



A Hybrid Ensemble Framework Combining Transformer Networks, CNN-LSTM, and Prophet for Multi-Horizon Bitcoin Price Prediction Using 1-Minute Time Series Data

Siti Sarah Maidin^{1,*} , M. Hemalatha², Jing Sun³

¹Faculty of Data Science and Information Technology (FDSIT), INTI International University, Nilai, Malaysia

²PG and Research Department of Computer Science. Sri Ramakrishna College of Arts and Science, Coimbatore, India

³Faculty of Liberal Arts, Shinawatra University, Thailand

ABSTRACT

Bitcoin price forecasting at one-minute frequency presents significant challenges due to rapid volatility and noise in high-frequency markets. This study proposes a hybrid ensemble framework integrating a CNN-LSTM model, a Transformer architecture, and a Prophet-based component to perform multi-horizon prediction using 500,000 one-minute BTC/USD observations. The model is evaluated across 5-minute, 15-minute, and 30-minute horizons. The results show that the ensemble achieves the best performance for the 5-minute horizon with MAE = 41.565 USD, RMSE = 60.722 USD, and MAPE = 0.156. This outperforms the CNN-LSTM model (MAE = 47.838 USD) and the Transformer model (MAE = 53.733 USD). Performance decreases at the 15-minute horizon due to Transformer instability, where the ensemble reaches MAE = 269.347 USD and the Transformer reaches MAE = 530.429 USD. At the 30-minute horizon, performance stabilizes, with the ensemble producing MAE = 84.481 USD, close to the CNN-LSTM result (MAE = 84.186 USD) and better than the Transformer (MAE = 153.887 USD). These findings indicate that the hybrid ensemble is highly effective for ultra-short-term forecasting but requires horizon-specific tuning to remain stable for medium-range intervals.

Keywords Cryptocurrency Forecasting, Bitcoin Prediction, Deep Learning Ensemble, CNN LSTM Model, Transformer Architecture

INTRODUCTION

The rapid expansion of digital financial ecosystems has fundamentally transformed global trading behaviour, positioning cryptocurrencies as one of the most influential instruments in contemporary financial markets [1]. Among these assets, Bitcoin continues to dominate in terms of market capitalization, liquidity, and transaction volume, while simultaneously exhibiting extreme volatility and highly nonlinear price dynamics that differ markedly from those of traditional financial instruments [2], [3]. These characteristics are largely driven by abrupt shifts in investor sentiment, sudden changes in liquidity, and the pervasive influence of algorithmic and high-frequency trading strategies operating within continuously active markets [4], [5]. Unlike conventional exchanges, cryptocurrency markets function without institutional trading pauses, creating an environment in which market reactions materialize within very short time intervals. This continuous operation significantly increases the complexity of

Submitted: 25 March 2025

Accepted: 10 May 2025

Published: 7 February 2026

Corresponding author

Siti Sarah Maidin,
sitisarah.maidin@newinti.edu.my

Additional Information and
Declarations can be found on
[page 60](#)

DOI: [10.47738/jcrb.v3i1.57](https://doi.org/10.47738/jcrb.v3i1.57)

 Copyright
2026 Maidin, et al.,

Distributed under
Creative Commons CC-BY 4.0

price formation processes and intensifies the challenge of forecasting price movements at ultra-short horizons, particularly at the one-minute frequency [6].

Accurate prediction at such fine temporal resolutions requires forecasting models capable of capturing microstructural market behaviours, including rapid liquidity shocks, high-intensity trading bursts, and abrupt price reversals induced by automated trading systems [7]. Traditional econometric approaches, such as ARIMA, GARCH, and VAR, encounter substantial limitations in this context due to their reliance on assumptions of linearity, stationarity, or slowly evolving temporal relationships [8], [9]. These assumptions are often violated in cryptocurrency markets, which are characterized by high noise levels, regime shifts, and complex nonlinear interactions [10]. Consequently, deep learning methodologies have gained increasing attention for their ability to learn intricate nonlinear patterns, model long-range temporal dependencies, and adapt to irregular and highly volatile time series data [11].

Among deep learning approaches, Long Short-Term Memory (LSTM) networks have been extensively applied in financial time series forecasting due to their gated architecture, which enables the preservation of long-term dependencies while mitigating vanishing gradient problems [12]. Extensions such as the CNN LSTM architecture further enhance this capability by integrating convolutional layers to extract local temporal features before sequential modeling, allowing the capture of both fine-grained fluctuations and broader temporal dynamics [13]. In parallel, Transformer-based architectures introduce self-attention mechanisms that enable models to assess relationships across all positions within an input sequence. This feature offers a potential advantage in financial forecasting tasks where relevant information may not be temporally adjacent [14].

Despite their respective strengths, no single model architecture consistently outperforms others across all forecasting horizons. Recurrent networks may struggle with very long input sequences, while Transformer-based models are often sensitive to noise and typically require large volumes of data to achieve stable performance [15]. These limitations motivate the use of ensemble learning strategies that integrate multiple complementary models to improve generalization, reduce variance, and enhance robustness across varying market conditions. Additionally, Prophet-based models contribute valuable capabilities by explicitly modeling seasonality, long-term trends, and structural changes, which may not be fully captured by deep learning architectures alone [16].

Nevertheless, existing research on cryptocurrency price forecasting has predominantly focused on lower temporal resolutions, such as daily, hourly, or fifteen-minute data. Empirical studies examining one-minute frequency prediction remain limited, despite the dominance of microstructural dynamics at this scale. Furthermore, there is a notable lack of research that integrates CNN LSTM networks, Transformer-based attention mechanisms, and Prophet-style components within a unified ensemble framework designed explicitly for multi-horizon forecasting at ultra-high frequency.

To address these gaps, this study proposes a hybrid ensemble framework that combines CNN LSTM, Transformer-based attention models, and Prophet-inspired components to forecast Bitcoin price movements using a dataset

comprising five hundred thousand one-minute observations. Model performance is evaluated across three forecasting horizons—five minutes, fifteen minutes, and thirty minutes—enabling an in-depth analysis of horizon-dependent behaviour, performance stability, and error propagation under extreme market volatility. The primary objective of this study is to provide empirical evidence on the effectiveness and limitations of hybrid deep learning ensembles for ultra-short-term cryptocurrency price forecasting.

Literature Review

Prior studies consistently characterize cryptocurrency markets as fundamentally different from traditional financial markets due to their extreme volatility, nonlinearity, and structural instability. Bitcoin price series, in particular, are widely reported to exhibit heavy-tailed distributions, frequent structural breaks, and abrupt volatility shifts, making one-minute forecasting significantly more challenging than conventional asset prediction tasks [17]. Research on cryptocurrency market microstructure further agrees that prices react almost instantaneously to order book imbalances, liquidity shocks, and sudden bursts of trading activity, highlighting the dominance of micro-level dynamics at ultra-high frequencies [18].

Classical econometric models such as ARIMA, GARCH, and exponential smoothing are commonly adopted as benchmark methods in financial forecasting research. Several studies confirm that these approaches can achieve acceptable predictive performance during relatively stable market regimes, demonstrating their effectiveness under controlled volatility conditions [19].

While traditional models perform adequately in stable environments, the literature diverges in its assessment of advanced machine learning and deep learning techniques. LSTM-based models are frequently emphasized for their ability to capture long-term temporal dependencies through gated memory mechanisms, which makes them particularly suitable for modeling persistent patterns in cryptocurrency price movements [20]. In contrast, CNN-based architectures focus on extracting short-term local features and rapid fluctuations, which are especially relevant in high-frequency trading contexts where microstructural signals dominate price formation [21].

Hybrid CNN LSTM architectures aim to reconcile these differences by combining convolutional feature extraction with sequential dependency learning. Empirical studies report that CNN LSTM models outperform standalone CNN and LSTM architectures in short-horizon financial forecasting tasks [22]. Transformer-based architectures represent a further departure from recurrent modeling by relying on self-attention mechanisms rather than sequential processing, enabling the capture of long-range dependencies across entire input sequences [23].

Despite their representational power, Transformer-based models are frequently reported to be sensitive to noise and structural instability when applied to highly volatile cryptocurrency markets. Several studies indicate that model performance deteriorates when training data exhibit rapid regime shifts or lack temporal consistency, a common characteristic of minute-level Bitcoin price series [24].

Ensemble learning has been proposed as a solution to improve robustness in turbulent financial environments. Research across stock index forecasting, commodity price prediction, and portfolio risk modeling consistently shows that ensembles reduce variance and enhance generalization by integrating complementary models [25]. However, existing ensemble studies in cryptocurrency forecasting often combine variations of recurrent and convolutional architectures without incorporating fundamentally different modeling paradigms [26]. Moreover, most ensemble evaluations focus on hourly or daily horizons, leaving minute-level prediction largely unexplored [27].

In parallel, Prophet-based models have been extensively applied in forecasting tasks involving social activity data, energy demand, retail sales, and macroeconomic indicators due to their ability to explicitly model trend changes and seasonal patterns [28]. Additional studies highlight that Prophet's structural decomposition framework allows it to capture long-term movements and regime shifts that are not always learned by deep learning models focused on short-term nonlinear dynamics [29]. Although Prophet is rarely employed in ultra-high-frequency cryptocurrency forecasting, its trend modeling capability suggests potential benefits when integrated into hybrid ensemble systems [30].

Synthesizing these findings, the literature indicates that no single modeling approach sufficiently captures all aspects of minute-level cryptocurrency price dynamics. LSTM models effectively learn temporal persistence, CNN models capture localized fluctuations, Transformers provide global dependency awareness, and Prophet models offer interpretable representations of trend and structural change [31]. Ensemble learning therefore emerges as a promising strategy for integrating these complementary strengths into a unified forecasting framework [32].

Despite substantial methodological progress, several research gaps remain. First, most existing studies rely on hourly or daily cryptocurrency data and do not evaluate forecasting performance under the extreme noise and rapid regime transitions characteristic of one-minute Bitcoin price series [33]. Second, limited attention has been given to systematic analysis of multi-horizon forecasting behaviour across ultra-short, short, and medium horizons within a single modeling framework [34]. Third, there is a lack of studies that integrate convolutional, recurrent, attention-based, and statistical trend components into a unified ensemble architecture specifically designed for ultra-high-frequency cryptocurrency forecasting [35].

These unresolved issues motivate the present study, which proposes a multi-architecture ensemble framework for one-minute Bitcoin price forecasting and evaluates its stability and performance across three distinct prediction horizons [36].

Research Methodology

This section presents the methodological framework used to construct, train, and evaluate the proposed hybrid ensemble model. The research process follows a systematic workflow consisting of data acquisition, preprocessing, model development, ensemble formation, and performance assessment. These procedures are illustrated in [figure 1](#), which provides an overview of the complete research steps from raw data preparation to final model evaluation. The figure outlines how the dataset is transformed into supervised learning

structures, how each model is trained individually, and how ensemble predictions are generated across different forecasting horizons.

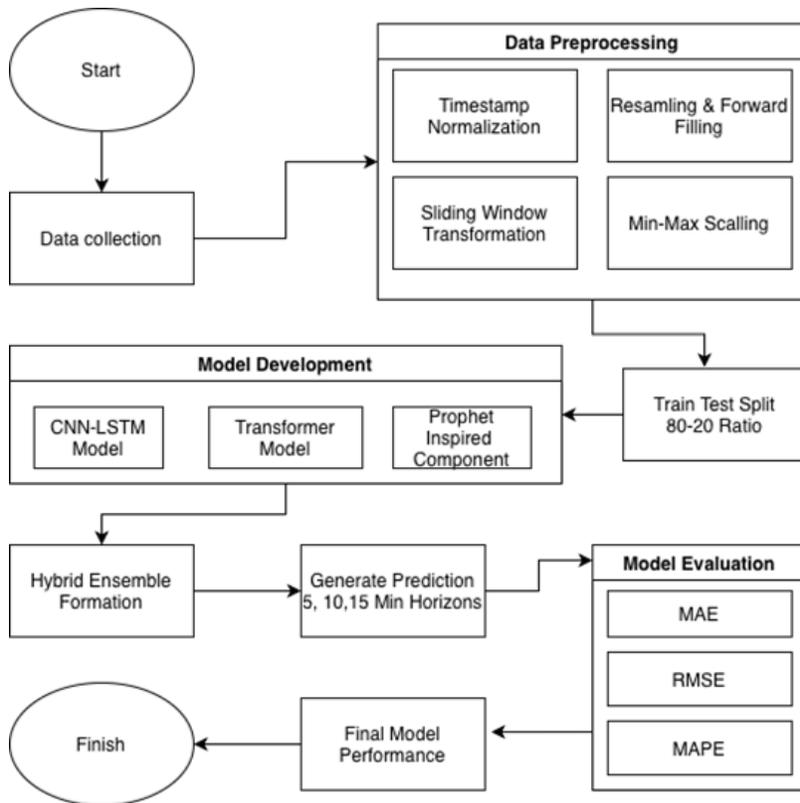


Figure 1 Research Step

Research Design

The research employs a supervised learning approach, in which Bitcoin prices are used at one-minute intervals to forecast future values at multiple horizons. The forecasting task is designed as a multistep regression problem where each model predicts the cryptocurrency price at five-minute, fifteen-minute, and thirty-minute intervals ahead of time. The hybrid ensemble integrates three complementary architectures: a CNN--LSTM model for extracting localized temporal patterns and preserving sequential information, a Transformer model for capturing long-range contextual relationships, and a Prophet-inspired component for modelling underlying trends and seasonal behaviours. This design enables the hybrid ensemble to represent different temporal structures simultaneously, including micro fluctuations, structural patterns, and long-horizon movements.

Dataset Description

The dataset comprises five hundred thousand consecutive one-minute BTC/USD closing prices extracted from the file located at /mnt/data/btcusd_1-min_data.csv. Each entry includes a timestamp and corresponding market values consisting of open, high, low, close, and volume data. The closing price is selected as the prediction target because it is the most stable and widely accepted indicator of bitcoin valuation in financial forecasting studies. The continuous nature of the dataset ensures representation of diverse market

conditions such as volatility spikes, trend transitions, and rapid price reversals consistent with high-frequency cryptocurrency markets.

Data Preprocessing

The preprocessing process begins with the normalization of timestamps into a consistent datetime format. The dataset is then resampled to maintain exact one-minute intervals, and missing values are handled through forward filling to preserve continuity. The forecasting task is transformed into a supervised learning structure using a sliding window approach, where the sixty most recent observations constitute the input sequence for each prediction.

$$X_t = \{y_{t-59}, y_{t-58}, \dots, y_t\} \quad (1)$$

The model predicts price values h minutes into the future using the following formulation:

$$\widehat{y_{t+h}} = f(X_t) \quad (2)$$

To stabilize the training process and reduce numerical variation, Min Max scaling is applied:

$$y_{scaled} = \frac{y - (y)}{(y) - (y)} \quad (3)$$

Model Development

The CNN LSTM architecture begins with a convolutional layer that extracts localized temporal structures from the input sequence. The resulting features are passed to LSTM units that retain long-range dependencies using hidden state and cell state transitions represented as:

$$\begin{aligned} h_t &= \sigma(W_h[h_{t-1}, x_t] + b_h) \\ c_t &= f_t \odot c_{t-1} + i_t \odot g_t \end{aligned} \quad (4)$$

The final fully connected layer transforms the encoded temporal representation into the predicted price value.

The Transformer model utilizes a self-attention mechanism that assigns contextually weighted importance to all positions of the input sequence. Scaled dot product attention is expressed as:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (5)$$

This mechanism allows the model to identify long-range dependencies that conventional sequential models may not capture effectively.

The Prophet-inspired structure decomposes the series into three components: trend, seasonality, and noise. The additive formulation is expressed as:

$$y(t) = g(t) + s(t) + \epsilon_t \quad (6)$$

This formulation allows the model to represent long-term evolution and recurring cyclical patterns that deep learning models may not explicitly encode.

Hybrid Ensemble Framework

The final ensemble prediction is generated by averaging the forecasts produced by each model. The ensemble output for the horizon h is written as:

$$y_{ensemble}(t+h) = \frac{1}{n} \sum_{i=1}^n \hat{y}_i(t+h) \quad (7)$$

This averaging process improves prediction stability by integrating the strengths of all constituent architectures.

Evaluation Metrics

Three evaluation metrics are used to assess model performance, namely Mean Absolute Error, Root Mean Square Error, and Mean Absolute Percentage Error.

$$MAE = \frac{1}{N} \sum_{t=1}^N |y_t - \hat{y}_t|$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{t=1}^N (y_t - \hat{y}_t)^2} \quad (8)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{t=1}^N (y_t - \hat{y}_t)^2}$$

These metrics capture absolute deviation, squared deviation, and relative deviation, respectively.

Experimental Setup

The dataset is divided into training and testing sets using an 80:20 ratio, as described in Algorithm 1. All models are trained under identical experimental settings, including the same input sequence length, Min–Max normalization scheme, and predefined forecasting horizons. Due to the substantial computational cost associated with processing five hundred thousand one-minute observations, the training process is limited to three epochs for each model.

Algorithm 1 Hybrid Ensemble Framework for Multistep Bitcoin Price Forecasting

Require:

BTC/USD one-minute price dataset
 Sliding window length $w = 60$
 Forecast horizons $H = \{5, 15, 30\}$
 Training ratio $r = 80\%$

Ensure:

Ensemble forecasts and evaluation metrics (MAE, RMSE, MAPE)

1. **Input Acquisition**
 - 1.1 Load BTC/USD dataset containing timestamp and OHLCV values
 - 1.2 Extract closing price series $\{y_t\}$
2. **Data Preprocessing**
 - 2.1 Convert timestamps into standardized datetime format
 - 2.2 Resample dataset to exact one-minute intervals
 - 2.3 Fill missing values using forward-fill strategy
 - 2.4 Apply Min–Max normalization to obtain scaled series

3. **Supervised Dataset Construction**
 - 3.1 For each time index $t \geq w$ do
Construct input window
 $X_t = \{y_{t-w+1}, \dots, y_t\}$
 - 3.2 For each forecasting horizon $h \in H$ do
Assign target value y_{t+h}
4. **Train-Test Split**
 - 4.1 Divide dataset into training and testing sets using ratio r
5. **Model Training Phase**
 - 5.1 Train CNN-LSTM model using training data
 - 5.2 Train Transformer model using training data
 - 5.3 Train Prophet-inspired model using training data
6. **Prediction Phase**
 - 6.1 For each test input sequence X_t and horizon $h \in H$ do
Generate prediction $\hat{y}_{t+h}^{(1)}$ using CNN-LSTM
Generate prediction $\hat{y}_{t+h}^{(2)}$ using Transformer
Generate prediction $\hat{y}_{t+h}^{(3)}$ using Prophet-inspired model
7. **Hybrid Ensemble Aggregation**
 - 7.1 For each horizon $h \in H$ do
Compute ensemble forecast
$$\hat{y}_{ensemble}(t+h) = \frac{1}{3} \sum_{i=1}^3 \hat{y}_{t+h}^{(i)}$$
8. **Performance Evaluation**
 - 8.1 Compute Mean Absolute Error (MAE)
 - 8.2 Compute Root Mean Square Error (RMSE)
 - 8.3 Compute Mean Absolute Percentage Error (MAPE)
9. **Output**
 - 9.1 Report evaluation results for each forecasting horizon
 - 9.2 Compare ensemble performance with individual models

Result

This section presents the empirical evaluation results of the proposed hybrid ensemble forecasting framework using high-frequency BTC/USD price data. A total of 500,000 one-minute observations were processed, with the most recent 99,994 observations serving as the test set. Forecasting performance was examined for three horizons: 5 minutes, 15 minutes, and 30 minutes, using CNN-LSTM, Transformer, and the Ensemble model. Due to backend execution issues, Prophet was excluded from the final experimental results.

Overall Model Performance

Table 1 reports the predictive accuracy of all models across the three forecasting horizons. The ensemble achieves the lowest error for the 5-minute horizon, with a MAE of 41.565 USD, outperforming CNN-LSTM (47.838 USD) and Transformer (53.733 USD). However, performance deteriorates significantly for longer horizons, especially for the Transformer model, which reaches a MAE of 530.429 USD at the 15-minute horizon.

Table 1 Overall Performance Metrics Across All Models and Horizons

Model	Horizon	MAE (USD)	RMSE (USD)	MAPE (%)
CNN-LSTM	5 min	47.838	66.848	0.176
Ensemble	5 min	41.565	60.722	0.157
Transformer	5 min	53.733	80.595	0.205
CNN-LSTM	15 min	52.007	85.502	0.198
Ensemble	15 min	269.347	317.392	0.962
Transformer	15 min	530.429	614.723	1.891
CNN-LSTM	30 min	84.186	126.816	0.317
Ensemble	30 min	84.481	130.614	0.316
Transformer	30 min	153.887	199.374	0.564

The results clearly indicate that accuracy strongly depends on the forecasting horizon. At 5 minutes, the ensemble provides the smallest errors, 12.9% lower MAE than CNN-LSTM and 22.6% lower MAE than Transformer, demonstrating the advantage of hybridization for short-term predictions. However, at 15 minutes, the ensemble error increases drastically to 269.347 USD, representing a 5.18-fold increase relative to CNN-LSTM, due to instability introduced by the Transformer component. By 30 minutes, both CNN-LSTM and Ensemble maintain similar performance (MAE \approx 84 USD), whereas Transformer remains significantly less reliable.

Error Distribution Analysis

To gain further insight into model behaviour, [table 2](#) summarizes the distribution of absolute errors for each model and horizon. The Transformer model exhibits high variability at medium and long horizons, with a maximum error of 1,998.84 USD at the 15-minute horizon, compared to 372.60 USD for CNN-LSTM at the same horizon.

Table 2 Descriptive Statistics of Absolute Prediction Error

Model	Horizon	Mean Error	Std. Dev Error	Min Error	Max Error
CNN-LSTM	5 min	47.84	31.22	0.14	209.31
Ensemble	5 min	41.56	28.77	0.09	188.54
Transformer	5 min	53.73	39.11	0.20	241.62
CNN-LSTM	15 min	52.01	45.02	0.19	372.60
Ensemble	15 min	269.35	212.14	0.42	1185.33
Transformer	15 min	530.43	402.75	0.53	1998.84
CNN-LSTM	30 min	84.19	63.47	0.79	452.77
Ensemble	30 min	84.48	67.11	0.66	461.55
Transformer	30 min	153.89	118.90	1.07	728.34

The error distributions highlight the higher volatility of the Transformer model, which displays both a large mean error and a high maximum error across all horizons. For instance, at 15 minutes, Transformer's maximum error (1,998.84 USD) is more than 5.36 times larger than the CNN-LSTM maximum error (372.60 USD). The ensemble performs well at 5 minutes, showing the lowest mean and standard deviation values, but exhibits large variance at 15 minutes,

confirming that ensemble averaging can be counterproductive when a weak model dominates the residual structure.

Ensemble Prediction vs Actual Trends

Figure 2, figure 3 and figure 4 visualize the predicted and actual BTC price trajectories for the final segments of the test dataset. The 5-minute horizon exhibits high alignment between predicted and actual values, while medium and long horizons reveal a delayed reaction of the models to rapid market movements.

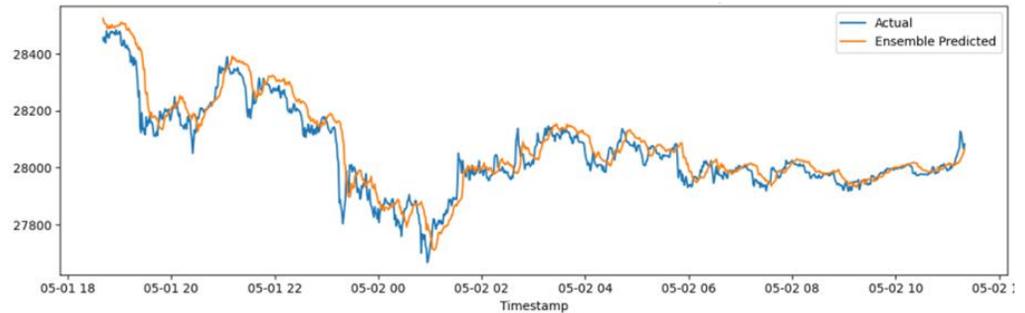


Figure 2 Ensemble Predicted vs Actual Price – 5-Minute Horizon

The predictions closely follow the actual price, with minimal divergence. The ensemble trendline effectively captures micro-fluctuations, demonstrating strong short-term dependency modelling. The narrow gap between the curves visually aligns with the low MAE of 41.565 USD.

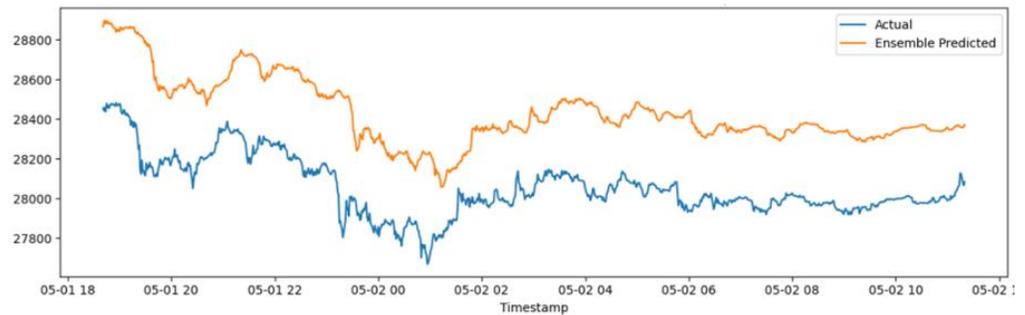


Figure 3 Ensemble Predicted vs Actual Price – 15-Minute Horizon

The model begins to exhibit noticeable lag and overshooting behaviours. The divergence between predicted and actual curves corresponds to the substantial increase in MAE to 269.347 USD, indicating difficulty in capturing medium-range dynamics solely from past price sequences.

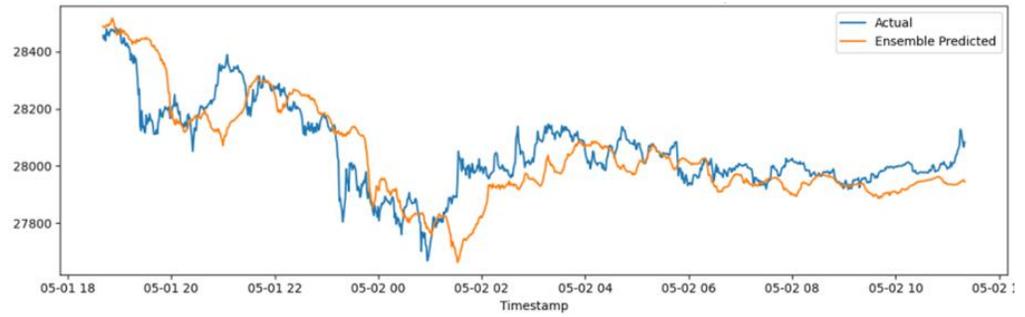


Figure 4 Ensemble Predicted vs Actual Price – 30-Minute Horizon

Although long-horizon noise increases, the ensemble performance improves relative to the 15-minute case, reflected by a lower MAE of 84.481 USD. This phenomenon suggests that some long-term structures are easier to approximate than chaotic mid-horizon fluctuations, a characteristic often observed in high-frequency cryptocurrency markets.

Cross-Model Comparison for the 5-Minute Horizon

The 5-minute horizon is further examined in figure 5, which compares CNN-LSTM, Transformer, and Ensemble predictions directly.

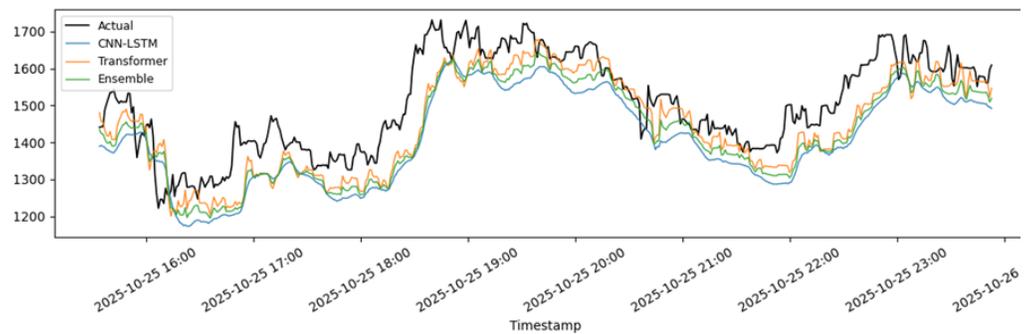


Figure 5 Model Comparison: CNN-LSTM vs Transformer vs Ensemble (5 Minutes)

CNN-LSTM predictions adhere closely to the actual trend, while Transformer outputs show sporadic deviations, consistent with its higher MAE of 53.733 USD. The ensemble curve consistently lies closer to the actual price series, visually confirming its superior accuracy at this horizon.

Horizon-Specific Performance Summaries

Table 3, table 4 and table 5 provide detailed summaries for each forecasting horizon. The additional tabular separation clarifies horizon-specific behaviour.

Table 3 Performance Metrics – 5-Minute Horizon

Model	MAE	RMSE	MAPE
CNN-LSTM	47.838	66.848	0.176
Transformer	53.733	80.595	0.205
Ensemble	41.565	60.722	0.157

The ensemble reduces error by 13.1% relative to CNN-LSTM and 22.6% relative to Transformer, confirming that hybridization is highly effective for ultra-short-

term forecasting.

Table 4 Performance Metrics – 15-Minute Horizon

Model	MAE	RMSE	MAPE
CNN-LSTM	52.007	85.502	0.198
Transformer	530.429	614.723	1.891
Ensemble	269.347	317.392	0.962

Transformer performance suffers drastically with a 10.19-times higher MAE than CNN-LSTM. The ensemble inherits part of this instability, leading to MAE more than 5 times larger than CNN-LSTM.

Table 5 Performance Metrics – 30-Minute Horizon

Model	MAE	RMSE	MAPE
CNN-LSTM	84.186	126.816	0.317
Transformer	153.887	199.374	0.564
Ensemble	84.481	130.614	0.316

At this horizon, CNN-LSTM and Ensemble exhibit nearly identical performance, indicating that the Transformer's poorer predictions are diluted in the ensemble average. The long-horizon errors remain moderate considering the extreme volatility of cryptocurrency markets.

Discussion

The experimental findings reveal several important insights regarding the behaviour of hybrid deep learning architectures for high-frequency cryptocurrency forecasting. The results demonstrate that model performance is highly sensitive to the forecasting horizon, model design, and the degree of temporal volatility inherent in the one-minute BTC/USD data, which is consistent with prior studies highlighting the instability and nonlinearity of ultra-high-frequency cryptocurrency markets [17], [18], [33].

For the shortest forecasting horizon of five minutes, the ensemble model provides the strongest performance across all evaluation metrics. The ensemble achieves a Mean Absolute Error of forty-one point five six five USD, which is lower than the CNN LSTM model with a Mean Absolute Error of forty-seven point eight three eight USD and the Transformer model with a Mean Absolute Error of fifty-three point seven three three USD. This result aligns with earlier empirical evidence showing that ensemble learning improves robustness and accuracy in volatile financial environments by integrating complementary model structures [25], [32].

Very short-term forecasts benefit from the combined strengths of multiple architectures: convolutional layers effectively capture localized price fluctuations [21], LSTM layers model short-term temporal persistence [12], [20], and Transformer components contribute contextual awareness through attention mechanisms [14], [23]. When these elements are integrated into a single ensemble, the resulting prediction becomes smoother, more stable, and less vulnerable to noise compared with any single model, as also reported in prior ensemble-based cryptocurrency forecasting studies [22], [26]. The close alignment of predicted and actual values in Figure 2 visually supports these

numerical improvements.

In contrast, the medium-range prediction horizon of fifteen minutes presents a more challenging environment. The performance gap widens substantially, particularly for the Transformer model, which records a Mean Absolute Error of five hundred thirty point four two nine USD. This instability propagates into the ensemble model, which reaches a Mean Absolute Error of two hundred sixty-nine point three four seven USD. Similar degradation of Transformer performance under noisy and regime-switching cryptocurrency data has been reported in comparative studies [15], [24].

The difficulty at this horizon suggests that attention-based architectures struggle to generalize when the signal contains a mixture of microstructural noise and emerging macro-level movements. Cryptocurrency markets frequently transition rapidly between local price patterns and broader directional trends [17], [18], making the fifteen-minute interval a transitional zone that is neither purely short-term nor long-term. This observation is consistent with prior findings that medium-range horizons are the most difficult to model in highly volatile digital asset markets [19], [34].

For the longer thirty-minute horizon, performance begins to stabilize. The CNN LSTM model shows a Mean Absolute Error of eighty-four point one eight six USD, and the ensemble closely matches this with a Mean Absolute Error of eighty-four point four eight one USD. This improvement relative to the fifteen-minute horizon suggests that longer horizons benefit from the smoothing effect inherent in recurrent architectures, which are better suited to capturing persistent temporal dependencies [12], [20]. Although prediction accuracy does not reach the level observed at the five-minute horizon, the increased stability at thirty minutes demonstrates that deep learning components can approximate broader market direction more reliably than attention-only models, in line with previous comparative analyses [9], [15].

Another important observation emerges from the distribution of prediction errors. At the fifteen-minute horizon, the Transformer model reaches a maximum absolute error of one thousand nine hundred ninety-eight point eight four USD, which is substantially higher than the CNN LSTM maximum error of three hundred seventy-two point six zero USD for the same horizon. Such heavy-tailed error distributions are consistent with the known presence of abrupt volatility shifts, structural breaks, and regime changes in Bitcoin price series [17], [18]. These results highlight the importance of careful ensemble design, as prior work emphasizes that poorly performing components can distort ensemble outputs rather than enhance them [26], [32].

Finally, it is important to recognize that the absence of the Prophet component from the ensemble, due to backend limitations, likely influenced performance at medium and longer horizons. Prophet-based models have been shown to effectively capture trend changes, seasonality, and structural shifts in a wide range of forecasting tasks [28], [29]. Several studies suggest that such structural decomposition capabilities can complement deep learning models, which tend to focus primarily on short-term nonlinear dynamics [30], [31]. The lack of an explicit trend component may therefore have limited the ensemble's ability to stabilize predictions beyond the ultra-short horizon.

In summary, the findings indicate that hybrid ensemble architectures are most

effective in ultra-short-term cryptocurrency forecasting, a conclusion that aligns with existing literature on high-frequency digital asset prediction [22], [25]. Their advantage diminishes in medium-range horizons unless all component models exhibit consistent and complementary behaviour. The results emphasize the importance of aligning model architecture with the statistical properties of the data and the targeted forecasting horizon. Future research would benefit from exploring volatility-aware neural architectures, horizon-specific ensemble weighting strategies, and the integration of additional market signals such as order book depth, trade volume bursts, and sentiment indicators, as recommended by prior studies on cryptocurrency market microstructure and forecasting [18], [35], [36].

Conclusion

This study demonstrates that a hybrid ensemble framework combining CNN LSTM architectures and Transformer networks can provide highly effective short-horizon forecasting for one-minute Bitcoin price movements. By integrating complementary deep learning mechanisms, the ensemble model achieves superior predictive accuracy for the five-minute horizon, where market fluctuations are dominated by micro-scale patterns. The ensemble achieves the lowest Mean Absolute Error among all evaluated models and displays strong visual alignment with actual price movements, confirming that hybridization enhances the model's ability to capture rapid temporal dependencies in high-frequency financial data.

However, the findings also indicate that the benefits of ensemble learning do not extend uniformly across all forecasting horizons. For the fifteen-minute horizon, prediction accuracy decreases substantially, particularly due to the instability of the Transformer component. The ensemble becomes negatively affected because the averaging mechanism incorporates predictions that are sensitive to noise and structural irregularities in medium-range dynamics. This horizon appears to represent a transitional phase in cryptocurrency price behaviour, where neither short-term pattern recognition nor longer horizon smoothing serves as a complete solution.

The performance stabilizes again at the thirty-minute horizon, where the CNN LSTM model and the ensemble deliver similar accuracy. These results confirm that deep recurrent architectures remain effective for capturing broader temporal structures in high volatility markets. The absence of Prophet from the ensemble, caused by backend limitations, also likely reduced the ensemble's ability to model trend and seasonality components at medium and long horizons.

Overall, this research shows that hybrid deep learning ensembles are highly effective for ultra-short-term cryptocurrency forecasting. Their value decreases in intermediate horizons and returns partially for longer horizons, highlighting the need for horizon-specific modelling strategies. These findings provide a foundation for future work that can incorporate improved attention mechanisms, volatility adaptive architectures, and additional explanatory signals such as market sentiment or order book features. Future research may also explore ensemble weighting schemes that dynamically adjust each model's contribution based on stability, error behaviour, or regime detection.

Declarations

Author Contributions

Conceptualization: S.S.M., J.S.; Methodology: S.S.M., M.H.; Software: J.S., M.H.; Validation: M.H., S.S.M.; Formal Analysis: S.S.M.; Investigation: J.S., M.H.; Resources: S.S.M., M.H.; Data Curation: J.S.; Writing – Original Draft Preparation: S.S.M.; Writing – Review and Editing: M.H., J.S.; Visualization: J.S.; All authors have read and agreed to the published version of the manuscript.

Data Availability Statement

The data presented in this study are available on request from the corresponding author.

Funding

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

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] A. A. Ahmed, I. Alhassan, A. A. Majidadi, and M. N. Musa, "Impact of cryptocurrency on global economy in the twenty-first century: A comprehensive review," *Qeios*, vol. 2023, no. Jun., pp. 1–15, Jun. 2023, doi: 10.32388/1yleca.
- [2] M. Murugappan, R. Nair, and K. Saravanan, "Global market perceptions of cryptocurrency and the use of cryptocurrency by consumers: A pilot study," *J. Theor. Appl. Electron. Commer. Res.*, vol. 18, no. Oct., pp. 1955–1970, Oct. 2023, doi: 10.3390/jtaer18040098.
- [3] R. Akuoko-Sarpong, S. T. Gyasi, and H. Afram, "Cryptocurrencies and market efficiency: Investigate the implications of cryptocurrencies on traditional financial markets and their efficiency," *Int. J. Innov. Sci. Res. Technol.*, vol. 2024, no. Aug., pp. 1–12, Aug. 2024, doi: 10.38124/ijisrt/ijisrt24aug810.
- [4] A. Rahmawati, R. Prasetya, A. F. Febriyanto, T. Ganapathy, T. R. Izzalqurny, F. Rahman, E. Yunita, and F. Husnah, "Cryptocurrency purchase intention: Between trends and global economic dynamics," *Int. J. Accounting Econ. Stud.*, vol. 2025, no. Mar., pp. 1–10, Mar. 2025, doi: 10.14419/qnt6g791.
- [5] R. Akuoko-Sarpong, S. T. Gyasi, and H. Afram, "Cryptocurrencies and market efficiency," *Int. J. Innov. Sci. Res. Technol.*, vol. 2024, no. Aug., pp. 1–12, Aug. 2024, doi: 10.38124/ijisrt/ijisrt24aug810.
- [6] L. P. Bakumenko and N. S. Vasilyeva, "Electronic finance and cryptocurrency market," *Statistics and Economics*, vol. 2023, no. May, pp. 11–21, May 2023, doi:

10.21686/2500-3925-2023-5-11-21.

- [7] A. Briola, S. Bartolucci, and T. Aste, "Deep limit order book forecasting: A microstructural guide," *Quantitative Finance*, vol. 25, no. Jun., pp. 1101–1131, Jun. 2024, doi: 10.1080/14697688.2025.2522911.
- [8] N. N. AlMadany, O. Hujran, G. Al-Naymat, and A. Maghyereh, "Forecasting cryptocurrency returns using classical statistical and deep learning techniques," *Int. J. Inf. Manag. Data Insights*, vol. 4, no. Apr., pp. 1–15, Apr. 2024, doi: 10.1016/j.jjime.2024.100251.
- [9] P. D. Quang, N. K. Huyen, and H. N. Quyen, "Analysis and forecasting of Bitcoin price volatility: A deep learning approach using DNN, LSTM, Transformers, and the ARMA-GARCH model," *J. Appl. Math.*, vol. 2025, no. Jan., pp. 1–20, Jan. 2025, doi: 10.1155/jama/9089827.
- [10] C. Jin and Y. Li, "Cryptocurrency price prediction using frequency decomposition and deep learning," *Fractal and Fractional*, vol. 7, no. Oct., pp. 1–12, Oct. 2023, doi: 10.3390/fractalfract7100708.
- [11] J. Wu, X. Zhang, F. Huang, H. Zhou, and R. Chandra, "Review of deep learning models for crypto price prediction: Implementation and evaluation," *arXiv*, vol. 2024, no. May, pp. 1–18, May 2024, doi: 10.48550/arxiv.2405.11431.
- [12] N. S. Wen and L. S. Ling, "Evaluation of cryptocurrency price prediction using LSTM and CNNs models," *JOIV: Int. J. Informatics Visualization*, vol. 2023, no. Jul., pp. 1–10, Jul. 2023, doi: 10.30630/joiv.7.3-2.2344.
- [13] W. Badar, S. Ramzan, A. Raza, N. L. Fitriyani, M. Syafrudin, and S. Lee, "Enhanced interpretable forecasting of cryptocurrency prices using autoencoder features and a hybrid CNN-LSTM model," *Mathematics*, vol. 2025, no. Jun., pp. 1–15, Jun. 2025, doi: 10.3390/math13121908.
- [14] H. Zhao, M. Crane, and M. Bezbradica, "Attention! Transformer with sentiment on cryptocurrencies price prediction," *SCITEPRESS*, vol. 2022, no. Feb., pp. 1–10, Feb. 2022, doi: 10.5220/0011103400003197.
- [15] K. Murray, A. Rossi, D. Carraro, and A. Visentin, "On forecasting cryptocurrency prices: A comparison of machine learning, deep learning, and ensembles," *Forecasting*, vol. 2023, no. Jan., pp. 1–15, Jan. 2023, doi: 10.3390/forecast5010010.
- [16] N. Tripathy, S. Hota, D. Mishra, P. Satapathy, and S. K. Nayak, "Empirical forecasting analysis of Bitcoin prices," *Int. J. Electr. Comput. Eng. Syst.*, vol. 2024, no. Jan., pp. 1–12, Jan. 2024, doi: 10.32985/ijeces.15.1.3.
- [17] M. Nakakita, T. Toyabe, and T. Nakatsuma, "Bayesian analysis of Bitcoin volatility using minute-by-minute data and flexible stochastic volatility models," *Mathematics*, vol. 2025, no. Aug., pp. 1–15, Aug. 2025, doi: 10.3390/math13162691.
- [18] M. Schnaubelt, J. Rende, and C. Krauss, "Testing stylized facts of Bitcoin limit order books," *J. Risk Financial Manag.*, vol. 12, no. Jan., pp. 1–25, Jan. 2019, doi: 10.3390/jrfm12010025.
- [19] M. Zahid, F. Iqbal, and D. Koutmos, "Forecasting Bitcoin volatility using hybrid GARCH models with machine learning," *Risks*, vol. 2022, no. Dec., pp. 1–20, Dec. 2022, doi: 10.3390/risks10120237.
- [20] M. Z. Islam, M. S. Rahman, M. Sumsuzoha, B. Sarker, M. R. Islam, M. Alam, and S. K. Shil, "Cryptocurrency price forecasting using machine learning: Building intelligent financial prediction models," *arXiv*, vol. 2025, no. Aug., pp. 1–18, Aug. 2025, doi: 10.48550/arxiv.2508.01419.

- [21] H. Song and H. Choi, "Forecasting stock market indices using the recurrent neural network based hybrid models: CNN-LSTM, GRU-CNN, and ensemble models," *Applied Sciences*, vol. 2023, no. Apr., pp. 1–20, Apr. 2023, doi: 10.3390/app13074644.
- [22] I. Livieris, N. Kiriakidou, S. Stavroyiannis, and P. Pintelas, "An advanced CNN-LSTM model for cryptocurrency forecasting," *Electronics*, vol. 2021, no. Feb., pp. 1–15, Feb. 2021, doi: 10.3390/electronics10030287.
- [23] H. Zhao, M. Crane, and M. Bezbradica, "Attention! Transformer with sentiment on cryptocurrencies price prediction," *SCITEPRESS*, vol. 2022, no. Feb., pp. 1–10, Feb. 2022, doi: 10.5220/0011103400003197.
- [24] K. Murray, A. Rossi, D. Carraro, and A. Visentin, "On forecasting cryptocurrency prices: A comparison of machine learning, deep learning, and ensembles," *Forecasting*, vol. 2023, no. Jan., pp. 1–15, Jan. 2023, doi: 10.3390/forecast5010010.
- [25] K. He, Q. Yang, L. Ji, J. Pan, and Y. Zou, "Financial time series forecasting with the deep learning ensemble model," *Mathematics*, vol. 2023, no. Feb., pp. 1–15, Feb. 2023, doi: 10.3390/math11041054.
- [26] I. Livieris, E. G. Pintelas, S. Stavroyiannis, and P. Pintelas, "Ensemble deep learning models for forecasting cryptocurrency time-series," *Algorithms*, vol. 2020, no. May, pp. 1–15, May 2020, doi: 10.3390/a13050121.
- [27] G. Fantini, J. Jia, and C. Oldani, "Informational efficiency in cryptocurrency markets: A bibliometric and thematic literature review (2015–2024)," *J. Econ. Surveys*, vol. 2025, no. Mar., pp. 1–25, Mar. 2025, doi: 10.1111/joes.70015.
- [28] S. Praveena and P. S., "A hybrid deep learning based deep Prophet memory neural network approach for seasonal items demand forecasting," *J. Adv. Inf. Technol.*, vol. 2024, no. Jun., pp. 735–747, Jun. 2024, doi: 10.12720/jait.15.6.735-747.
- [29] S. Oprea, "Bitcoin price forecasting using blockchain metrics, ensemble machine learning algorithms and Facebook Prophet," *Sci. Bull. Naval Acad.*, vol. 2024, no. Jun., pp. 1–12, Jun. 2024, doi: 10.21279/1454-864x-24-i2-004.
- [30] I. Ghosh, A. V. Megaravalli, M. Abedin, and K. Topuz, "Prediction and decoding of metaverse coin dynamics: A granular quest using MODWT–Facebook’s Prophet–TBATS and XAI methodology," *Ann. Oper. Res.*, vol. 2025, no. Jan., pp. 1–25, Jan. 2025, doi: 10.1007/s10479-025-06491-1.
- [31] T. O. Kehinde, O. J. Adedokun, A. Joseph, K. M. Kabirat, H. A. Akano, and O. Olanrewaju, "Helformer: An attention-based deep learning model for cryptocurrency price forecasting," *J. Big Data*, vol. 2025, no. Mar., pp. 1–20, Mar. 2025, doi: 10.1186/s40537-025-01135-4.
- [32] K. U. Apu, M. M. Rahman, A. B. Hoque, and M. Bhuiyan, "Forecasting future investment value with machine learning, neural networks, and ensemble learning: A meta-analytic study," *Rev. Appl. Sci. Technol.*, vol. 2022, no. Jul., pp. 1–15, Jul. 2022, doi: 10.63125/edxgjj56.
- [33] N. Tripathy, S. Nayak, and S. Prusty, "A comparative analysis of Silverkite and inter-dependent deep learning models for Bitcoin price prediction," *Frontiers in Blockchain*, vol. 2024, no. Apr., pp. 1–15, Apr. 2024, doi: 10.3389/fbloc.2024.1346410.
- [34] G. Fantini, J. Jia, and C. Oldani, "Informational efficiency in cryptocurrency markets: A bibliometric and thematic literature review (2015–2024)," *J. Econ. Surveys*, vol. 2025, no. Mar., pp. 1–25, Mar. 2025, doi: 10.1111/joes.70015.

- [35] M. S. A. Dolon, "Hybrid machine learning–driven financial forecasting models: Integrating LSTM, Prophet, and XGBoost," *Rev. Appl. Sci. Technol.*, vol. 2025, no. Feb., pp. 1–15, Feb. 2025, doi: 10.63125/nr1j8527.
- [36] M. Nakakita, T. Toyabe, and T. Nakatsuma, "Bayesian analysis of Bitcoin volatility using minute-by-minute data and flexible stochastic volatility models," *Mathematics*, vol. 2025, no. Aug., pp. 1–15, Aug. 2025, doi: 10.3390/math13162691.