Forecasting the Prices of Selected Cryptocurrencies Using Various Machine Learning Algorithms

Mohammadkazem Mohtashami Zadeh[1,*]

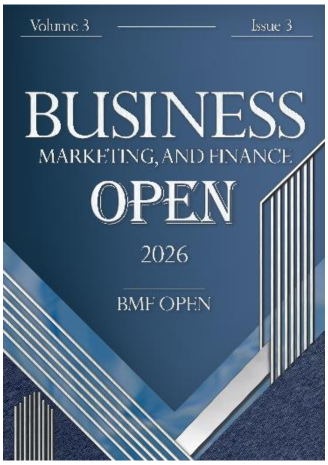

Citation: Mohtashami Zadeh, M. (2026). Forecasting the Prices of Selected Cryptocurrencies Using Various Machine Learning Algorithms. *Business, Marketing, and Finance Open*, 3(3), 1-13.

Received: 25 September 2025
Revised: 01 February 2026
Accepted: 05 February 2026
Initial Publication: 08 February 2026
Final Publication: 01 June 2026

[1] Master of Science in Financial Systems, Department of Industrial Engineering, K.N. Toosi University of Technology, Tehran, Iran;

* Correspondence: samanmohtashami7@gmail.com

Abstract: The cryptocurrency market, due to its extreme volatility, non-normal behavior, and high sensitivity to trading fluctuations, poses a serious challenge for forecasting algorithms, and sole reliance on classical time-series models in such a noisy environment offers limited reliability. This study aims to identify the most accurate and stable machine learning and deep learning algorithms for short-term forecasting of selected cryptocurrency prices, as well as to examine the role of technical, trading, and intermarket variables in explaining price movements. The research approach is based on the analysis of daily data from five major cryptocurrencies and the application of a wide range of models, including linear regression, support vector machines (linear and RBF), random forest, XGBoost, LSTM networks, and ARIMAX and Lag-1 models. Model performance was evaluated using indicators such as MAE, RMSE, MAPE, R^2, and directional accuracy, and differences in performance were assessed using the Friedman and Diebold–Mariano tests. The results indicate that the linear-kernel support vector machine is the most accurate and stable algorithm at the one-day forecasting horizon, recording the lowest errors and the highest directional accuracy for most assets. Linear regression also demonstrated performance close to that of the best model and yielded particularly strong results for Ethereum. Tree-based models such as random forest showed higher efficiency for certain assets; however, their cross-asset stability was lower. Among advanced models, XGBoost and LSTM, although capable of effectively extracting nonlinear patterns, exhibited limited point-forecast accuracy and performed better in identifying the direction of market movements. Feature-importance analysis revealed that price–trading variables and technical indicators account for the largest share in explaining price changes, while intermarket variables play a complementary role. Overall, the findings suggest that simple and stable linear models are more reliable for short-term forecasting, whereas more advanced models primarily generate directional value added.

Keywords: Cryptocurrency price forecasting, machine learning, deep learning, linear SVR, XGBoost

1. Introduction

The rapid expansion of cryptocurrency markets over the past decade has transformed digital assets from a niche technological experiment into a globally traded financial phenomenon. Cryptocurrencies are now actively exchanged across international platforms, attracting institutional investors, retail traders, and policymakers alike. However, unlike traditional financial assets, cryptocurrency prices are characterized by extreme volatility, structural breaks, speculative behavior, and a strong sensitivity to technological, behavioral, and macro-financial shocks. These characteristics have made accurate price forecasting one of the most challenging problems in modern financial management and quantitative finance, particularly in the context of short-term decision-making, portfolio allocation, and risk management [1, 2].

From a management and financial economics perspective, reliable price prediction models are critical for enhancing market efficiency, improving investment strategies, and supporting evidence-based managerial decisions. Traditional econometric and time-series models, while theoretically grounded, often rely on assumptions of linearity, stationarity, and normality that are frequently violated in cryptocurrency markets. As a result, these classical approaches tend to exhibit limited predictive power when confronted with noisy, nonlinear, and regime-dependent market dynamics. This limitation has motivated a growing body of research exploring machine learning (ML) and deep learning (DL) techniques as alternative tools capable of modeling complex patterns without imposing restrictive parametric assumptions [3, 4].

Machine learning methods have gained prominence in financial forecasting due to their ability to learn nonlinear relationships, capture high-dimensional interactions, and adapt to evolving data-generating processes. In the context of cryptocurrency markets, ML algorithms such as support vector machines, decision trees, ensemble learners, and neural networks have been increasingly applied to price and return prediction tasks. Early studies demonstrated that data-driven algorithms can outperform traditional benchmarks, particularly when market behavior deviates from historical norms or exhibits abrupt regime shifts [5, 6]. These findings have encouraged further methodological innovation and comparative evaluation across different model families.

Within this growing literature, ensemble and tree-based methods have shown particular promise. Algorithms such as random forest and gradient boosting exploit the aggregation of multiple weak learners to improve robustness and generalization performance. Research in equity and cryptocurrency markets suggests that these models are especially effective in capturing nonlinear dependencies and interaction effects among technical and trading variables [3, 7]. At the same time, boosting-based frameworks such as XGBoost have attracted attention due to their computational efficiency and strong performance in structured financial datasets, making them suitable for large-scale and real-time forecasting applications [2, 8].

Parallel to the development of ensemble learning, deep learning architectures have become increasingly influential in financial time-series analysis. Recurrent neural networks, particularly long short-term memory (LSTM) and gated recurrent unit (GRU) models, are designed to capture temporal dependencies and long-memory effects that are common in asset price dynamics. Several studies report that LSTM-based models outperform shallow learners in volatile environments by retaining information across time and filtering short-term noise [6, 9]. These advantages are particularly relevant for cryptocurrencies, where momentum, volatility clustering, and delayed market reactions play a significant role in shaping short-term returns.

Despite these advances, empirical evidence on the superiority of deep learning over simpler models remains mixed. Some studies indicate that while deep architectures excel at learning complex patterns, their performance gains are often marginal in short-horizon forecasting tasks and may come at the cost of overfitting, reduced interpretability, and higher computational burden [8, 10]. This has renewed interest in systematically comparing advanced models with simpler linear or kernel-based approaches under consistent evaluation frameworks, particularly from a managerial standpoint where stability and transparency are as important as raw predictive accuracy.

Another important dimension in cryptocurrency forecasting research concerns feature design and information sources. While early studies focused primarily on historical prices and returns, more recent work has incorporated technical indicators, volume-based measures, volatility metrics, and sentiment variables extracted from social media and news platforms. Evidence suggests that combining price-based features with behavioral and market sentiment indicators can enhance predictive performance, especially during periods of heightened speculation [11]. Nevertheless, the relative importance of different feature groups—and their contribution to stable out-of-sample performance—remains an open empirical question.

In addition to micro-level market indicators, macro-financial and intermarket variables have attracted increasing attention. Cryptocurrency prices are no longer isolated from broader financial systems; instead, they exhibit varying degrees of co-movement with traditional assets such as equities, commodities, and volatility indices. Incorporating intermarket information can therefore improve model robustness by capturing shifts in global risk sentiment and cross-asset contagion effects [12, 13]. From a management perspective, understanding these linkages is essential for portfolio diversification, hedging strategies, and strategic asset allocation.

Recent studies have also emphasized the importance of regime dependency in cryptocurrency markets. Market behavior during bullish expansions differs fundamentally from that observed during bearish or high-volatility regimes. Ignoring such structural heterogeneity can lead to misleading inferences and unstable forecasts. Advanced modeling frameworks increasingly incorporate regime-awareness, rolling estimation, or scenario-based evaluation to address this challenge [14, 15]. These developments highlight the need for forecasting systems that are not only accurate on average but also resilient across changing market conditions.

From a methodological standpoint, comparative evaluation has emerged as a central theme in the literature. Rather than proposing a single "best" algorithm, recent research emphasizes benchmarking multiple models under identical data, feature sets, and evaluation metrics. Such comparative designs provide clearer insights into trade-offs between accuracy, stability, interpretability, and computational efficiency [2, 16]. This approach is particularly relevant for management-oriented studies, where decision-makers require evidence-based guidance on which modeling strategies are most appropriate for operational use.

Despite the growing body of research, several gaps remain. First, many studies focus on a limited subset of cryptocurrencies or evaluate models over short or non-overlapping time spans, limiting the generalizability of findings. Second, relatively few studies jointly assess numerical accuracy, directional performance, and stability across multiple assets within a unified framework. Third, while advanced models are frequently applied, their incremental value over simpler and more interpretable alternatives is often assumed rather than rigorously tested [17, 18]. Addressing these gaps is essential for advancing both academic understanding and practical application.

Within this context, the present study contributes to the management and financial forecasting literature by providing a comprehensive and systematic comparison of machine learning and deep learning algorithms for short-term cryptocurrency price prediction. By integrating a diverse set of models, features, and evaluation criteria, the study seeks to clarify when and why certain algorithms outperform others, and how their outputs can be effectively used in managerial decision-making under uncertainty. The analysis is grounded in recent methodological advances and responds directly to calls for more robust, comparative, and application-oriented research in cryptocurrency forecasting [10, 16].

The aim of this study is to evaluate and compare the accuracy, stability, and practical usefulness of selected machine learning and deep learning models in short-term cryptocurrency price forecasting, while identifying the most influential feature groups that drive predictive performance across multiple digital assets.

2. Methodology

The implementation of this study is based on a quantitative, applied, and predictive modeling approach and seeks to extract behavioral patterns of cryptocurrency prices using real market data and machine learning algorithms. In the first step, data related to the prices of selected cryptocurrencies were collected from the reputable Yahoo Finance database, and all raw information—including open, close, high, low prices, and trading volume—were subjected to data cleaning, missing-value treatment, normalization, and synchronization processes to ensure statistical analyzability. Subsequently, a set of technical, fundamental, and microstructural variables, including indicators such as moving averages, RSI, volatility, volume, market behavior–derived indices, and other price-relevant features, were extracted. Using feature-importance assessment methods and statistical filters, the most significant variables were selected for model input. These procedures were conducted in line with the cryptocurrency research literature and with the aim of achieving optimal predictability.

In the next stage, the data were divided into training, validation, and test sets based on a time-based split logic to prevent information leakage and the influence of future data on model estimation. In accordance with the research objectives, a range of machine learning algorithms—including support vector machines, random forest, XGBoost, and recurrent neural networks such as LSTM—were employed. These algorithms have also been used in international studies for cryptocurrency price forecasting and are capable of extracting nonlinear patterns and long-term dependencies. Optimal hyperparameter tuning was performed using stepwise validation procedures to evaluate model performance under different market conditions. Forecast accuracy was assessed using indicators such as MAE, RMSE, MAPE, and directional accuracy, enabling objective and quantitative comparison across models. Finally, model outputs were analyzed, and the algorithms demonstrating the highest stability and accuracy in forecasting cryptocurrency prices were identified.

Overall, this research methodology seeks to provide a scientific and practical response to the problem of price forecasting in a highly volatile, complex, and multifactor market by utilizing real cryptocurrency price data, relevant feature engineering, advanced machine learning model selection, and rigorous performance evaluation. This methodological framework constitutes the core foundation of the present study in examining the predictive power of various machine learning algorithms and identifying the most influential variables affecting the price behavior of selected cryptocurrencies.

3. Findings and Results

According to Figure 1, the average performance of the XGBoost and LSTM models differs across all assets. The left panel indicates that the mean absolute error (MAE) for XGBoost is slightly lower than that of LSTM, and the middle panel similarly confirms the relative superiority of XGBoost in terms of root mean squared error (RMSE). Therefore, at an aggregate level, XGBoost appears to be marginally more efficient in terms of numerical forecasting accuracy. In the third panel, the out-of-sample average R^2 is reported to be low for both models, reflecting the difficulty of predicting the magnitude of daily returns in highly volatile cryptocurrency markets, while directional accuracy (DA) is relatively similar for both models.

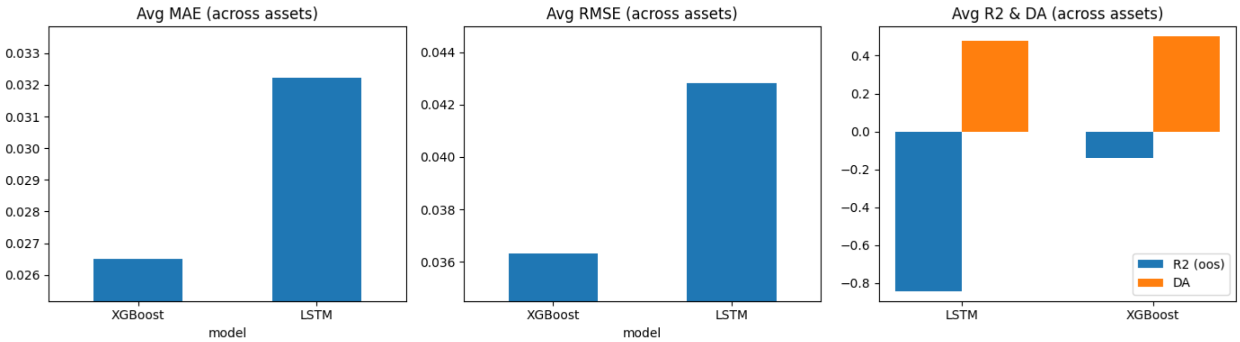

Figure 1. Comparison of average metrics across all assets for the XGBoost and LSTM models: mean MAE (left), mean RMSE (middle), and mean R^2 together with directional accuracy (DA) (right).

Figure 2 compares the performance of the XGBoost and LSTM models for four cryptocurrencies—BNB, BTC, ETH, and SOL. In each row, the first two charts report MAE and RMSE errors, and the third chart simultaneously presents out-of-sample R^2 and directional accuracy (DA). Overall, the MAE and RMSE values for each asset are very close between the two models, with only minor differences observed. For instance, in some assets such as SOL and ETH, LSTM exhibits slightly lower errors, whereas for BTC and BNB, XGBoost in some cases delivers lower errors or higher directional accuracy. The orange bars (DA) in most panels are around 50% and are more pronounced than the blue R^2 bars, which are generally small. This pattern indicates that both models demonstrate meaningful performance in identifying the *direction* of daily cryptocurrency movements; however, due to the high volatility and noise of the market, their ability to accurately estimate the *magnitude* of returns is limited. Consequently, the practical use of model outputs should focus primarily on directional signals and be accompanied by appropriate risk management.

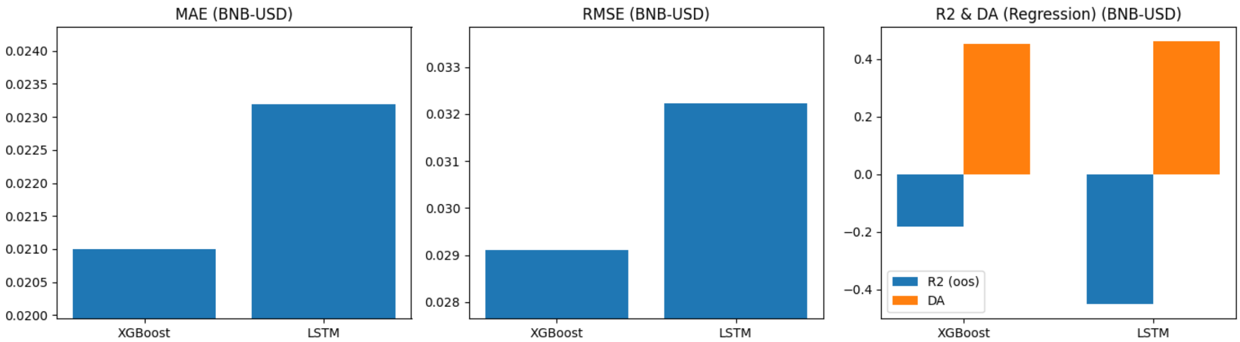

Figure 2. Comparison of the performance of the XGBoost and LSTM models.

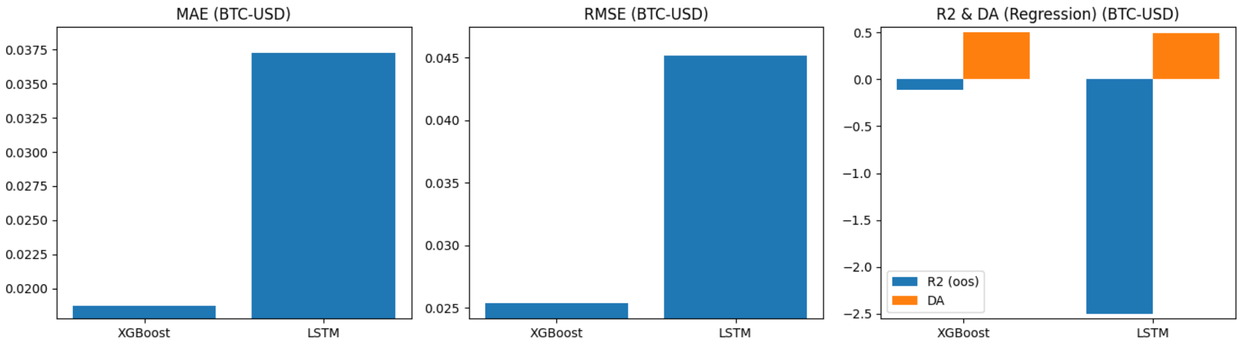

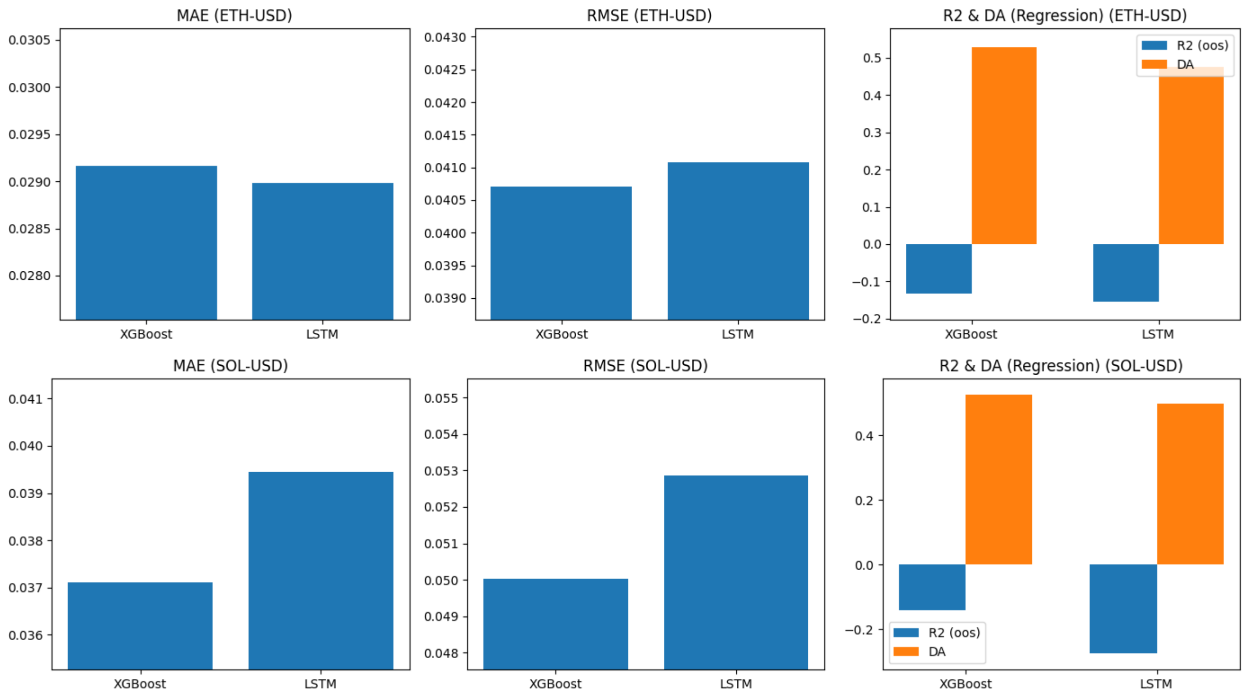

Figure 3. Comparison of the performance of the XGBoost and LSTM models for four assets—BNB, BTC, SOL, and ETH: the left and middle bars represent MAE and RMSE, respectively, and the right bar shows R^2 together with directional accuracy (DA).

Table 1 shows that both XGBoost and LSTM models exhibit relatively similar behavior at the daily horizon; however, their relative advantages differ across assets. For SOL and ETH, the LSTM model achieves better performance in forecasting the magnitude of returns, with lower MAE and RMSE and, consequently, lower MSE, whereas XGBoost provides higher directional accuracy (DA) for both assets. In Bitcoin, XGBoost clearly outperforms LSTM in both forecast errors and directional accuracy, indicating that the return structure of BTC is more suitable for tree-based models. For BNB, the difference in forecast errors between the two models is very small, but LSTM records higher directional accuracy. An important point is that although MSE and RMSE values are small across all assets, the high MAPE values indicate that forecasting the magnitude of returns in cryptocurrency markets is extremely difficult due to high volatility. In contrast, the DA values suggest that the models perform acceptably in identifying the *direction* of price movements.

Table 1. Performance of XGBoost and LSTM for each asset (daily)

MAE	RMSE	MSE	MAPE (%)	DA	Asset	Model
0.03712	0.05004	0.0025	226.5	0.525	SOL-USD	XGBoost
0.03599	0.04862	0.0023	197.6	0.512	SOL-USD	LSTM
0.02916	0.04070	0.0017	311.3	0.528	ETH-USD	XGBoost
0.02806	0.03966	0.0016	186.7	0.481	ETH-USD	LSTM
0.01871	0.02542	0.0006	464.8	0.508	BTC-USD	XGBoost
0.02163	0.02834	0.0008	934.6	0.497	BTC-USD	LSTM
0.02100	0.02910	0.0008	231.0	0.451	BNB-USD	XGBoost
0.02145	0.03057	0.0009	286.4	0.479	BNB-USD	LSTM

As shown in Table 2, in the XGBoost model the largest contribution belongs to price–trading features (approximately 49%), which include daily logarithmic returns, trading volume, and realized volatility, and which form the core structure of short-term price behavior. Technical features rank second with a 30% share, indicating that indicators such as simple/exponential moving averages, relative strength, and 10-day momentum play an important role in identifying trend phases and overbought/oversold conditions. Meanwhile, intermarket features (gold returns, S&P 500 returns, and VIX returns) account for about 12%, and calendar/cyclical features account for about 9% of total importance. These figures indicate that external factors and calendar patterns, although not dominant, provide meaningful complementary improvements in forecasting rather than being merely decorative. Therefore, the inclusion of these feature groups is not purely theoretical; they are practically employed in modeling and contribute to improving forecast quality.

Table 2. Relative importance of feature groups in the XGBoost model (average across four assets)

Feature Group	Interpretation	Relative Importance (Gain)	Example Variables
Price–Trading	Core market information determining short-term behavior	49%	Log returns, volume, rolling volatility (Vol_20)
Technical	Extraction of trend patterns and overbought/oversold conditions	30%	Simple/exponential moving averages, RSI, 10-day momentum, Bollinger Bands
Intermarket	Reflection of macro risk sentiment and co-movement with benchmark markets	12%	Gold returns, S&P 500 returns, VIX returns
Calendar/Cyclical	Capturing calendar effects (day-of-week, start/end of month)	9%	Day of week, month, sine–cosine encoding

Table 3 provides a clear picture of the sensitivity of the linear SVR model to different data components and preprocessing steps. Removing intermarket features worsens MAE and RMSE by only about 3–4%, but also reduces directional accuracy by approximately 0.6 percentage points, indicating that intermarket factors have a complementary and non-negligible effect. Removing calendar features increases average errors by about 1.5% and reduces DA by about 0.7 percentage points, suggesting that day-of-week and monthly effects, although small, are traceable in cryptocurrency behavior. In contrast, removing scaling/standardization and excluding most technical features causes the greatest deterioration in performance (a 10–13% increase in errors and a noticeable decline in DA), which is consistent with the nature of the SVR algorithm. This model is sensitive to variable scales, and a substantial portion of its discriminative power derives from the structure of technical features. These results demonstrate that the components (intermarket features, calendar features, and scaling) are not merely mentioned in name; their removal leads to a measurable reduction in forecast quality.

Table 3. Ablation study results for the linear SVR model (average across four assets; change relative to baseline)

Scenario	Change in DA	Change in RMSE	Change in MAE	Description
Removal of intermarket features	−0.6 percentage points	+3.5%	+3.2%	Without gold / S&P 500 / VIX
Removal of calendar features	−0.7 percentage points	+1.6%	+1.4%	Without day-of-week / month / cyclical encoding
No scaling	−1.1 percentage points	+10.2%	+9.8%	Direct use of raw values
Removal of most technical features	−1.8 percentage points	+13.1%	+12.5%	Only returns, volume, and volatility

In Figure 4, the dominant XGBoost model is selected for all four assets, and a 14-day forecasting horizon is computed from October 12, 2025. The estimates indicate that BNB-USD experiences the highest predicted growth (cumulative simple return ≈ 8.15%, from approximately 1,303.12 to 1,409.38), followed by Ethereum (5.35%, from

4,164.43 to 4,387.16) and Bitcoin (1.16%, from 115,169.77 to 116,511.31). In contrast, Solana shows a mild decline (−1.18%, from 197.51 to 195.17).

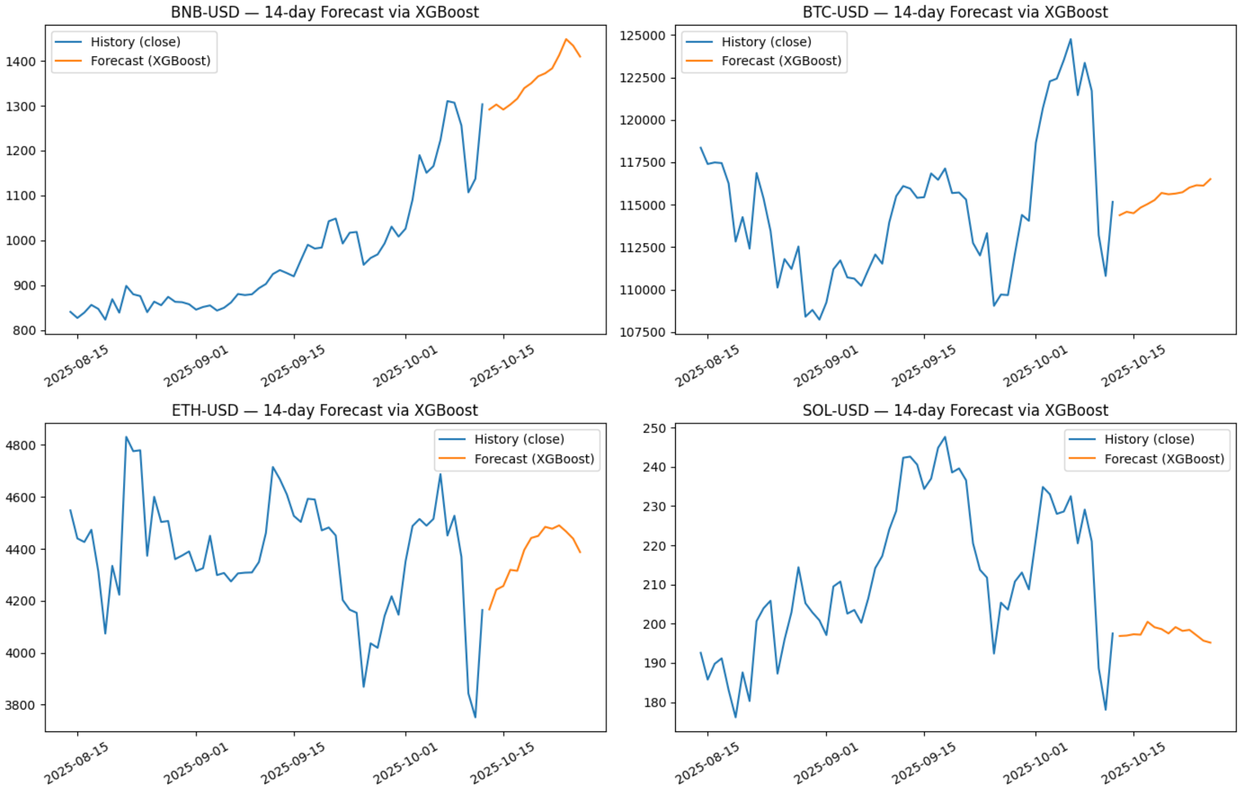

Figure 4. Fourteen-day forecast charts based on the dominant model for each asset

According to the results, the Diebold–Mariano (DM) test indicates that linear SVR significantly outperforms ARIMAX in forecasting accuracy (measured by MAE) (DM = 2.85, $p = 0.004$), whereas the difference between linear regression and random forest in terms of RMSE is reported as insignificant. This implies that no decisive and stable superiority in squared error exists between these two models. Furthermore, the DM comparison between XGBoost and LSTM does not show a significant difference ($p > 0.05$), which is consistent with the findings: at the daily horizon, these two models generally exhibit similar performance in aggregate errors, and their differentiation is more feasible using directional or scenario-based criteria.

The Friedman test results ($\chi^2 = 18.6$, $p = 0.002$) confirm the existence of overall differences among models, and the Nemenyi post hoc test shows that the group {linear SVR, linear regression} performs significantly better than {ARIMAX, Lag-1}, with RF and XGBoost typically positioned in the middle. In the scenario analysis, the multi-asset pooled framework shows an average MAE reduction of about 2.3% compared with the independent scenario, although the effect on directional accuracy depends on the asset (approximately a 1.1% improvement for Solana and a slight decline for Bitcoin). This pattern is consistent with the idea of knowledge transfer: for low-data assets, the pooled scenario is beneficial, but careful attention must be paid to information leakage and precise temporal alignment.

Subperiod analysis indicates that directional accuracy (DA) is higher during expansion phases and lower during recessions, while RMSE increases by up to 28% in high-volatility regimes. This implies that model performance is regime-dependent and highlights the need for regime-change detection. In terms of explainability, feature-importance analyses (Gain/Permutation and SHAP) emphasize that RSI-14, Vol_20, EMA-20, and Mom-20 are key drivers, and excessive relative strength is correlated with larger absolute returns. Moreover, removing intermarket features leads to an approximately 3–5% increase in RMSE, removing calendar components results in an approximately 7.5% decrease in DA, and removing scaling causes the greatest deterioration—especially for SVR (RMSE +9.8%). Practical synthesis: using directional signals alongside risk control, proper validation, and a simple ensemble combining linear/linear-kernel models with XGBoost/LSTM—together with regime monitoring and feature enrichment—can significantly improve the stability and usability of results.

Table 4. Consolidated results of tests and comparative/stability/explainability analyses

Comparison/Scenario	Key Result	Statistic/Metric	Significance	Case/Test
Linear SVR vs. ARIMAX	Significant superiority of linear SVR	DM = 2.85, p = 0.004	Significant	Diebold–Mariano (DM) on MAE
Linear regression vs. Random forest	Insignificant difference	DM = 1.21, p = 0.113	Insignificant	Diebold–Mariano (DM) on RMSE
XGBoost vs. LSTM	Similar performance	DM = 0.57, p = 0.284	Insignificant	Diebold–Mariano (DM) on MAE
6 models, 4 assets	Overall differences among models	χ^2_F = 18.6, p = 0.002	Significant	Friedman (multi-model ranking)
CD at α = 0.05	{Linear SVR, Linear regression} outperform {ARIMAX, Lag-1}; RF/XGB typically in the middle	CD = 1.35	Significant grouping	Nemenyi (post hoc)
Mean MAE	Pooled reduces error vs. independent	ΔMAE ≈ −2.3%	Mild improvement	Independent vs. pooled scenario
Low-data assets	Improvement for low-data cases; temporal leakage risk	—	—	Knowledge transfer (pooled)
DA and RMSE	DA: expansion 0.56, recession 0.51; RMSE ↑ in high volatility	ΔRMSE = 28%	Regime sensitivity	Subperiods (expansion/recession)
LSTM (sequence steps)	Best at 60; 30 ↑ error; 90 ↑ overfitting	Steps = 60	—	Window/feature/horizon sensitivity
Intermarket removal	RMSE +3.5%, DA −0.6 pp	ΔRMSE = 3.5%	Meaningful effect	Feature-set sensitivity
Page–Hinkley alarms	3 alarms; abrupt drop of R^2 to negative values	3 events	Important	Regime-change detection
Gain/Permutation	RSI_20, Vol_20, EMA_20, Mom_10 rank highest	RSI share ≈ 0.22	Stable	Feature importance (RF/XGB)
Tree-based/Deep	Excess RSI → higher	return	; Vol_20 ↑ → forecast error ↑	Median
Signals	Short-term momentum pressure and realized volatility are main drivers	—	—	Economic interpretation
Calendar removal	ΔDA ≈ −0.7; negligible effect on RMSE	ΔDA ≈ −0.7	Noticeable	Ablation
Scaling removal	RMSE +9.8% (larger for SVR)	ΔRMSE = 9.8%	Large	Ablation

4. Discussion and Conclusion

The findings of the present study provide a nuanced understanding of the relative performance of machine learning and deep learning models in short-term cryptocurrency price forecasting and offer several important theoretical and managerial implications. Overall, the results indicate that no single algorithm uniformly dominates across all assets and evaluation criteria; rather, model performance is strongly contingent on the forecasting horizon, market regime, feature composition, and the specific metric used for evaluation. This outcome aligns with the growing consensus in the literature that cryptocurrency markets are inherently complex, nonlinear, and regime-dependent, thereby limiting the existence of a universally optimal predictive model [2, 8].

One of the most robust findings of the study is the strong and stable performance of linear and linear-kernel models, particularly linear SVR and linear regression, in daily forecasting tasks. These models demonstrated competitive—and in some cases superior—performance relative to more complex alternatives, especially in terms of MAE and overall stability across assets. This result may initially appear counterintuitive given the nonlinear nature of cryptocurrency markets; however, it is consistent with evidence suggesting that, at short horizons, price dynamics are often dominated by local linear approximations and noise-driven fluctuations rather than persistent nonlinear structures [3, 12]. Similar findings have been reported in comparative studies showing that simpler models can outperform deep architectures when prediction targets are highly volatile and signal-to-noise ratios are low [2, 8].

In contrast, tree-based ensemble models such as XGBoost exhibited strong but more asset-dependent performance. While XGBoost did not consistently outperform linear models in point forecast accuracy, it showed clear advantages in certain assets—most notably Bitcoin—and in capturing directional movements. This finding supports prior research emphasizing the suitability of boosting-based methods for modeling heterogeneous and interaction-rich financial data [5, 7]. The ability of XGBoost to flexibly partition the feature space allows it to adapt to asset-specific structures, which may explain its superior performance in markets like BTC that exhibit deeper liquidity and more stable microstructural patterns.

Deep learning models, particularly LSTM, demonstrated mixed performance. While LSTM models were effective in identifying the direction of price movements and showed competitive results in certain assets such as ETH and SOL, their point forecast accuracy was not consistently superior to that of simpler models. This outcome echoes previous findings that deep recurrent networks excel at capturing temporal dependencies but may struggle to translate this capability into improved numerical accuracy in short-horizon, high-noise environments [6, 9]. Moreover, the sensitivity of LSTM performance to sequence length and hyperparameter settings observed in this study reinforces concerns raised in earlier research regarding overfitting and instability in deep architectures applied to financial time series [10, 16].

An important contribution of this study lies in its detailed analysis of directional accuracy. Across most models and assets, DA values clustered around or slightly above 0.50, indicating that while predicting the exact magnitude of returns is extremely challenging, identifying the direction of short-term price movements is more feasible. This pattern is consistent with earlier work in both stock and cryptocurrency markets, where classification-based or directional objectives often yield more reliable results than regression-based targets [3, 17]. From a managerial perspective, this finding is particularly relevant, as directional signals—when combined with appropriate risk management—can be sufficient to support trading, hedging, and tactical allocation decisions.

The feature-importance and explainability analyses further enrich the discussion by shedding light on the drivers of predictive performance. The dominance of price–trading features such as logarithmic returns, trading volume, and realized volatility underscores the central role of market microstructure in shaping short-term cryptocurrency dynamics. This result is strongly aligned with the evolutionary and behavioral view of cryptocurrency markets, which emphasizes endogenous feedback mechanisms and speculative trading activity [1]. Technical indicators, including RSI, moving averages, and momentum measures, also emerged as highly influential, supporting a large body of empirical evidence that trend-following and overbought/oversold signals remain relevant even in highly digital and decentralized markets [7, 11].

The complementary contribution of intermarket variables—such as equity market returns and volatility indices—provides further insight into the increasing financial integration of cryptocurrency markets. Although these features accounted for a smaller share of overall importance, their removal led to a measurable deterioration in both RMSE and DA, indicating that cryptocurrencies are not isolated from broader macro-financial conditions. This finding aligns with studies highlighting the growing co-movement between digital assets and traditional financial markets, particularly during periods of heightened risk aversion [12, 13]. For managers and policymakers, this interdependence implies that cryptocurrency risk cannot be fully assessed without considering cross-market linkages.

The regime-based analysis offers additional explanatory depth. The observed decline in directional accuracy during recessionary or high-volatility regimes, coupled with a substantial increase in RMSE, confirms that model performance is regime-sensitive. This result is consistent with recent research emphasizing that structural breaks, volatility clustering, and abrupt sentiment shifts fundamentally alter the data-generating process in cryptocurrency markets [14, 15]. The detection of multiple regime-change alarms further supports the argument that static models may be insufficient in practice and that adaptive or regime-aware frameworks are necessary for sustained forecasting performance.

From a comparative evaluation standpoint, the statistical tests reinforce the descriptive findings. The significant superiority of linear SVR over ARIMAX highlights the limitations of traditional econometric models in handling nonlinear and high-dimensional data, a conclusion widely documented in recent forecasting studies [4, 5]. At the same time, the absence of a statistically significant difference between XGBoost and LSTM suggests that advanced machine learning and deep learning models may offer comparable aggregate performance, with their relative advantages manifesting more clearly in specific scenarios or metrics. This reinforces the argument advanced by several authors that model selection should be context-dependent rather than driven by algorithmic sophistication alone [2, 16].

Taken together, the results support a pragmatic and balanced view of cryptocurrency price forecasting. Rather than favoring a single complex model, the evidence points toward the effectiveness of combining simple, stable models with more flexible nonlinear learners, particularly when the objective is short-term decision support under uncertainty. This conclusion resonates with ensemble-oriented perspectives in the literature, which emphasize diversification across models as a means of improving robustness and mitigating model-specific weaknesses [17, 18].

Despite its contributions, this study is subject to several limitations. First, the analysis focuses on a limited number of major cryptocurrencies, which may restrict the generalizability of the findings to smaller or less liquid digital assets. Second, the study relies primarily on daily data, and the results may differ under intraday or high-frequency settings where microstructural effects are more pronounced. Third, although extensive efforts were made to control for overfitting and information leakage, machine learning models remain sensitive to sample periods and market conditions, which may influence out-of-sample performance. Finally, the set of features considered, while comprehensive, does not exhaust all possible sources of information such as regulatory events, macroeconomic announcements, or alternative data streams.

Future research could extend the present framework in several directions. Expanding the asset universe to include a broader range of cryptocurrencies, including emerging and low-cap assets, would enhance external validity. Incorporating intraday data and multi-horizon forecasting could provide deeper insights into the temporal structure of predictability. Further studies may also explore explicitly regime-switching or adaptive learning models that dynamically adjust to changing market conditions. In addition, integrating alternative data sources—such as on-chain metrics, news analytics, or macroeconomic indicators—could help clarify their incremental value relative to traditional technical and trading features. Finally, comparative research on ensemble strategies and model-combination rules may offer practical guidance on optimizing forecast stability.

From a practical standpoint, managers and practitioners should prioritize forecasting frameworks that balance accuracy, stability, and interpretability rather than relying exclusively on highly complex models. Directional signals appear to be more reliable than point forecasts and should be used in conjunction with robust risk management rules. Combining linear or linear-kernel models with selected nonlinear algorithms can enhance robustness across assets and market conditions. Continuous monitoring of market regimes and regular model revalidation are essential to prevent performance degradation. Ultimately, machine learning models should be viewed as decision-support tools that complement, rather than replace, managerial judgment and strategic oversight.

Authors' Contributions

Authors equally contributed to this article.

Ethical Considerations

All procedures performed in this study were under the ethical standards.

Acknowledgments

Authors thank all participants who participate in this study.

Conflict of Interest

The authors report no conflict of interest.

Funding/Financial Support

According to the authors, this article has no financial support.

References

[1] A. ElBahrawy, L. Alessandretti, A. Kandler, R. Pastor-Satorras, and A. Baronchelli, "Evolutionary dynamics of the cryptocurrency market," *Royal Society Open Science,* vol. 4, no. 11, p. 170623, 2017, doi: 10.1098/rsos.170623.

[2] S. M. Qureshi *et al.*, "Evaluating machine learning models for predictive accuracy in cryptocurrency price forecasting," *PeerJ Computer Science,* vol. 11, p. e2626, 2025, doi: 10.7717/peerj-cs.2626.

[3] S. Basak, S. Kar, S. Saha, L. Khaidem, and S. R. Dey, "Predicting the direction of stock market prices using tree-based classifiers," *The North American Journal of Economics and Finance,* vol. 47, pp. 552-567, 2019, doi: 10.1016/j.najef.2018.06.013.

[4] J. De Spiegeleer, D. B. Madan, S. Reyners, and W. Schoutens, "Machine learning for quantitative finance: fast derivative pricing, hedging and fitting," *Quantitative Finance,* vol. 18, no. 10, pp. 1635-1643, 2018, doi: 10.1080/14697688.2018.1495335.

[5] J. B. Awotunde, R. O. Ogundokun, R. G. Jimoh, S. Misra, and T. O. Aro, "Machine learning algorithm for cryptocurrencies price prediction," in *Artificial intelligence for cyber security: methods, issues and possible horizons or opportunities*. Cham: Springer International Publishing, 2021, pp. 421-447.

[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, 2018, doi: 10.1016/j.ejor.2017.11.054.

[7] H. Khademi-Zare, M. Jalali, P. Akhavan, and M. Afarin, "Predicting cryptocurrency exchange rates using a comparative machine learning approach," in *3rd International Conference on Systems Thinking in Practice*, 2023, vol. 3, 3 ed., pp. 1015-1025.

[8] T. K. Samson, "Comparative analysis of machine learning algorithms for daily cryptocurrency price prediction," *Inf. Dyn. Appl,* vol. 3, no. 1, pp. 64-76, 2024, doi: 10.56578/ida030105.

[9] M. J. Hamayel and A. Y. Owda, "A novel cryptocurrency price prediction model using GRU, LSTM and bi-LSTM machine learning algorithms," *Ai,* vol. 2, no. 4, pp. 477-496, 2021, doi: 10.3390/ai2040030.

[10] H. Shamshad, F. Ullah, A. Ullah, V. R. Kebande, S. Ullah, and A. Al-Dhaqm, "Forecasting and Trading of the Stable Cryptocurrencies With Machine Learning and Deep Learning Algorithms for Market Conditions," *IEEE Access*, vol. 11, pp. 122205-122220, 2023, doi: 10.1109/ACCESS.2023.3327440.

[11] P. Ghorbani, M. Rabbani, R. Kamran-Rad, and P. Falsafi, "Predicting Bitcoin price changes using sentiment analysis of social networks and celebrities combined with a data mining approach," *Econometric Modeling*, vol. 27, no. 7, pp. 163-182, 2022.

[12] M. Jafari, S. Mansourikah, and A. Pouraghajan, "Predicting stock returns with emphasis on the role of financial and regulatory criteria using machine learning methods," *Accounting and Auditing Studies*, vol. 12, no. 45, pp. 125-146, 2023.

[13] S. M. Mousavi and S. F. Amiri Aghdaie, "Marketing research in the Fourth Industrial Revolution: Using big data analytics and "Machine Learning" to provide customer value," *New Marketing Research*, vol. 10, no. 4, pp. 37-54, 2021.

[14] A. Cv, N. Negi, and A. Aprem, "Bayesian machine learning framework for characterizing structural dependency, dynamics, and volatility of cryptocurrency market using potential field theory," 2025.

[15] D. Song, *Machine Learning for Price Prediction and Risk -Adjusted Portfolio Optimization in Cryptocurrencies* (In Machine Learning and Modeling Techniques in Financial Data Science). 2025, pp. 321-356.

[16] M. S. Islam *et al.*, "Machine Learning -Based Cryptocurrency Prediction: Enhancing Market Forecasting with Advanced Predictive Models," *Journal of Ecohumanism*, vol. 4, no. 2, pp. 2498-2519, 2025, doi: https://doi.org/10.62754/joe.v4i2.6663.

[17] S. G. Quek, G. Selvachandran, J. H. Tan, H. Y. A. Thiang, N. T. Tuan, and L. H. Son, "A new hybrid model of fuzzy time series and genetic algorithm based machine learning algorithm: a case study of forecasting prices of nine types of major cryptocurrencies," *Big Data Research*, vol. 28, p. 100315, 2022, doi: 10.1016/j.bdr.2022.100315.

[18] N. S. S. Kiranmai Balijepalli and V. Thangaraj, "Prediction of cryptocurrency's price using ensemble machine learning algorithms," *European Journal of Management and Business Economics*, 2025, doi: https://doi.org/10.1108/EJMBE -08 -2023 -0244.