A Comparative Evaluation of SARIMAX, LSTM, and Prophet Models for Cryptocurrency Price Trend Prediction

Drissia Ennagoura
Laboratory of Computer Science, Innovation and Artificial Intelligence, Faculty of Science Dhar El Mahraz, Sidi Mohamed Ben Abdellah University, Fez, Morocco | Smartilab Laboratory, Moroccan School of Engineering Sciences (EMSI), Rabat, Morocco
drissia.ennagoura@usmba.ac.ma (corresponding author)

Kamal El Kehal
Laboratory of Computer Science, Innovation and Artificial Intelligence, Faculty of Science, Dhar El Mahraz, Sidi Mohamed Ben Abdellah University, Fez, Morocco
kamal.elkehal@usmba.ac.ma

Abdelhamid Berdai
Laboratory of Geo-Environmental Analysis, Planning, and Sustainable Development (LAGEA-DD), Sidi Mohamed Ben Abdellah University, Fez, Morocco
abdelhamid.berdai@usmba.ac.ma

Safae Merzouk
Smartilab Laboratory, Moroccan School of Engineering Sciences (EMSI), Rabat, Morocco
s.merzouk@emsi.ma

Khalid El Fahssi
Laboratory of Computer Science, Innovation and Artificial Intelligence, Faculty of Science, Dhar El Mahraz, Sidi Mohamed Ben Abdellah University, Fez, Morocco
khalid.elfahssi@usmba.ac.ma

Mohamed El Mahjouby
Department of Computer Science, Laboratory of Information Systems and Software Engineering (SIGL), Abdelmalek Essaadi University, Larache, Morocco
m.elmahjouby@uae.ac.ma

Mohamed El Far
Laboratory of Computer Science, Innovation and Artificial Intelligence, Faculty of Science, Dhar El Mahraz, Sidi Mohamed Ben Abdellah University, Fez, Morocco
mohamed.elfar@usmba.ac.ma

Mohamed Taj Bennani
Laboratory of Computer Science, Innovation and Artificial Intelligence, Faculty of Science, Dhar El Mahraz, Sidi Mohamed Ben Abdellah University, Fez, Morocco
taj.bennani@usmba.ac.ma

Received: 15 October 2025 | Revised: 21 November 2025 and 19 December 2025 | Accepted: 21 December 2025

 DOI: https://doi.org/10.48084/etasr.15554

ABSTRACT

Cryptocurrency price prediction is challenging due to strong nonlinearity and high volatility. This paper comparatively evaluates three forecasting models for Ethereum (ETH): SARIMAX with exogenous technical indicators, Long Short-Term Memory (LSTM) networks, and Facebook Prophet. Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and Exponential Moving Average (EMA) are incorporated to enhance signal quality. Empirical results reveal clear trade-offs between predictive accuracy, profitability, and risk. SARIMAX achieves the highest directional accuracy (75.00%) with limited profitability, while LSTM yields the highest cumulative profit (23.84%) at the cost of higher drawdown. Prophet provides a balanced compromise between accuracy and risk. The study contributes by jointly evaluating statistical forecasting accuracy and trading-oriented performance metrics, offering practical insights into model suitability for different investor risk profiles.

Keywords-cryptocurrency; Ethereum (ETH); price prediction; SARIMAX; LSTM; Facebook Prophet; technical indicators; algorithmic trading

I. INTRODUCTION

Algorithmic trading relies heavily on accurate price trend forecasting, a task that is particularly challenging in cryptocurrency markets due to strong nonlinearity, high volatility, and frequent regime changes. Unlike traditional financial markets, cryptocurrency markets operate continuously, lack centralized regulation, and are strongly influenced by speculative behavior and rapid shifts in investor sentiment. These characteristics result in complex price dynamics with abrupt volatility bursts and nonlinear dependencies that significantly limit the effectiveness of classical forecasting techniques. Ethereum (ETH), one of the most actively traded cryptocurrencies, exhibits such complex dynamics, making it an ideal benchmark asset for evaluating advanced time-series forecasting models.

This study comparatively evaluates three widely adopted forecasting models: Seasonal Autoregressive Integrated Moving Average with Exogenous variables (SARIMAX), Long Short-Term Memory (LSTM) neural networks, and Facebook Prophet. SARIMAX is a classical statistical model capable of capturing trend and seasonality while incorporating external regressors, making it suitable for interpretable time-series modelling [1-4]. LSTM networks are specifically designed to learn nonlinear and long-term temporal dependencies and have demonstrated strong performance in financial and cryptocurrency forecasting tasks [5-8]. Prophet is an additive regression model known for its robustness to seasonality changes, missing data, and structural breakdowns, although its application in highly volatile cryptocurrency markets remains relatively limited [9, 10].

From a methodological perspective, existing cryptocurrency forecasting approaches can be broadly categorized into three groups: statistical models, deep learning-based models, and hybrid or additive approaches. Statistical models such as ARIMA and SARIMA remain attractive due to their simplicity, interpretability, and low computational cost, but their performance often deteriorates under strong nonlinear and non-stationary conditions. Deep learning models, particularly recurrent neural networks such as LSTM and GRU, have shown superior ability to capture nonlinear patterns and long-term dependencies in financial time series. More recently, hybrid models and additive frameworks—such as CNN-LSTM architectures and Prophet-based extensions—have been proposed to combine the strengths of different modelling paradigms and improve robustness under volatile market conditions [11-13]. Recent studies have reported promising results for deep learning and hybrid models in cryptocurrency forecasting, including LSTM-based architectures and CNN-LSTM hybrids [11]. Although these approaches often improve predictive accuracy, most existing works primarily focus on statistical error metrics such as the Root Mean Squared Error (RMSE) or Mean Absolute Error (MAE). However, a low prediction error does not necessarily translate into profitable or stable trading strategies. In practice, a forecasting model can achieve strong statistical accuracy while still generating economically unviable or highly risky trading signals.

To address this limitation, this study adopts a dual-evaluation framework that jointly assesses the accuracy of forecasting and trading-oriented performance. The three models are compared under identical experimental conditions using both statistical metrics and financial indicators, including signal accuracy, cumulative profit, profit factor, and maximum drawdown. This framework enables a realistic assessment of the trade-offs between accuracy, profitability, and risk, and provides actionable insights into model suitability for different algorithmic trading risk profiles.

Recent advances in attention-based and Transformer architectures have shown great potential in time-series forecasting across various domains, including finance, energy, and macroeconomics [14-16]. Despite their expressive power, such models typically require large datasets, extensive regularization, and careful hyperparameter tuning, which can limit their robustness when applied to relatively short and noisy cryptocurrency time series such as ETH. In this context, this study serves as a practical benchmark by comparing three parsimonious yet widely used forecasting models, against which more complex architectures and alternative data sources—such as sentiment indicators or on-chain metrics—can be evaluated in future research.

II. METHODS AND MATERIALS

A. Dataset and Feature Construction

Ethereum (ETH) historical price data were collected from the KuCoin cryptocurrency exchange using its official REST API and consist of daily OHLCV (Open, High, Low, Close, Volume) observations covering the period from May 29, 2022, to February 29, 2024, resulting in 2,629 data points. The daily frequency was selected, as it provides a suitable balance between noise reduction and the ability to capture medium-term market dynamics and is widely adopted in cryptocurrency forecasting studies. Higher-frequency data, although richer in information, are more sensitive to market microstructure noise and transaction costs, whereas lower-frequency data may fail to capture timely trading signals.

The Ethereum OHLCV data used in this study were retrieved programmatically from the official KuCoin REST API. While the exact processed dataset is not publicly hosted due to institutional constraints, all data are publicly accessible via the same API, allowing full reproducibility of the experiments.

Each observation includes open, high, low, and close prices, as well as trading volume. The closing price series was used as the primary target variable for forecasting and trading signal generation.

B. Feature Engineering

Based on the OHLCV data, three widely used technical indicators were computed: Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and Exponential Moving Average (EMA). These indicators were selected due to their complementary characteristics in capturing momentum, trend reversals, and longer-term price direction, while maintaining interpretability and computational efficiency.

RSI provides information on overbought and oversold market conditions, MACD captures momentum and trend-following behavior through moving-average differentials, and EMA emphasizes recent price movements, allowing the models to respond more quickly to market changes. These indicators served as exogenous variables for SARIMAX and as input features for the learning-based models.

Figure 1 illustrates the overall forecasting and trading evaluation workflow adopted in this study. Raw OHLCV data are first preprocessed, followed by feature construction through technical indicator computation. Forecasts generated by SARIMAX, LSTM, and Prophet are subsequently transformed into trading signals, which are evaluated using both statistical and trading-oriented performance metrics.

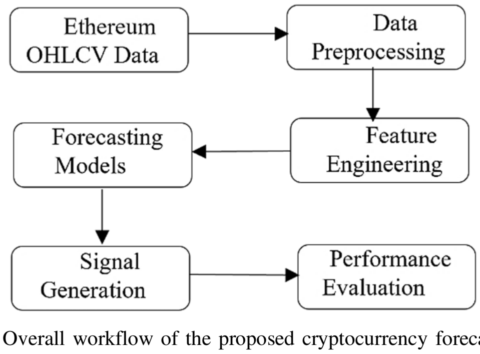

Fig. 1. Overall workflow of the proposed cryptocurrency forecasting and trading evaluation framework.

C. Data Preprocessing

Data preprocessing involved removing missing values and ensuring temporal consistency of the time series. Stationarity was assessed using the Augmented Dickey-Fuller (ADF) and KPSS tests. Since the ETH price series exhibited non-stationary behavior, first-order differencing was applied for models requiring stationarity assumptions, such as SARIMAX.

For learning-based models, all input features were normalized using Min-Max scaling to the range [0, 1], which improves training stability and accelerates convergence. The dataset was then split chronologically into training (90%) and testing (10%) subsets to preserve the temporal structure of the time series.

D. Forecasting Models

1) SARIMAX

The SARIMAX model was employed to capture linear dependencies, trend, and seasonality in the ETH price series. Technical indicators (RSI, MACD, and MACD signal line) were incorporated as exogenous regressors. The model parameters (p, d, q) and seasonal components (P, D, Q, s) were selected based on information criteria and diagnostic tests, with a weekly seasonality ($s = 7$). Forecasted prices were converted into buy and sell signals based on directional changes, and model performance was evaluated through back testing.

2) LSTM Model

LSTM networks were used to model nonlinear temporal dependencies in the ETH price series. The input features consisted of normalized RSI, MACD, and EMA indicators, which were transformed into fixed-length sequences to capture temporal relationships. The final architecture consisted of a single LSTM layer with 100 hidden units, followed by a dropout layer ($rate$ = 0.8) and a dense output layer. This configuration was selected after preliminary experiments with deeper, stacked, and attention-enhanced LSTM architectures, which consistently led to overfitting and did not improve forecasting or trading performance due to the relatively short and noisy nature of the ETH time series. Training was performed using the Adam optimizer and MSE loss for 20 epochs, with early stopping based on validation loss to prevent overfitting. Trading signals were generated from the sign of consecutive predicted price changes.

3) Prophet Model

Prophet is an additive regression model designed to capture trend and multiple seasonal components in time series data. In this study, the ETH closing prices were formatted according to Prophet's requirements, and RSI, MACD, and the MACD signal line were included as additional regressors. The model was trained on the training dataset, and future regressor values were approximated using their most recent observations. Directional trading signals were derived from changes in predicted prices and evaluated using the same backtesting framework as the other models to ensure methodological consistency.

E. Forecasting Models

For all models, the forecasted prices were transformed into buy and sell signals based on the direction of the predicted price movements. A long position was opened when an upward movement was predicted and closed when a downward signal was generated. Model performance was evaluated using both forecasting- and trading-oriented metrics. Forecasting accuracy was assessed through directional signal accuracy, while trading performance was evaluated using cumulative profit, profit factor, and Maximum Drawdown (MDD). This dual-evaluation framework enables a realistic comparison of the models in terms of accuracy, profitability, and risk under identical experimental conditions.

III. RESULTS AND DISCUSSION

This section presents a comparative evaluation of the forecasting and trading performance of the SARIMAX, LSTM, and Prophet models applied to Ethereum (ETH) price prediction. All models were trained on 90% of the dataset and evaluated on the remaining 10% under identical experimental conditions to ensure a fair and consistent comparison. Model performance is assessed using both directional accuracy and trading-oriented metrics, including cumulative profit, maximum drawdown, and profit factor.

SARIMAX achieved the highest directional signal accuracy (75.00%), reflecting strong stability and reliable trend-following behavior. However, this high accuracy was accompanied by relatively low profitability, with a total return of 3.02% and an MDD of –4.07%. These results highlight the inherent limitations of linear statistical models in capturing the highly nonlinear, volatile, and rapidly changing dynamics of cryptocurrency markets. Although SARIMAX effectively filters noisy signals and provides interpretable outputs, its conservative nature limits its ability to exploit large price movements during strong market trends.

In contrast, the LSTM model generated the highest cumulative profit (23.84%), demonstrating its superior ability to exploit nonlinear temporal dependencies in ETH price movements. This enhanced profitability came at the cost of increased risk exposure, as reflected by a higher maximum drawdown of –9.66%. Nevertheless, the profit factor of 2.36 confirms that total gains exceeded total losses, validating the practical viability of the LSTM-based trading strategy. These results suggest that deep learning models are more suitable for aggressive trading strategies that prioritize return maximization over strict risk control.

Figure 2 illustrates the real versus predicted ETH closing prices generated by the LSTM model, showing its ability to closely track price movements and capture major directional changes under volatile market conditions. The profitability of the LSTM-based strategy is further illustrated in Figure 3, which depicts the evolution of trading capital starting from an initial capital of 1000 USD. Despite experiencing intermittent drawdowns, the capital trajectory shows sustained growth over successive trades, confirming the model's strong profit-generating capability under controlled risk.

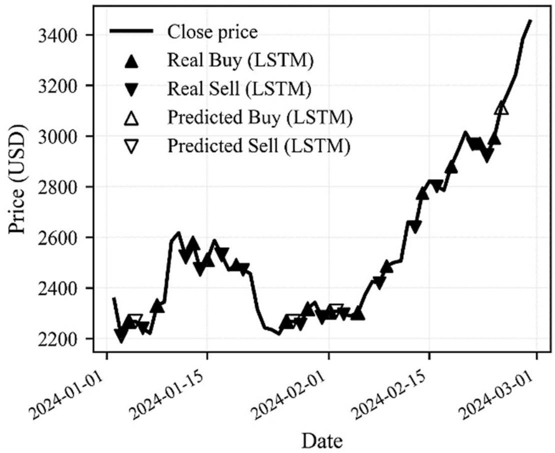

Fig. 2. Comparison of real and predicted trading signals generated by the LSTM model.

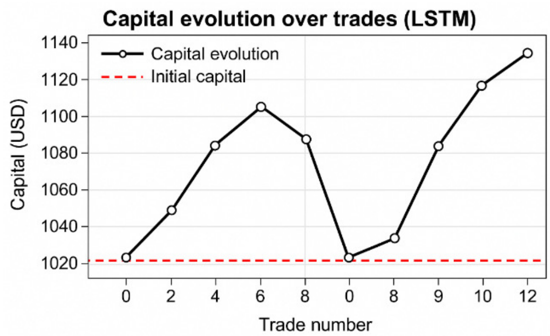

Fig. 3. Capital evolution over trades for the LSTM-based trading strategy with an initial capital of 1000 USD.

Prophet provided intermediate performance between SARIMAX and LSTM. It achieved a total profit of 17.64% with moderate risk exposure (MDD of −5.54%) and a profit factor of 5.24. While Prophet was able to capture medium-term trend reversals effectively, its additive structure resulted in reduced robustness during periods of increased volatility. Consequently, its profitability remained lower than that of LSTM, although higher than SARIMAX, positioning Prophet as a balanced alternative for traders seeking moderate risk–return profiles.

Table I summarizes the comparative performance of the three models. The results reveal clear trade-offs between predictive accuracy, profitability, and risk. SARIMAX is most suitable for stability-oriented strategies, LSTM favors aggressive profit-seeking strategies, and Prophet offers a balanced compromise between return and drawdown.

TABLE I. SUMMARY COMPARISON OF MODELS

Model	Signal Accuracy (%)	Total Profit (%)	Max Drawdown (%)	Profit Factor
SARIMAX	75.00	3.02	-4.07	6.04
LSTM	71.19	23.84	-9.66	2.36
Prophet	63.93	17.64	-5.54	5.24

A. *Practical Implications for Algorithmic Trading*

Beyond numerical performance metrics, the comparative results indicate that no single forecasting model consistently outperforms the others across all evaluation criteria. Instead, model selection should be guided by the trader's risk tolerance, investment horizon, and strategic objectives. Statistical models such as SARIMAX favor signal stability and risk control, whereas deep learning approaches like LSTM prioritize profitability at the expense of higher volatility. Prophet offers an intermediate solution by balancing return and drawdown, making it suitable for moderate risk-return trading strategies.

From an applied perspective, these findings emphasize the importance of aligning forecasting models with portfolio management constraints. Conservative investors may prefer SARIMAX due to its interpretability and lower drawdowns, while active traders seeking higher returns may favor LSTM, provided that adequate risk management mechanisms—such as position sizing or stop-loss rules—are implemented. Prophet can serve as a compromise solution in systematic trading frameworks that require both robustness and transparency.

B. *Practical Implications for Algorithmic Trading*

It should be noted that transaction costs were not included in the backtesting framework, as the primary objective was comparative model evaluation rather than deployment-level profitability. Although transaction fees and slippage would reduce absolute returns, they are unlikely to alter the relative ranking of the models due to similar trading frequencies.

To assess whether the observed performance differences are statistically significant, additional tests were conducted. McNemar's test indicated that SARIMAX differs significantly from both LSTM ($\chi^2 = 5.82$, $p = 0.0159$) and Prophet ($\chi^2 = 10.23$, $p = 0.0014$) in terms of directional accuracy, while no significant difference was found between LSTM and Prophet ($\chi^2 = 2.12$, $p = 0.1456$). Furthermore, the Diebold-Mariano test revealed highly significant differences across all model pairs ($p < 0.001$), with LSTM yielding the lowest forecast error, followed by Prophet and SARIMAX. These results confirm that the evaluated models exhibit statistically distinct performance characteristics, strengthening the robustness and credibility of the comparative conclusions.

IV. CONCLUSION

This study presented a comparative evaluation of three widely used time-series forecasting models—SARIMAX, LSTM, and Prophet—for Ethereum (ETH) price prediction using technical indicators such as RSI, MACD, and EMA. Model performance was assessed through a dual-evaluation framework that combined directional accuracy and trading-oriented metrics, including cumulative profit, maximum drawdown, and profit factor. The results highlight clear trade-offs between accuracy, profitability, and risk. SARIMAX achieved the highest signal accuracy (75.00%) with the lowest drawdown, making it suitable for stability-oriented strategies. LSTM delivered the highest profitability (23.84%) at the cost of higher risk exposure, while Prophet provided a balanced compromise between profit and drawdown, demonstrating effective medium-term trend detection.

This study is limited by its reliance on price-based technical indicators and its focus on a single asset (Ethereum). Future work will extend the analysis to additional cryptocurrencies and incorporate alternative data sources such as sentiment indicators and on-chain metrics. Moreover, hybrid and Transformer-based architectures will be explored within the same dual-evaluation framework to further enhance forecasting robustness and trading performance.

REFERENCES

[1] C. Chatfield, *Time-Series Forecasting*. Chapman and Hall/CRC, 2000.

[2] H. Li, Numerical Methods Using Kotlin: For Data Science, Analysis, and Engineering. Apress, 2023.

[3] N. Deretić, D. Stanimirović, M. A. Awadh, N. Vujanović, and A. Djukić, "SARIMA Modelling Approach for Forecasting of Traffic Accidents," *Sustainability*, vol. 14, no. 8, Apr. 2022, Art. no. 4403, https://doi.org/10.3390/su14084403.

[4] T. Falatouri, F. Darbanian, P. Brandtner, and C. Udokwu, "Predictive Analytics for Demand Forecasting – A Comparison of SARIMA and LSTM in Retail SCM," *Procedia Computer Science*, vol. 200, pp. 993–1003, 2022, https://doi.org/10.1016/j.procs.2022.01.298.

[5] S. Siami-Namini, N. Tavakoli, and A. S. Namin, "A Comparison of ARIMA and LSTM in Forecasting Time Series," in *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, Dec. 2018, pp. 1394–1401, https://doi.org/10.1109/ICMLA.2018.00227.

[6] T. Fischer and C. Krauss, "Deep learning with long short-term memory networks for financial market predictions," *European Journal of Operational Research*, vol. 270, no. 2, pp. 654–669, Oct. 2018, https://doi.org/10.1016/j.ejor.2017.11.054.

[7] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997, https://doi.org/10.1162/neco.1997.9.8.1735.

[8] G. P. Zhang, "Time series forecasting using a hybrid ARIMA and neural network model," *Neurocomputing*, vol. 50, pp. 159–175, Jan. 2003, https://doi.org/10.1016/S0925-2312(01)00702-0.

[9] S. J. Taylor and B. Letham, "Forecasting at Scale," *The American Statistician*, vol. 72, no. 1, pp. 37–45, Jan. 2018, https://doi.org/10.1080/00031305.2017.1380080.

[10] J. Sailaja, K. N. Bhargavi, G. L. N. Vanguri, N. Suryakala, and S. Thiruveedula, "Crypto Currency Price Prediction on Ethereum Using Time Series Forecasting Models Arima and Facebook Prophet Models," in *International Conference on Computational Innovations and Emerging Trends (ICCIET- 2024)*, July 2024, pp. 1074–1084, https://doi.org/10.2991/978-94-6463-471-6_102.

[11] O. M. Ahmed, L. M. Haji, A. M. Ahmed, and N. M. Salih, "Bitcoin Price Prediction using the Hybrid Convolutional Recurrent Model Architecture," *Engineering, Technology & Applied Science Research*, vol. 13, no. 5, pp. 11735–11738, Oct. 2023, https://doi.org/10.48084/etasr.6223.

[12] S. Smyl, "A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting," *International Journal of Forecasting*, vol. 36, no. 1, pp. 75–85, Jan. 2020, https://doi.org/10.1016/j.ijforecast.2019.03.017.

[13] S. Lahmiri and S. Bekiros, "Cryptocurrency forecasting with deep learning chaotic neural networks," *Chaos, Solitons & Fractals*, vol. 118, pp. 35–40, Jan. 2019, https://doi.org/10.1016/j.chaos.2018.11.014.

[14] G. Demosthenous, C. Georgiou, and E. Polydorou, "From On-chain to Macro: Assessing the Importance of Data Source Diversity in Cryptocurrency Market Forecasting." arXiv, 2025, https://doi.org/10.48550/ARXIV.2506.21246.

[15] M. L. de Prado, *Advances in Financial Machine Learning*. John Wiley & Sons, 2018.

[16] A. Farooq, M. I. Uddin, M. Adnan, A. A. Alarood, E. Alsolami, and S. Habibullah, "Interpretable multi-horizon time series forecasting of

cryptocurrencies by leverage temporal fusion transformer," *Heliyon*, vol. 10, no. 22, Nov. 2024, https://doi.org/10.1016/j.heliyon.2024.e40142.