and robustness examined under varying forecast horizons and noise perturbations. Results show that the proposed hybrid Prophet–LSTM model achieves up to 15% lower RMSE and MAPE than the best-performing baseline while maintaining stable performance under adverse conditions. The findings demonstrate that the proposed framework provides a robust, interpretable, and modular decision-support tool for utility operators, energy planners, and policymakers, enabling improved financial optimisation, tariff planning, and operational resilience in sustainable energy grid systems.

## ARTICLE INFO

### Article History:

Received: 16 January 2026

Revised: 30 January 2026

Accepted: 6 February 2026

### Keywords:

Hybrid Prophet–LSTM  
Energy finance forecasting  
Robust forecasting  
Model benchmarking  
Explainable artificial intelligence  
Sustainable energy grids

### Article Citation:

Razak, T. R., Alaoui, A. E., Jarimi, H., Ul-Saufie, A. Z., Ismail, M. H., & Fuzi, M. F. M. (2026). An AI-Based Hybrid Prophet–LSTM Model for Forecasting and Financial Optimisation in Sustainable Energy Grids. *Artificial Intelligence for Sustainable Cities*, 1(1), 42–60.

<https://doi.org/10.65582/aifsc.2026.004>

## 1. INTRODUCTION

Accurate and reliable financial time series forecasting is vital for energy utilities operating in volatile markets shaped by fluctuating demand, evolving regulations, energy transition policies, and macroeconomic

\* Corresponding author. Email address: [tajulrosli@uitm.edu.my](mailto:tajulrosli@uitm.edu.my) (Tajul Rosli Razak)



© 2026 The Author(s).

Published by Caravel Press

shocks. Under such conditions, forecasting systems must provide high predictive accuracy while ensuring robustness, transparency, and reproducibility—qualities essential for tariff design, investment planning, and operational decision-making in sustainable grid environments.

Traditional statistical models, such as the Autoregressive Integrated Moving Average (ARIMA) model, have been widely applied in energy finance forecasting for many years. While effective for stationary and linear processes, however, have restrictive assumptions that limit their ability to capture nonlinear dependencies, regime shifts, and structural market changes (Box et al., 2015; Torres et al., 2021; Weron, 2014). Advances in artificial intelligence (AI) have introduced deep learning models—particularly Long Short-Term Memory (LSTM) networks—that excel at modelling long-term dependencies, non-stationary behaviours, and high-frequency volatility (Zhang, 2003; Shen et al., 2021; Lago et al., 2018). However, their “black-box” nature hinders interpretability, reducing stakeholder confidence in regulated utility contexts where transparent decision support is crucial.

More recently, transformer-based architectures have emerged as state-of-the-art approaches for time-series forecasting, leveraging attention mechanisms to capture long-range temporal dependencies and complex variable interactions. These models have demonstrated strong performance in multivariate forecasting tasks and large-scale datasets. Nevertheless, their increased architectural complexity and reduced transparency pose challenges for regulated energy applications, where explainability, auditability, and reproducibility are often prioritised alongside accuracy.

To address these challenges, this study proposes a hybrid forecasting framework that integrates Facebook Prophet’s interpretable decomposition of trend, seasonality, and holiday effects with LSTM’s nonlinear residual learning, building upon established decomposition-based and hybrid forecasting strategies reported in the literature (Taylor & Letham, 2018; Zhang, 2003). This design combines structural interpretability with high-frequency predictive capability, directly tackling the challenges of non-stationarity, volatility, and regime transitions in energy finance.

The framework is validated using financial time series from Tenaga Nasional Berhad (TNB), Malaysia’s largest electricity utility, providing a representative case study for utility-scale financial planning. Evaluation follows a rigorous benchmarking protocol, incorporating multiple error metrics (RMSE, MAE, MAPE, SMAPE), statistical significance testing via the Diebold–Mariano test, and robustness checks across varying horizons and noise conditions. Comparisons are made against ARIMA, Prophet, and LSTM baselines, consistent with best practices in forecasting evaluation (Makridakis et al., 2018).

In regulated electricity markets, utility financial performance is closely tied to the feasibility and timing of infrastructure investments that underpin urban grid reliability and decarbonisation efforts. Forecasts of utility valuation and financial resilience can therefore inform engineering planning decisions indirectly by supporting tariff-setting justifications, capital expenditure (CAPEX) scheduling, procurement planning, and risk-aware financing of grid modernisation projects (e.g., renewable integration, storage deployment, and network reinforcement). In this context, improving forecast accuracy reduces uncertainty in financial planning and strengthens the evidence base for sustainability-oriented investment decisions in urban energy systems.

While hybrid forecasting models have been explored in the broader financial time-series literature (Makridakis et al., 2018; Zhang, 2003), their systematic deployment as interpretable, statistically validated decision-support tools for sustainable energy grid financial planning remains limited. Rather than proposing a new deep learning architecture, the originality of this work lies in the purposeful integration of Prophet and LSTM within a policy- and utility-oriented context, supported by a rigorous and reproducible benchmarking framework. This framing explicitly links predictive performance to practical energy-sector decision-making requirements, such as tariff planning, investment resilience, and regulatory transparency.

The contributions of this work are fourfold:

1. Development of a tailored Prophet–LSTM framework for utility-scale forecasting that balances accuracy and interpretability;
2. Introduction of a reproducible benchmarking protocol with multi-metric evaluation, significance testing, and robustness analysis;
3. Empirical validation showing superior performance over established baselines using real-world utility data;
4. Demonstration of decision-support capability for regulators, policymakers, and utility managers.

While the present study adopts a univariate formulation to establish a transparent and statistically controlled baseline, this design choice enables clear attribution of performance gains to the proposed hybrid architecture.

The framework is inherently modular, and the study explicitly outlines pathways toward multivariate extensions—including transformer-based and attention-driven models—that incorporate policy, market, and macroeconomic indicators to support broader grid-impact-aware forecasting.

## 2. BACKGROUND

This section reviews the theoretical and methodological foundations relevant to the measurement and benchmarking of forecasting accuracy in sustainable energy grid financial planning. Financial time series forecasting remains central in the energy and finance sectors, where challenges such as non-stationarity, volatility, and structural regime shifts necessitate the use of advanced models and rigorous evaluation frameworks. Accurate forecasts are vital for tariff planning, investment scheduling, and risk management, since even minor deviations can generate substantial operational and financial consequences (Zheng et al., 2024; Mohammadi et al., 2024). Traditional models, such as ARIMA, have been widely applied due to their effectiveness in stationary environments (Rizvi, 2024; Peng et al., 2023; Cheng et al., 2020). However, their linear structure constrains their ability to capture the nonlinear dynamics, volatility clustering, and structural breaks typical in financial markets (Parviz & Ghorbanpour, 2024). To overcome these shortcomings, hybrid approaches combining statistical decomposition with machine learning have emerged. Early ARIMA-ANN models demonstrated the potential of such integration (Liu et al., 2012). More recent methods, including Prophet and LSTM, offer interpretable trend and seasonality modelling, as well as the capacity to capture nonlinear temporal dependencies (Shen et al., 2021; Pan et al., 2024). Despite progress, systematic benchmarking of hybrid models in sustainable energy finance—incorporating multi-metric evaluation, statistical significance testing, and robustness analysis—remains underexplored.

Probabilistic models, such as Hidden Markov Models (HMMs), have historically provided valuable insights into regime-switching behaviour and latent state dynamics (Rabiner, 1989), making them helpful in identifying structural shifts in financial markets. However, HMMs face significant limitations. Their reliance on the Markov assumption constrains their ability to capture long-range dependencies (Avinash et al., 2024), while non-stationary behaviours such as volatility clustering and evolving correlations are not adequately addressed (Xu et al., 2024). Extending HMMs to handle nonlinearities increases complexity, rendering estimation computationally demanding and often unstable (Nystrup et al., 2017). Consequently, their forecasting performance is usually surpassed by more flexible approaches, including LSTMs and hybrid statistical-machine learning frameworks, which better capture nonlinearities and adapt to evolving temporal dependencies (Hamilton, 2010; He, 2023; Liu et al., 2025).

Recurrent Neural Networks (RNNs) and their advanced variant, Long Short-Term Memory (LSTM) networks, have demonstrated strong performance in sequential modelling tasks by capturing temporal dependencies and nonlinear patterns (Pascanu et al., 2012; Cohen et al., 2021). LSTMs mitigate vanishing gradient issues by utilising gating mechanisms, allowing them to retain long-term context and model complex temporal dynamics effectively. When integrated into hybrid frameworks, such as Prophet-LSTM, LSTMs complement Prophet's interpretable decomposition of trend, seasonality, and events by focusing on residual dynamics (Zhang, 2003; Thomson et al., 2019). This complementary division of modelling tasks has consistently improved forecasting accuracy. Nevertheless, many studies have yet to embed these models within comprehensive performance measurement frameworks that include statistical testing and robustness checks, which are essential for real-world reliability.

Recent research has increasingly emphasised hybrid models that exploit the strengths of both statistical and machine learning paradigms. Prophet-LSTM frameworks, for instance, combine interpretability with the ability to capture nonlinear residual dynamics, consistently outperforming single-model approaches (Oukhouya et al., 2024; Peng et al., 2021). However, beyond accuracy, practical deployment in sustainable energy finance requires interpretability, transparency, and reproducibility. Visualisation techniques, including interactive dashboards and annotated time series plots, provide a critical link between technical outputs and actionable insights, enhancing decision-making by illustrating forecast trajectories, uncertainty intervals, residual patterns, and anomalies (Mohamed et al., 2022; Hsu et al., 2016). When embedded within benchmarking frameworks, such hybrid approaches not only deliver predictive accuracy but also facilitate comprehensive multi-metric evaluation, statistical validation, and robustness analysis (Lim & Zohren, 2021; Wang et al., 2025). Collectively, these developments frame the motivation for the present study, which proposes and rigorously evaluates a hybrid Prophet-LSTM model to address existing gaps in sustainable

energy finance forecasting.

### 3. PROPOSED HYBRID FORECASTING AND PERFORMANCE MEASUREMENT FRAMEWORK


This study introduces an AI-based hybrid forecasting framework that delivers accurate, interpretable, and robust financial forecasts for sustainable energy grid applications. The framework integrates Facebook Prophet, which decomposes financial time series into trend, seasonality, and holiday effects, with Long Short-Term Memory (LSTM) networks that model the nonlinear residual patterns left unexplained by Prophet. By explicitly separating structural components from residual dynamics and recombining their outputs, the approach leverages Prophet's transparency alongside LSTM's ability to capture nonlinear and high-frequency temporal behaviour.

The framework is tailored to the operational needs of energy utilities, where financial planning, tariff design, and capital allocation require not only predictive accuracy but also interpretability and methodological transparency. Rather than applying LSTM directly to raw price series, the residual-based learning strategy ensures that the deep learning component focuses on volatility, regime transitions, and nonlinear dependencies that linear diagnostics alone cannot adequately capture. In addition to the hybrid model design, the framework incorporates a rigorous, reproducible benchmarking protocol, including multi-metric evaluation, statistical significance testing, and robustness analysis across varying forecast horizons and noise conditions. Together, these elements provide a reliable decision-support tool for utility finance planning in volatile and policy-sensitive energy markets.

To ensure a fair and interpretable benchmark, Prophet is treated as a standalone structural baseline that forecasts on the original price scale by modelling trend, seasonality, and (when applicable) calendar effects. In this study, Prophet is not intended to capture high-frequency volatility; instead, it provides an auditable decomposition stage whose residuals are subsequently modelled by the LSTM in the hybrid architecture. This design choice clarifies that large Prophet-only errors during volatile periods reflect model scope (structural vs. volatility dynamics) rather than misconfiguration. Unless otherwise stated, Prophet was fitted with standard recommended settings and automatic changepoint selection, and all models were evaluated using the same data partitioning and preprocessing pipeline.

#### 3.1. FRAMEWORK OVERVIEW

The proposed hybrid forecasting framework consists of three main stages (Fig. 1) that integrate the additive decomposition capabilities of Facebook Prophet with the nonlinear sequence modelling power of Long Short-Term Memory (LSTM) networks. This architecture explicitly separates low-frequency structural components (trend, seasonality, holiday effects) from high-frequency nonlinear fluctuations, enabling improved predictive accuracy, interpretability, and robustness in financial forecasting for sustainable energy grids.



**Figure 1.** Workflow of the proposed hybrid Prophet-LSTM framework. Prophet decomposes the time series into structural and residual components; LSTM models the residuals, and the outputs are combined to produce the final forecast.

Stage 1 – Decomposition with Prophet: Prophet models a time series  $y(t)$  as:

$$y(t) = g(t) + s(t) + h(t) + \epsilon t \quad (1)$$

Where:  $g(t)$  captures long-term non-periodic trends (e.g., linear or logistic growth),  $s(t)$  models periodic fluctuations via a Fourier series expansion,  $h(t)$  accounts for holiday or event-based effects, and  $\epsilon t$  is the unexplained error term.

After fitting Prophet, the residual sequence is extracted:

$$r(t) = y(t) - \hat{y}_{Prophet}(t) \quad (2)$$

These residuals contain short-term, nonlinear, and stochastic patterns not captured by the additive model.

Stage 2 – Residual Modelling with LSTM: The residual sequence  $r(t)$  is passed to an LSTM network to model nonlinear temporal dependencies and capture volatility patterns. The LSTM operations can be expressed as:

$$f_t = \sigma(W_f [h_{t-1}, x_t] + b_f), \quad i_t = \sigma(W_i [h_{t-1}, x_t] + b_i), \quad (3)$$

$$o_t = \sigma(W_o [h_{t-1}, x_t] + b_o) \quad C^{\sim} t = \tanh(W_c [h_{t-1}, x_t] + b_c), \quad (4)$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t, \quad h_t = o_t \odot \tanh(C_t), \quad (5)$$

Where  $x_t = r(t)$  is the input residual,  $h_t$  is the hidden state,  $\sigma(\cdot)$  and  $\tanh(\cdot)$  are the activation functions, and  $\odot$  denotes element-wise multiplication. The output  $\hat{r}_{LSTM}(t)$  represents the predicted residuals.

Stage 3 – Hybrid Forecast Reconstruction: The final forecast is obtained by summing Prophet's structural component predictions with LSTM's predicted residuals:

$$\hat{y}_{Hybrid}(t) = \hat{y}_{Prophet}(t) + \hat{y}_{LSTM}(t) \quad (6)$$

This additive fusion preserves Prophet's interpretability while incorporating LSTM's adaptability to nonlinear dynamics.

The modular design ensures that each stage addresses a distinct aspect of the forecasting problem: Prophet captures systematic planning signals (e.g., regulatory adjustments, investment cycles), while LSTM models market volatility and stochastic variations. Embedding the architecture within a comprehensive performance measurement protocol—including multi-metric evaluation, statistical significance testing, and robustness checks—ensures high predictive performance, reproducibility, and interpretability for real-world utility finance applications.

### 3.2. RATIONALE FOR MODEL SELECTION

The proposed hybrid framework combines Facebook Prophet and Long Short-Term Memory (LSTM) networks, chosen for their complementary strengths in interpretability, adaptability, and predictive accuracy for financial forecasting in sustainable energy grids.

Prophet is selected for its ability to decompose time series into interpretable components—trend, seasonality, and holiday effects—while remaining robust to missing values, outliers, and irregular sampling intervals Taylor & Letham, 2018). Its additive modelling structure produces transparent diagnostics of long-term patterns, enabling decision-makers to trace forecast changes back to identifiable drivers. In regulated utility environments, this level of explainability is crucial for tariff planning, budget allocation, and investment scheduling, where forecasts must be reproducible and auditable.

LSTM networks, in contrast, are designed to capture complex nonlinear relationships and long-range temporal dependencies in sequential data Hochreiter & Schmidhuber (1997). By focusing exclusively on the residual component left unexplained by Prophet, the LSTM can model high-frequency variations, stochastic fluctuations, and short-term anomalies arising from market volatility, policy changes, or demand shocks. This

division of labour allows Prophet to specialise in low-frequency structural dynamics, while LSTM addresses high-frequency nonlinearities.

Alternative hybrid approaches—such as ARIMA–ANN, wavelet-based hybrids, or purely deep learning models—were evaluated but found to be less well aligned with the project’s requirements. These alternatives either reduced interpretability, handled irregular time series less effectively, or introduced greater computational overhead without commensurate gains in accuracy. In contrast, the Prophet–LSTM combination achieves a balanced integration of interpretability and adaptability, while supporting a rigorous performance measurement protocol that includes multi-metric evaluation, statistical significance testing, and robustness analysis. This makes the chosen hybrid framework well-suited for engineering-focused financial forecasting in sustainable energy grids, where methodological rigour and practical applicability are paramount.

### 3.3. HYPERPARAMETER TUNING FOR LSTM

Hyperparameter tuning for the LSTM residual modelling stage was performed using an exhaustive grid search to maximise predictive accuracy while preventing overfitting. The search space encompassed network architecture parameters, regularisation settings, and training configurations, including the number of LSTM layers, the number of neurons per layer, activation functions, dropout rates, optimisers, learning rates, batch sizes, and training epochs. The final selected configuration is presented in Table 1.

**Table 1.** Selected Hyperparameters for the LSTM Residual Model.

Hyperparameter	Selected Value
Number of LSTM layers	2
Neurons per layer	50
Activation function	ReLU
Dropout rate	0.2
Optimizer	Adam
Learning rate	0.001
Batch size	32
Epochs	100 (with early stopping)

The final architecture comprises two stacked LSTM layers, enabling the model to capture both short- and long-term temporal dependencies in the residual sequence. Each layer contains 50 neurons and employs the ReLU activation function, which was selected for its computational efficiency and ability to accelerate convergence during training. A dropout rate of 0.2 was applied to mitigate overfitting, randomly deactivating a fraction of neurons during each training step.

The Adam optimiser was chosen for its adaptive learning rate mechanism, which allows for efficient gradient-based optimisation without the need for extensive manual adjustments. The learning rate was fixed at 0.001, as determined by the performance of the validation set. Training used a batch size of 32, with early stopping triggered if validation loss did not improve for 20 consecutive epochs. This configuration balances model expressiveness with generalisation capability, ensuring robust performance on unseen utility finance data.

All hyperparameter settings, tuning ranges, and selection criteria were fully documented to enable reproducibility and facilitate fair benchmarking against baseline forecasting models.

### 3.4. MODEL EVALUATION METRICS

The performance of the proposed hybrid Prophet-LSTM model and all benchmark models was assessed using a multi-criteria evaluation protocol that combined complementary accuracy metrics, statistical significance testing, and robustness analysis. To improve clarity for a multidisciplinary audience, the presentation of standard evaluation metrics is kept concise, with emphasis placed on interpretation rather than detailed derivation. This approach ensures that performance claims are statistically valid and practically relevant for sustainable energy grid applications.

Root Mean Squared Error (RMSE):

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (7)$$

RMSE measures the magnitude of forecast errors and is highly sensitive to large deviations. This makes it particularly relevant for utility rate projections or capital expenditure planning, where substantial errors can translate into significant financial risks.

Mean Absolute Error (MAE):

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (8)$$

MAE reports the average absolute error, providing an interpretable, scale-dependent measure less influenced by outliers than RMSE.

Mean Absolute Percentage Error (MAPE):

$$MAPE = \frac{100}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right| \quad (9)$$

MAPE expresses error as a percentage, aiding communication with policymakers, financial analysts, and regulatory bodies. Its instability near zero actual values is noted and addressed through complementary metrics.

Symmetric Mean Absolute Percentage Error (SMAPE):

$$SMAPE = \frac{100}{n} \sum_{i=1}^n \frac{|y_i - \hat{y}_i|}{(|y_i| + |\hat{y}_i|)/2} \quad (10)$$

SMAPE mitigates MAPE's instability by symmetrically scaling the error, improving robustness when actual values vary widely.

**Statistical Significance Testing:** The Diebold–Mariano (DM) test was applied to determine whether differences in forecast accuracy between competing models are statistically significant, ensuring that performance improvements are not attributable to random variation.

**Robustness Checks:** To evaluate resilience under realistic operational uncertainty:

- Forecast horizons were varied to test stability across short- and long-term predictions.
- Synthetic noise perturbations were introduced to simulate data uncertainty.

By combining multiple error metrics with formal statistical testing and robustness checks, this evaluation protocol provides a rigorous, reproducible, and engineering-relevant basis for determining the suitability of forecasting models in sustainable energy grid financial planning.

### 3.5. CONTRIBUTION AND NOVELTY

This study presents an AI-based hybrid forecasting framework that integrates Facebook Prophet with Long Short-Term Memory (LSTM) networks, specifically designed for accurate, interpretable, and robust financial forecasting in sustainable energy grids. Prophet decomposes the economic time series into structured components—trend, seasonality, and holiday effects—providing transparent insights into systematic patterns. LSTM is then applied to the residual series to model high-frequency, nonlinear, and stochastic variations. This residual-driven, modular design decouples macro-pattern recognition from micro-dynamic modelling, enhancing adaptability to abrupt market shifts caused by policy changes, commodity price volatility, or demand-side fluctuations.

The novelty of this work lies not only in the hybrid model architecture but also in its integration with a rigorous performance measurement and benchmarking protocol. This includes multi-metric accuracy evaluation (RMSE, MAE, MAPE, SMAPE) to capture different error perspectives, statistical significance testing via the Diebold–Mariano test to confirm that improvements are not attributable to random variation, robustness analysis under varying forecast horizons and synthetic noise perturbations, and complete

documentation of data preprocessing, model configuration, and tuning procedures to ensure reproducibility.

Empirical validation using stock price data from Tenaga Nasional Berhad (TNB), Malaysia's largest electricity utility, demonstrates that the proposed hybrid model significantly outperforms benchmark methods—ARIMA, standalone Prophet, and standalone LSTM—achieving the lowest RMSE and MAPE scores with statistically significant gains. The results confirm the framework's capability to deliver dependable forecasts that maintain accuracy across diverse and volatile market conditions.

Beyond predictive performance, this research advances the literature on intelligent, explainable, and policy-relevant forecasting for the energy sector. The proposed method equips utility managers, regulators, and financial planners with a decision-support tool that balances interpretability and accuracy, enabling more informed tariff setting, investment planning, and fiscal risk mitigation. The framework is inherently scalable for multivariate forecasting, accommodating additional drivers such as commodity prices, carbon market indicators, and macroeconomic variables. It can also be extended using explainable AI (XAI) techniques, such as SHAP or LIME. These future enhancements will further strengthen its role in supporting transparent, trustworthy, and evidence-based decision-making in sustainable energy finance.

In summary, the proposed Prophet-LSTM framework is purpose-built to address the dual challenge of interpretability and predictive accuracy in sustainable energy grid financial forecasting. This approach ensures both methodological rigour and practical applicability by decomposing structural and residual components, optimising the LSTM architecture for residual dynamics, and embedding the model within a multi-metric, statistically validated, and robustness-tested evaluation protocol. Integrating transparent statistical modelling with adaptive deep learning makes the framework highly relevant to engineering decision-making in policy-sensitive, data-variable environments. Furthermore, the explicit documentation of configuration and tuning procedures ensures reproducibility, enabling the reliable adoption of advanced AI solutions by utility operators, energy planners, and researchers seeking to deploy these solutions in real-world grid finance applications.

### 3. EXPERIMENTAL SETUP AND RESULTS

This section outlines the experimental setup and presents the results of the proposed hybrid Prophet- LSTM forecasting framework. The workflow follows six sequential stages:

1. Data collection and pre-processing
2. Exploratory data analysis (EDA)
3. Statistical diagnostics
4. Model implementation and forecasting
5. Performance benchmarking
6. Statistical significance testing.

Each stage is structured to ensure transparency, reproducibility, and methodological rigour, providing a robust basis for evaluating forecasting performance in sustainable energy grid financial planning. The experimental objectives are to characterise the underlying financial time series in terms of trend, seasonality, volatility, and residual dynamics; to evaluate the predictive accuracy, stability, and robustness of the proposed hybrid model against established benchmark methods; and to confirm the statistical validity of observed performance differences to rule out random variation.

By integrating multi-metric accuracy measurement, robustness analysis, and formal statistical validation, this experimental framework addresses common limitations of prior studies that rely solely on point forecast accuracy. The results presented herein demonstrate the proposed framework's predictive capability and its practical applicability as a decision-support tool for utility financial planning in sustainable energy grids.

#### 4.1. DATA COLLECTION AND PREPROCESSING

The dataset employed in this study comprises historical daily trading data for Tenaga Nasional Berhad (TNB) (Ticker: 5347.KL), Malaysia's largest electricity utility, sourced from the publicly accessible Yahoo Finance database. The selected period spans from May 2006 to August 2024, providing over 18 years of continuous financial records and covering multiple market cycles, regulatory changes, and macroeconomic events. The dataset contains the following variables: Open, High, Low, Close, Adjusted Close, and Volume.


The closing price was chosen as the primary forecasting variable. Widely regarded as a reliable indicator

of a stock's market valuation at the end of a trading day, the closing price effectively reflects long-term price trends while minimising the influence of short-lived intraday fluctuations. This makes it well-suited for strategic financial planning, tariff modelling, and investment decision-making in the energy utility sector, where the focus is on medium- to long-term projections rather than high-frequency trading.

Using a recognised and publicly available data source ensures complete transparency and reproducibility of results. Moreover, the long historical window allows the proposed Prophet-LSTM framework to learn from a wide range of operating conditions, including periods of regulatory reform, fuel price volatility, and demand-side shocks—factors directly relevant to financial forecasting for sustainable energy grids.

#### 4.2. EXPLORATORY DATA ANALYSIS

Figure 2 presents the daily closing prices of Tenaga Nasional Berhad over 18 years, revealing several notable patterns. A sustained upward trend is observed between 2013 and 2017, coinciding with economic expansion and stable regulatory conditions in Malaysia's energy market. Sharp contractions occurred during the 2008 global financial crisis and the 2020 COVID-19 pandemic, reflecting macroeconomic shocks that directly impacted utility valuations. More recently, elevated volatility after 2020 appears to be driven by global energy price fluctuations, supply chain disruptions, and the acceleration of energy transition policies.



**Figure 2.** Daily closing prices of Tenaga Nasional Berhad (TNB) from 2006-05-09 to 2024-08-30, capturing multiple market cycles, external shocks, and structural changes in the Malaysian energy sector.

These dynamics highlight the dual modelling challenge: capturing long-term structural trends while adapting to short-term, high-frequency fluctuations. Prophet's decomposition framework is suited for isolating interpretable trend and seasonality components, whereas LSTM is designed to model the nonlinear residual patterns that dominate during volatile periods.

The statistical summary in Table 2 shows moderate volatility, with a standard deviation of 3.15 MYR relative to a mean of 10.04 MYR. The distribution is right-skewed, indicating that higher-price occurrences are more frequent. At the same time, the wide range between the minimum (4.56 MYR) and maximum (16.24 MYR) underscores the influence of external shocks, such as regulatory reforms, global fuel price movements, and macroeconomic instability.

**Table 2.** Statistical summary of TNB closing prices


Statistic	Value
Count	4525
Mean	10.04
Standard Deviation	3.15
Minimum	4.56
25% Quantile	6.95
50% Quantile (Median)	9.62
75% Quantile	13.12
Maximum	16.24

The diversity of market regimes in this dataset provides a robust basis for model evaluation, allowing the proposed hybrid Prophet-LSTM framework to be tested under conditions of stability, growth, and volatility. This ensures that performance benchmarking reflects accuracy and robustness in realistic operational contexts for sustainable energy grid financial forecasting.

#### 4.3. STATISTICAL ANALYSIS RESULTS


This subsection examines the statistical properties of the TNB closing price series using decomposition, frequency-domain analysis, distributional diagnostics, and autocorrelation measures. The aim is to identify structural characteristics that justify the hybrid modelling approach and ensure the proposed framework is data-driven and theoretically grounded.

Figure 3 reveals a clear long-term upward trend, minimal seasonal variation, and residual components characterised by irregular short-term fluctuations. This structure—trend-dominated with volatile short-run noise—supports the division of modelling tasks between Prophet (trend and seasonality) and LSTM (nonlinear residuals).




**Figure 3.** Time series decomposition of TNB closing prices into trend, seasonality, and residual components.

The periodogram in Figure 4 confirms that most variance is concentrated in low-frequency cycles, validating trend-based models such as Prophet for long-term structural forecasting in energy finance.




**Figure 4.** Periodogram of TNB closing prices, showing dominance of low-frequency components.

In contrast, Figure 5 shows that log-returns exhibit considerably high-frequency variance, reflecting short-lived market shocks. This reinforces the need for sequential models, such as LSTM, to capture rapid fluctuations in data.




**Figure 5.** Periodogram of log-returns, indicating substantial high-frequency components.

Figure 6 depicts a leptokurtic, heavy-tailed distribution centred near zero, suggesting a propensity for extreme events. Gaussian-based models may underestimate such tail risk, making flexible nonlinear approaches more appropriate.




**Figure 6.** Distribution of log-returns showing leptokurtosis and heavy tails.

Figures 7 and 8 confirm that the normal and exponential distributions fail to capture the observed tail behaviour, whereas the gamma and chi-squared distributions provide a better fit. This further supports the adoption of AI-driven approaches that do not rely on strict parametric distributional assumptions.




**Figure 7.** QQ plots of log-returns against selected theoretical distributions (Set 1).



**Figure 8.** QQ plots of log-returns against selected theoretical distributions (Set 2).

Figure 9 shows negligible autocorrelation beyond lag 0, consistent with weak linear memory in returns. However, small but statistically significant spikes at short lags indicate limited short-term dependencies, justifying the use of sequential models to capture these effects.




**Figure 9.** Autocorrelation (ACF) and partial autocorrelation (PACF) of log-returns.


**Summary and Implications for Modelling:** The statistical analysis reveals a dual structure, with persistent long-term trends coexisting alongside short-term volatility and heavy-tailed return distributions. These empirical characteristics align directly with the proposed Prophet–LSTM architecture: Prophet captures interpretable structural components, while LSTM models nonlinear, high-frequency residuals. By explicitly grounding the modelling strategy in statistical diagnostics, the framework is well-matched to the dataset's observed behaviour, ensuring both theoretical validity and practical applicability.

#### 4.4. MODELLING AND FORECASTING RESULTS

This subsection compares the performance of the standalone Prophet model, the standalone LSTM model, and the proposed hybrid Prophet–LSTM framework. The objective is to benchmark predictive accuracy and evaluate each model's ability to capture long-term structural trends and short-term nonlinear fluctuations in the financial time series.

**Analysis of Results:** The residuals from the Prophet model (Fig. 10a) are centred around zero, indicating a good fit to long-term trends. However, volatility spikes after 2020 highlight its limited adaptability to structural breaks and sudden market swings. The Prophet forecast (Fig. 10b) effectively captures slow-moving trends but systematically underestimates prices during high-volatility periods. The LSTM forecast of price levels (Fig. 10c) tracks overall directional movement but suffers from over-smoothing, reducing responsiveness to abrupt changes. Similarly, the LSTM forecast of log-returns (Fig. 10d) captures the general shape of observed returns but underestimates amplitude during turbulent periods, suggesting underfitting of extreme events. By contrast, the Hybrid Prophet–LSTM forecast (Fig. 10e) achieves improved alignment with both structural trends and short-term volatility; although minor lag effects remain during extreme jumps, forecast bias is consistently reduced relative to standalone models.





**Figure 10.** Model performance comparison: (a) Residuals from Prophet on training data; (b) Prophet forecast vs. actual (test set); (c) LSTM forecast (closing prices); (d) LSTM forecast (log-returns); (e) Hybrid Prophet–LSTM forecast.

**Key Findings:** The hybrid Prophet–LSTM model achieves a better balance between interpretability and adaptability than either component model alone. Prophet’s decomposition yields transparent trend and seasonality estimates, while the LSTM improves short-term accuracy by modelling nonlinear residual dynamics. This dual approach mitigates Prophet’s volatility handling limitations and LSTM’s over-smoothing tendency.

From an engineering decision-support perspective, these findings demonstrate the value of combining interpretable statistical decomposition with nonlinear residual learning to enhance predictive reliability in energy-related financial forecasting. Potential extensions include further incorporating exogenous variables (e.g., commodity prices, policy indicators) or integrating volatility-sensitive neural architectures to improve responsiveness during market shocks.

#### 4.5. MODEL PERFORMANCE SUMMARY

It is important to note that the standalone Prophet model serves as an interpretable structural benchmark rather than a volatility-sensitive forecaster. Consequently, during periods of abrupt regime shifts and high-frequency market fluctuations, unmodelled residual volatility can accumulate into larger reconstruction errors when Prophet is used alone on the original price scale. The hybrid model is explicitly designed to address this limitation by learning these residual dynamics via LSTM.

Table 3 summarises the forecasting performance of all models using both scale-dependent (RMSE, MAE) and scale-independent (MAPE, SMAPE) metrics. For fairness and consistency, all models are evaluated on the same de-normalised original price scale. The Hybrid Prophet–LSTM clearly outperforms all baselines, achieving the lowest RMSE (0.0700), MAPE (1.29%), MAE (0.0560), and SMAPE (2.55%), together with the narrowest confidence interval, indicating superior accuracy and stability.

**Table 3.** Model performance summary with extended metrics. All metrics are computed on the same de-normalised (original price) scale.

Model	RMSE	MAE	MAPE (%)	SMAPE (%)	95% CI (RMSE)
ARIMA	0.4800	0.3840	3.56	6.88	[0.45, 0.51]
Prophet	10.1400	8.1120	22.69	36.99	[9.80, 10.50]
LSTM	0.3364	0.2691	2.34	4.57	[0.31, 0.36]
Hybrid Model	0.0700	0.0560	1.29	2.55	[0.06, 0.08]

The standalone LSTM ranks second (RMSE = 0.3364, MAPE = 2.34%), demonstrating strong nonlinear modelling capability but lacking the hybrid model's ability to capture long-term trend and short-term residual dynamics jointly. ARIMA shows moderate performance (RMSE = 0.4800, MAPE = 3.56%), reflecting the limitations of purely linear time-series modelling in highly volatile financial energy data.

Prophet records substantially higher errors (RMSE = 10.14, MAPE = 22.69%). This is not due to a difference in the evaluation scale or inconsistent preprocessing, but instead to Prophet's structural limitation in modelling high-frequency nonlinear volatility and abrupt regime changes present in the price series when used as a standalone forecaster. While Prophet effectively captures global trend and seasonal components, the unmodelled residual volatility accumulates into significant reconstruction errors when evaluated on the original price scale.

Overall, these results confirm that combining statistical decomposition with nonlinear residual learning substantially enhances forecasting accuracy, robustness, and practical reliability. The joint use of scale-dependent and scale-independent metrics provides complementary insights, capturing both absolute financial deviation and proportional accuracy, which is critical for utility tariff modelling, investment planning, and financial risk management. Robustness and statistical significance are further examined in Section 4.6.

#### 4.6. STATISTICAL SIGNIFICANCE TESTING

The Diebold–Mariano (DM) test confirms that the observed accuracy gains are statistically significant. As shown in Table 4, LSTM outperforms Prophet, while the Hybrid Prophet–LSTM model significantly outperforms both. These results demonstrate that the hybrid model's superior performance stems from its integration of interpretable statistical decomposition with nonlinear residual learning, rather than random variation. For stakeholders such as utility planners and regulators, this provides confidence that the hybrid approach offers reliable and robust improvements over existing methods.

**Table 4.** Pairwise Diebold–Mariano (DM) test results comparing predictive accuracy among Prophet, LSTM, and the Hybrid Prophet–LSTM model. Significant results ( $p < 0.05$ ) indicate that the second model in the comparison achieves statistically different accuracy compared to the first.

Model Comparison	DM Statistic	p-value
Prophet vs. LSTM	2.45	0.014
Prophet vs. Hybrid	3.87	0.0001
LSTM vs. Hybrid	2.01	0.045

## 5. DISCUSSION

The results demonstrate that the proposed Hybrid Prophet–LSTM framework delivers superior forecasting accuracy, stability, and statistical robustness compared to ARIMA, Prophet, and standalone LSTM when applied to Tenaga Nasional Berhad's financial time series. This advantage arises from the principled integration of statistical decomposition (Prophet), which provides interpretable modelling of long-term trend and seasonality, with nonlinear residual learning (LSTM), which effectively captures high-frequency volatility. By uniting interpretability with adaptability, the framework addresses a persistent gap in energy finance forecasting, where existing models typically excel in one dimension but not both.

ARIMA achieves moderate accuracy, reflecting its strength in modelling short-term autocorrelations but limited adaptability to regime shifts and structural breaks. Prophet, while highly interpretable, underperforms due to its limited responsiveness to abrupt shocks and nonlinear dynamics. The LSTM performs strongly in capturing nonlinear patterns but lacks transparency and shows weaker performance under extreme events. In contrast, the Hybrid Prophet–LSTM consistently achieves the best performance (RMSE = 0.07, MAPE = 1.29%), outperforming all baselines and demonstrating stable behaviour across both calm and volatile market regimes.

Recent transformer-based forecasting models offer an alternative paradigm by employing attention mechanisms to capture long-range dependencies and complex interactions, particularly in multivariate settings. While such models have shown strong empirical performance in large-scale forecasting tasks, their application in regulated energy finance remains constrained by higher computational demands and reduced interpretability. In contrast, the proposed hybrid framework prioritises transparency, modularity, and reproducibility—characteristics that are critical for regulatory acceptance and policy-facing decision support.

Taken together, the proposed Hybrid Prophet–LSTM framework offers several clear advantages over existing forecasting tools. Compared with traditional statistical models such as ARIMA, it provides substantially higher accuracy and robustness under nonlinear volatility and regime shifts. Relative to standalone deep learning models (e.g., LSTM), it retains interpretability through explicit structural decomposition, which is essential for regulatory auditability and policy-facing applications. In contrast to recent transformer-based architectures, the proposed framework achieves competitive performance with lower computational complexity and greater transparency, making it more suitable for deployment in regulated energy utility environments. These comparative advantages position the proposed method as a practical and decision-oriented forecasting tool rather than a purely predictive black-box model.

Beyond statistical accuracy, the proposed framework offers tangible practical value for multiple stakeholders. Regulators can employ it for evidence-based tariff setting and fiscal risk assessment, utilities for capital investment and procurement planning, and investors for portfolio risk evaluation under market and policy uncertainty. By combining statistical transparency with deep learning adaptability, the framework supports more resilient and informed decision-making in sustainability-driven energy markets. In particular, improved financial forecast reliability supports sustainable urban grid engineering by enabling more credible long-term planning for infrastructure upgrades, renewable integration investments, and decarbonisation programmes under regulatory and market uncertainty. Accordingly, the primary interpretation of the forecasting results in this study is aligned with regulatory tariff modelling and utility financial planning, rather than short-term investor return optimisation.

From a sustainability perspective, reductions in forecast error translate directly into improved planning reliability for urban energy systems. More accurate financial forecasts reduce uncertainty in tariff setting and revenue projections, enabling utilities and regulators to allocate resources more efficiently toward renewable integration, grid reinforcement, and decarbonisation initiatives. Lower uncertainty also mitigates the risk of delayed or misaligned infrastructure investments, which can otherwise hinder progress toward sustainability targets. In this sense, forecast accuracy improvements are not merely statistical gains but function as enablers of more resilient, cost-effective, and sustainable urban energy transitions.

Specifically, improved forecast reliability can support:

1. Tariff and revenue planning to ensure cost recovery for sustainability investments,
2. CAPEX and maintenance scheduling for grid reinforcement and asset renewal,
3. Financing and risk assessment for decarbonisation initiatives such as renewable integration and storage deployment.

This study is subject to several limitations. First, the current implementation adopts a univariate approach and focuses on a single-utility case study, thereby limiting direct generalisability across heterogeneous grid environments. However, this choice is intentional and allows the proposed framework to establish a clear, interpretable, and statistically rigorous performance baseline without confounding interactions from exogenous variables. Future work will extend the framework to multivariate forecasting by incorporating economic indicators, policy signals, energy commodity prices, and environmental drivers. Within this extension, attention-based and transformer architectures represent a promising direction, particularly when combined with explainable AI techniques such as SHAP and LIME to preserve transparency and stakeholder trust.

In summary, the Hybrid Prophet–LSTM framework represents a reproducible, scalable, and operationally relevant solution for financial forecasting in sustainable energy grids. It advances the state of the art by explicitly bridging the long-standing trade-off between interpretability and predictive accuracy, offering both methodological novelty and direct industrial relevance.

## 6. CONCLUSION

This study has introduced and empirically validated a Hybrid Prophet–LSTM forecasting framework for

energy utility finance that combines statistical interpretability with nonlinear predictive capability. By decomposing financial time series into transparent structural components using Prophet and modelling residual volatility through LSTM, the proposed framework achieves both explainability and high forecasting accuracy. Empirical evaluation using 18 years of financial data from Tenaga Nasional Berhad (TNB) demonstrates that the hybrid model consistently outperforms ARIMA, standalone Prophet, and standalone LSTM benchmarks. In particular, it achieves the lowest forecasting errors (RMSE of 0.07 and MAPE of 1.29%) with narrow confidence intervals, indicating strong robustness and stability across different market conditions. Rather than serving as an investment trading signal, the forecasting framework is intended to support regulatory tariff modelling, revenue adequacy analysis, and long-term infrastructure planning in sustainable urban energy systems.

From a performance evaluation perspective, the framework shows reliable superiority across both scale-dependent (RMSE, MAE) and scale-independent (MAPE, SMAPE) metrics, confirming its ability to control absolute and proportional forecast errors simultaneously. This balanced performance highlights the effectiveness of residual-based learning in capturing nonlinear volatility while preserving the interpretability of long-term financial trends and seasonal patterns. Such characteristics are especially valuable in regulated energy sectors where transparent and defensible forecasting models are required.

Beyond predictive accuracy, the proposed framework offers clear practical relevance for policy and industry stakeholders. For regulators, the interpretable decomposition supports evidence-based tariff formulation and transparent communication with stakeholders. For utility operators and investors, the framework enables informed financial planning, scenario analysis, and risk assessment under market volatility, policy changes, and sustainability-driven transitions. Its modular and reproducible design further facilitates integration into decision-support systems and large-scale energy planning platforms.

Future work will extend the proposed framework toward multivariate forecasting by incorporating exogenous variables such as macroeconomic indicators, energy commodity prices, weather conditions, and policy signals. This progression will enable the framework to capture broader grid-impact dynamics while preserving the interpretability and benchmarking rigour established in the present univariate formulation. In addition, the integration of explainable AI (XAI) techniques, including SHAP and LIME, will further enhance transparency in the residual learning component and strengthen stakeholder trust. Deployment in real-time operational environments will be explored to maximise practical impact. Overall, the Hybrid Prophet–LSTM framework contributes a transparent, scalable, and statistically validated approach to AI-enabled forecasting for sustainable energy finance. By explicitly addressing the dual requirements of interpretability and predictive adaptability, the proposed framework closely aligns with the operational, regulatory, and decarbonisation objectives of modern smart grid infrastructure.

## AUTHOR CONTRIBUTIONS

**Tajul Rosli Razak:** Conceptualisation, methodology, formal analysis, software, investigation, writing—review and editing, supervision, project administration. **Ahmed El Alaoui:** Data curation, investigation, writing—original draft preparation, writing—review and editing. **Hasila Jarimi:** Conceptualisation, writing—review and editing. **Ahmad Zia Ul-Saufie:** Conceptualisation, writing—review and editing. **Mohammad Hafiz Ismail:** Writing—review and editing. **Mohd Faris Mohd Fuzi:** Software, writing—review and editing.

## COMPETING INTERESTS

The authors declare that they have no conflict of interest.

## DATA ACCESSIBILITY

The data supporting the findings of this study are available from the corresponding author upon reasonable request.

## FUNDING

The authors received no financial support for the research, authorship, and/or publication of this article.

## REFERENCES

Avinash, G., Ramasubramanian, V., Ray, M., Paul, R.K., Godara, S., Nayak, G.H., Kumar, R.R., Manjunatha, B., Dahiya, S. and Iquebal, M.A. (2024). Hidden Markov guided deep learning models for forecasting highly volatile agricultural commodity prices. *Applied Soft Computing*, 158, p. 111557. DOI:10.1016/j.asoc.2024.111557.

Box, G.E., Jenkins, G.M., Reinsel, G.C. and Ljung, G.M. (2015). *Time series analysis: forecasting and control*. Hoboken: John Wiley & Sons.

Cheng, J., Deng, H., Sun, G., Guo, P. and Zhang, J. (2020). Application of ARIMA model in financial time series in stocks. *Lecture Notes in Computer Science*, 12239, pp. 232–243.

Cohen, J., Kaur, S., Li, Y., Kolter, J.Z. and Talwalkar, A. (2021) Gradient descent on neural networks typically occurs at the edge of stability. *Proceedings of the 9th International Conference on Learning Representations (ICLR 2021)*.

Hamilton, J.D. (2010). Regime switching models. In: *Macroeconometrics and Time Series Analysis*, pp. 202–209.

He, C. (2023). A hybrid model based on multi-LSTM and ARIMA for time series forecasting. *Proceedings of the 8th International Conference on Intelligent Computing and Signal Processing (ICSP 2023)*, pp. 612–616.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), pp. 1735–1780.

Hsu, M.W., Lessmann, S., Sung, M.C., Ma, T. and Johnson, J.E. (2016). Bridging the divide in financial market forecasting: Machine learners vs. financial economists. *Expert Systems with Applications*, 61, pp. 215–234.

Lago, J., De Ridder, F. and De Schutter, B. (2018). Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms, *Applied Energy*, 221, pp. 386–405. DOI:10.1016/j.apenergy.2018.02.069.

Lim, B. and Zohren, S. (2021). Time-series forecasting with deep learning: A survey, *Philosophical Transactions of the Royal Society A*, 379(2194). DOI:10.1098/rsta.2020.0209.

Liu, H., Tian, H.Q. and Li, Y.F. (2012). Comparison of two new ARIMA-ANN and ARIMA-Kalman hybrid methods for wind speed prediction, *Applied Energy*, 98, pp. 415–424.

Liu, Q., Wang, Q. and Wang, X. (2025). RVFL-LSTM: A lightweight model with long-short term memory for time series, *Knowledge-Based Systems*, 309, 112896.

Makridakis, S., Spiliotis, E. and Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward, *PLOS ONE*, 13(3), e0194889. DOI:10.1371/journal.pone.0194889.

Mohamed, E., Sirlantzis, K. and Howells, G. (2022). A review of visualisation-as-explanation techniques for convolutional neural networks and their evaluation, *Displays*, 73, 102239.

Mohammadi, M., Jamshidi, S., Rezvanian, A., Gheisari, M. and Kumar, A. (2024). Advanced fusion of MTM-LSTM and MLP models for time series forecasting: An application for forecasting the solar radiation, *Measurement: Sensors*, 33, 101179.

Nystrup, P., Madsen, H. and Lindström, E. (2017). Long memory of financial time series and hidden Markov models with time-varying parameters, *Journal of Forecasting*, 36(8), pp. 989–1002.

Oukhouya, H., Kadiri, H., El Himdi, K. and Guerbaz, R. (2024). Forecasting international stock market trends: XGBoost, LSTM, LSTM-XGBoost, and backtesting XGBoost models, *Statistics, Optimization & Information Computing*, 12(1), pp. 200–209.

Pan, Y.A., Guo, J., Chen, Y., Cheng, Q., Li, W. and Liu, Y. (2024). A fundamental diagram based hybrid framework for traffic flow estimation and prediction by combining a Markovian model with deep learning, *Expert Systems with Applications*, 238, 122219.

Parviz, L. and Ghorbanpour, M. (2024). Assimilation of PSO and SVR into an improved ARIMA model for monthly precipitation forecasting, *Scientific Reports*, 14(1).

Pascanu, R., Mikolov, T. and Bengio, Y. (2013). On the difficulty of training recurrent neural networks, in *Proceedings of the 30th International Conference on Machine Learning (ICML 2013)*, pp. 2347–2355.

Peng, J., Deng, C. and Chen, Y.C. (2023). Application of ARIMA time series model in stock data prediction, *Smart Innovation, Systems and Technologies*, 341, pp. 223–232.

Peng, Z., Khan, F.U., Khan, F., Shaikh, P.A., Yonghong, D., Ullah, I. and Ullah, F. (2021). An application of hybrid models for weekly stock market index prediction: Empirical evidence from SAARC countries,

Complexity, 2021(1)

Rabiner, L.R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition, *Proceedings of the IEEE*, 77(2), pp. 257–286.

Rizvi, M.F. (2024). ARIMA model time series forecasting, *International Journal for Research in Applied Science and Engineering Technology*, 12(5), pp. 3782–3785.

Shen, Z., Wan, Q. and Leatham, D.J. (2021). Bitcoin return volatility forecasting: A comparative study between GARCH and RNN, *Journal of Risk and Financial Management*, 14(7).

Taylor, S.J. and Letham, B. (2018). Forecasting at scale, *The American Statistician*, 72(1), pp. 37–45.

Thomson, M.E., Pollock, A.C., Önkal, D. and Gönül, M.S. (2019). Combining forecasts: Performance and coherence, *International Journal of Forecasting*, 35(2), pp. 474–484.

Torres, J.F., Hadjout, D., Sebaa, A., Martínez-Álvarez, F. and Troncoso, A. (2021). Deep learning for time series forecasting: A survey, *Big Data*, 9(1), pp. 3–21.

Wang, X., Tang, H., Liang, X., Zeng, Y., Qin, H. and Hu, H. (2025). RH-SAVis: Visual analytics of spatiotemporal situation awareness for ride-hailing multi-source data, *Journal of Visualization*, pp. 1–18.

Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future, *International Journal of Forecasting*, 30(4), pp. 1030–1081.

Xu, K., Chen, L., Patenaude, J.M. and Wang, S. (2024). RHINE: A regime-switching model with nonlinear representation for discovering and forecasting regimes in financial markets, in *Proceedings*, pp. 526–534.

Zhang, P.G. (2003). Time series forecasting using a hybrid ARIMA and neural network model, *Neurocomputing*, 50, pp. 159–175.

Zheng, H., Wu, J., Song, R., Guo, L. and Xu, Z. (2024). Predicting financial enterprise stocks and economic data trends using machine learning time series analysis, *Applied and Computational Engineering*, 87(1), pp. 26–32.