JURNAL FASILKOM
Volume 16 No. 2 | Agustus 2026: 431-436

P-ISSN : 2089-3353
E-ISSN : 2808-9162

Multi-Horizon Bitcoin Volatility Forecasting Using Support Vector Machines

Syaiful Hidayat[1], Mohammad Faris Al Fatih[2], Rizky Parlika[3]
[1]Software Engineering Study Program, Faculty of Science and Technology, Anwar Medika University
[2,3]Informatics, Faculty of Computer Science, University of Pembangunan Nasional "Veteran" Jawa Timur
[1]syaiful.hidayat@uam.ac.id*, [2]22081010277@student.upnjatim.ac.id, [3]rizkyparlika.if@upnjatim.ac.id

Abstract

This study uses Support Vector Regression (SVR) to forecast Bitcoin realized volatility across four timeframes (2h, 4h, 8h, and daily) and four forecast horizons using Binance OHLCV data from January 2023 to June 2026. Realized volatility is estimated using a 5-period rolling standard deviation of log returns. We construct 14 features from lagged volatility, returns, trading volume, RSI, MACD, Bollinger Bandwidth, and high-low spread. We evaluate linear and radial basis function (RBF) kernels using walk-forward validation with a 100-sample rolling window. To prevent look-ahead bias, the training window is adjusted for each forecast horizon so that only targets observable at the forecast time are included. Performance is evaluated using MAE, RMSE, MAPE, and R^2. The linear kernel outperforms the RBF kernel across all 16 timeframe-horizon combinations, achieving the highest R^2 of 0.5140 for daily volatility at the shortest horizon. Short horizons ($h = 1, 2$) generally produce positive R^2, whereas longer horizons ($h = 4, 12$) yield near-zero or negative R^2. The look-ahead bias correction produces only modest changes in performance, indicating that the initial bias had limited influence on the results. Permutation importance identifies lagged volatility as the dominant predictor, supporting volatility persistence. Overall, the findings demonstrate that linear SVR provides a simple, effective, and interpretable approach to Bitcoin volatility forecasting.

Keywords: bitcoin, volatility forecasting, SVR, realized volatility, cryptocurrency

1. Introduction

Over the past decade, crypto markets have grown explosively, with Bitcoin leading in market cap, trading volume, and institutional adoption. Bitcoin started life as an experimental niche digital currency in 2009 but has since become a mainstream financial instrument worldwide, attracting retail investors, institutional traders, and even sovereign wealth funds. The total market cap for cryptocurrencies is in the several-trillion-dollar range. By mid-2026, Bitcoin makes up a large chunk of that market cap. But one of the biggest headaches for market participants, risk managers, and policymakers alike remains Bitcoin's price volatility. Bitcoin is very volatile. Daily moves of 5-10% are common, and 20-30% drawdowns are usual. Hence, an accurate forecast of volatility is an important tool for informed decision-making [1].

Volatility is a statistical measure of the spread of returns for a given security or market index. It is an important input for option pricing models, portfolio optimization, value at risk (VaR), and risk-adjusted performance evaluation. However, classical econometric volatility models, such as Generalized Autoregressive Conditional Heteroskedasticity (GARCH) and its variants, are increasingly less able to capture the complex temporal patterns in cryptocurrency markets [2]. The Bitcoin price dynamics are non-stationary, non-linear, and highly turbulent. Shakourloo and Azimli [2] report regime-switching behavior in Bitcoin volatility, shifting from calm to turbulent phases, triggered by macroeconomic uncertainty, regulatory announcements, and changes in market sentiment. They used Markov-switching GARCH and hidden Markov copula models to show that Bitcoin's volatility pattern differs significantly from the standard assumption during crisis periods.

Forecasting cryptocurrency prices has been a subject of much research using machine and deep learning approaches. Ennagoura et al. [3] compared the SARIMAX, LSTM, and Facebook Prophet models to forecast the Ethereum price trend. Adding technical indicators such as RSI, MACD, and EMA improved signal quality. Results showed a clear trade-off between predictive accuracy and profitability. The SARIMAX model had the highest directional accuracy (75.00%). In this line of research, Kaygın et al. [4] provide convincing evidence that simple machine learning models can outperform complex deep architectures in the cryptocurrency prediction task. The authors compare SVM, LSTM, and GRU models on daily data for five major cryptocurrencies from October 2020 to September 2025. They found that the parsimonious SVM model always performed better out-of-sample, particularly for the very volatile assets. In addition, SHAP analysis [4] confirmed the role of immediate price-based variables and the weak explanatory power of many lagged technical indicators.

The selection of informative features has been discussed as a key factor in determining model performance for financial time series analysis [5]. Rakshitha K. et al. [5] proposed a feature selection framework based on mutual information (MI) and the Hilbert-Schmidt independence criterion (HSIC) for ranking 20 technical indicators for the NIFTY 50 index.

P-ISSN : 2089-3353
E-ISSN : 2808-9162

The most informative features found were the Ichimoku Conversion Line, WMA_{14}, and EMA_{14} for mutual information (MI) and the SuperTrend and Ichimoku Conversion Line for the Hilbert-Schmidt independence criterion (HSIC). The results underscore the importance of systematic feature engineering and selection in building robust volatility-forecasting models.

To improve the forecasting accuracy, several hybrid and ensemble approaches have been proposed that combine different model architectures. Oprea and Bâra [6] proposed an adaptive, regime-aware framework that combined ARIMAX, SARIMAX, and NeuralProphet models with an XGBoost meta-learner to identify hidden Markov model regimes, achieving an annual R^2 of 0.93. Maingo et al. [7] proposed a hybrid approach that combines the ARMA (3,2) - EGARCH (1,1) model with XGBoost as the feature for the volatility of the JSE Top40 Index. They found that the hybrid approach outperformed the standalone models across all accuracy metrics. Yan et al. [8] proposed an ensemble model with transfer learning and dynamic time warping as a flexible target-prediction framework and found that transfer learning performed better at predicting momentum indicators for American technology stocks. Bitcoin benefits from a similar approach [9].

The influence of external information sources on the predictability of technical indicators in cryptocurrency markets has been investigated. Ghosh et al. [10] used LightGBM and dynamic ensemble selection to explore the dynamic relationship between media coverage of political personalities and cryptocurrency movements. Using data from the RavenPack media tracker, the authors showed that Bitcoin is strongly dependent on media chatter about political personalities. Zhao et al. [11] proposed a hybrid model for the CSI 300 Index that combines wavelet denoising, a multi-head attention BiLSTM, ARIMA, and XGBoost. The ensemble method with inverse-error weights is much better than the individual models for all metrics.

The expanding literature on cryptocurrency forecasting highlights several gaps that require attention. Most studies focus on price prediction rather than volatility forecasting, which may be more important for risk management and derivative pricing. Second, most existing work uses a single time frame, usually daily data, and does not examine how time granularity affects model performance. Third, the systematic comparison of linear against non-linear SVR kernels for multiple timeframes and forecast horizons has not been extensively investigated in the context of Bitcoin volatility. This work fills these gaps by providing a detailed multi-horizon forecasting analysis of Bitcoin volatility using the SVR model across four time horizons [12].

The main aims of this work are (1) to construct realized volatility targets from the Bitcoin OHLCV data at 2-hour, 4-hour, 8-hour, and daily time scales; (2) to engineer a rich set of features from price- and volume-based technical indicators; (3) to evaluate the performance of linear and RBF SVR kernels with a walk-forward validation scheme for multiple forecast horizons; and (4) to provide empirical evidence of the best combination of time scales and horizons for Bitcoin volatility prediction.

2. Methodology

This study follows a systematic pipeline as summarized in Figure 1.

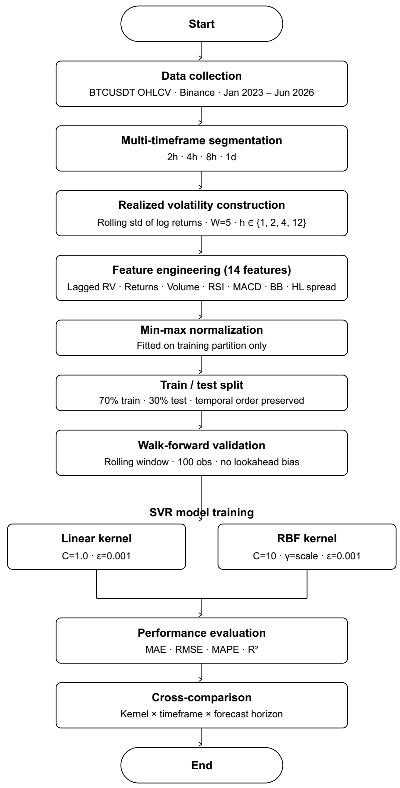

Figure 1. Proposed methodology flowchart

2.1. Data Description

The study utilizes historical OHLCV (Open, High, Low, Close, Volume) data for the Bitcoin/USDT pair from the cryptocurrency exchange Binance. The dataset spans Jan 1, 2023, through June 1, 2026, covering about 3.5 years of continuous trading data. We consider four temporal granularities: 2h, 4h, 8h, and 1d. The 2-hour dataset has 14,968 observations; the 4-hour dataset has 7,484 observations; the 8-hour dataset has 3,742 observations; and the daily dataset has 1,248 observations. The multi-timeframe approach allows us to investigate how temporal resolution affects volatility predictability. The high-frequency data capture intraday volatility patterns, while the low-frequency data provide a more stable signal at the expense of granularity.

P-ISSN : 2089-3353
E-ISSN : 2808-9162

2.2. Volatility Target Construction

Realized volatility (RV) is the standard deviation of log returns over a rolling window of 5 consecutive periods for each time frame. Formally, for a price series $\{P_t\}$, the log return is defined as $r_t = \ln(P_t/P_{t-1})$, and the realized volatility at time t is computed as:

$$\mathrm{RV_t} = \sqrt{\frac{1}{W-1}\sum_{i=t-W+1}^{t}(r_i - \overline{r_t})^2} \quad (1)$$

where $W = 5$ is the window size and r̄_t is the mean log return in the window. The forecast target for horizon h is defined as $\mathrm{RV_t} + h$, the realized volatility h periods ahead. Four forecast horizons, h, are considered, $h \in \{1, 2, 4, 12\}$, corresponding to short-term (next and two-period ahead), medium-term (four-period ahead), and long-term (twelve-period ahead) forecasting scenarios.

2.3. Feature Engineering

Fourteen features are built from raw data (open, high, low, close, and volume data) to be used as inputs to the SVR model. The feature set is designed to capture both the time persistence of volatility and the broader market dynamics that could affect future volatility levels.

Lagged volatility features: To exploit the well-known persistence property of financial volatility, i.e., that current levels of volatility can be partially predicted by recent historical volatility, we add five lagged values of the realized volatility ($\mathrm{RV_{t-1}}$, $\mathrm{RV_{t-2}}$, $\mathrm{RV_{t-3}}$, $\mathrm{RV_{t-5}}$, $\mathrm{RV_{t-10}}$) The autoregressive form is inspired by the GARCH family of models, which also includes lagged variance terms.

We include three lagged log returns ($r_{t-1}, r_{t-2}, r_{t-3}$) to capture short-term price momentum, as sharp price moves tend to be followed by volatility regime shifts.

Lagged features of volume: We add two lagged volume variables, (V_{t-1}, V_{t-2}) Research has demonstrated a positive relationship between trading volume and volatility, particularly during periods of market distress or significant price discovery events.

Technical Indicators: The 14-day Relative Strength Index (RSI) determines overbought and oversold levels that could indicate a trend change. The MACD signal line (a 9-period signal line calculated from the 12-period and 26-period exponential moving averages) indicates momentum and trend direction. Bollinger Bandwidth (20 periods, 2 standard deviations). A measure of volatility expansion and contraction. The last one is the high-low spread $\frac{\mathrm{High_t - Low_t}}{\mathrm{Close_t}}$ This is the range normalized by the closing price during the day. It shows the magnitude of price movements in each period.

2.4. Data Normalization

The min-max scaling scales each feature to a range between [0, 1] and normalizes all the features:

$$x_{\mathrm{norm}} = \frac{x - x_{min}}{x_{max} - x_{min}} \quad (2)$$

The scaler is fitted only on the first training partition to avoid information leakage from the test set, which is crucial to preserve the integrity of out-of-sample evaluation.

2.5. Support Vector Regression

The prediction model used is Support Vector Regression (SVR) with the ε-insensitive loss function. Given a training set $\{(x_i, y_i)\}_{i=1}^{N}$, SVR seeks a function $f(x) = \langle w, \phi(x)\rangle + b$ that deviates by at most ε from the observed targets, while remaining as flat as possible. The optimization problem is written as:

$$\min_{w,b,\xi,\xi^*} \frac{1}{2}|w|^2 + C\sum_{i=1}^{N}(\xi_i + \xi_i^*) \quad (3)$$

with the following restrictions:

$$y_i - f(x_i) \le \varepsilon + \xi_i, f(x_i) - y_i \le \varepsilon + \xi_i^*,$$
$$\xi_i, \xi_i^* \ge 0 \quad (4)$$

The two kernel configurations evaluated in this study are the following: The linear kernel $K(x, x') = \langle x, x'\rangle$ directly maps the input into a linear decision function, capturing additive relationships between features and the target.

The RBF kernel:

$$K(x, x') = exp\left(-\gamma\left||x - x'|\right|^2\right) \quad (5)$$

maps the input into an infinite-dimensional feature space, allowing it to represent non-linear relationships. The hyperparameters are set as follows; $C = 1.0$ and $\varepsilon = 0.001$ for the linear kernel; $C = 10.0, \gamma = scale$, and $\varepsilon = 0.001$ for the RBF kernel.

2.6. Walk-Forward Validation and Bias Correction

To evaluate out-of-sample performance while preserving temporal ordering, we employ a walk-forward validation scheme in which the model is re-estimated at each forecast origin i using a 100-observation training window and then predicts the target RV[i + h]. To prevent look-ahead bias in multi-horizon forecasting, training samples are restricted to those whose targets are already observable at time i, satisfying t + h ≤ i; thus, the training window is defined as [max(0, i − h − W + 1), i − h]. This excludes 0, 1, 3, and 11 observations for h = 1, 2, 4, and 12, respectively, ensuring that no future information enters the training process. The correction has minimal effects on shorter horizons: R^2 shifts remain within ±0.005 for h = 2 and reach 0.017 for h = 4, while the largest shift for h = 12 is 0.13. Importantly, the best-performing configuration, 1-day RV with h = 1 and a linear kernel ($R^2 = 0.5140$), is unaffected. These results indicate that the original findings are not materially driven by look-ahead bias, while the corrected procedure strengthens the validity of the reported performance rankings.

P-ISSN : 2089-3353
E-ISSN : 2808-9162

2.7. Evaluation Metrics

We employ four standard evaluation metrics to evaluate the model performance comprehensively. The Mean Absolute Error (MAE) is the mean of the absolute differences between predictions and actual values:

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{6}$$

The Root Mean Squared Error (RMSE) penalizes larger errors more than smaller errors:

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

MAPE (Mean Absolute Percentage Error) expresses the error as a percentage of the actual value:

$$\mathrm{MAPE} = \frac{100\%}{n}\sum_{i=1}^{n}\left|\frac{y_i-\hat{y}_i}{y_i}\right| \tag{8}$$

The coefficient of determination (R^2) measures the proportion of variance of the target explained by the model:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i-\hat{y}_i)^2}{\sum_{i=1}^{n}(y_i-\bar{y})^2} \tag{9}$$

3. Results and Discussion

3.1. Bitcoin Price Overview

Figure 2 displays the daily closing price of Bitcoin (BTCUSDT) during the study period, which spans from January 2023 to June 2026. The price pattern illustrates the typical volatility of crypto markets, with long periods of gradual appreciation interspersed with steep corrections and quick recoveries. This forecasting study is compelling because of the high volatility and appreciation, e.g., from a minimum of roughly \$16,500 in early January 2023 to peaks above \$100,000 in the subsequent months.

Figure 2. BTC/USDT Daily Closing Price

3.2. Evaluation Metrics Summary

Table 1 presents evaluation results across all timeframes, forecast horizons, and SVR kernel configurations, revealing consistent patterns.

Table 1. SVR Volatility Forecasting Performance

TF	H	K	MAE	$RMSE$	$MAPE$	R^2
2h	1	Lin	0.00168	0.002514	42.35	0.4888
2h	1	RBF	0.002387	0.003704	56.74	-0.1095
2h	2	Lin	0.001989	0.002904	50.23	0.3175
2h	2	RBF	0.003335	0.005196	82.54	-1.1844
2h	4	Lin	0.002816	0.003976	76.35	-0.279
2h	4	RBF	0.00469	0.006934	127.88	-2.8895
2h	12	Lin	0.003129	0.004327	86.49	-0.5155
2h	12	RBF	0.005379	0.008283	154.74	-4.5536
4h	1	Lin	0.002282	0.003268	39.47	0.4435
4h	1	RBF	0.003272	0.004808	54.33	-0.2047
4h	2	Lin	0.002735	0.003784	48.98	0.254
4h	2	RBF	0.004257	0.005975	77.79	-0.8601
4h	4	Lin	0.00369	0.004927	70.5	-0.2643
4h	4	RBF	0.00616	0.008699	118.19	-2.9416
4h	12	Lin	0.004053	0.005397	75.36	-0.517
4h	12	RBF	0.00663	0.009252	132.54	-3.458
8h	1	Lin	0.003265	0.004522	40.72	0.4571
8h	1	RBF	0.004103	0.005531	50.09	0.1876
8h	2	Lin	0.003811	0.005216	47.85	0.2783
8h	2	RBF	0.005655	0.007589	71.21	-0.528
8h	4	Lin	0.005057	0.006743	67.21	-0.2063
8h	4	RBF	0.008139	0.01085	106.91	-2.123
8h	12	Lin	0.005948	0.008229	79.1	-0.7922
8h	12	RBF	0.008913	0.011903	116.8	-2.7492
1d	1	Lin	0.005594	0.008228	34.1	0.514
1d	1	RBF	0.007284	0.010877	42.98	0.1506
1d	2	Lin	0.006927	0.010069	42.31	0.2734
1d	2	RBF	0.009845	0.013542	60.4	-0.3141
1d	4	Lin	0.00927	0.013373	55.54	-0.2814
1d	4	RBF	0.013551	0.018557	85.12	-1.4676
1d	12	Lin	0.00988	0.014453	61.14	-0.4907
1d	12	RBF	0.016835	0.024747	100.66	-3.3702

3.3. Kernel Comparison

An intriguing and robust result is that the linear kernel outperforms the RBF kernel in all 16 combinations of timeframe and horizon. The linear kernel obtains the lowest RMSE in each combination, with an improvement from moderate (e.g., 18% RMSE reduction for the 8h/h=1 case) to substantial (e.g., 48% RMSE reduction for the 2h/h=12 case). This result is consistent with Kaygın et al. [4], which showed that simpler models are often more effective than complex architectures for forecasting cryptocurrencies due to the inherent noise and non-stationarity of digital asset markets. In many cases, the RBF kernel yields negative R^2 values, implying that the model's predictions are worse than simply predicting the mean observed volatility. This degradation is even more prominent at longer horizons, where the R^2 of the RBF kernel reaches -4.5536 for the 2h/h=12 combination. This analysis considers the RBF kernel's overfitting behavior when applied to noisy financial data.

3.4. Horizon Effect

The forecast horizon affects accuracy greatly. The h=1 forecast outperforms others, with R^2: 0.4888 (2h), 0.4435 (4h), 0.4571 (8h), and 0.5140 (1d). The two-period horizon (h=2) shows moderate accuracy: 0.3175, 0.2540, 0.2783, and 0.2734, but accuracy drops sharply for h=4 and h=12, with negative R^2 values

P-ISSN : 2089-3353
E-ISSN : 2808-9162

indicating unreliable predictions beyond the short term. The h=2 R^2 values are slightly lower than those in the initial uncorrected analysis due to the removal of look-ahead bias, though the shifts remain within ±0.005, confirming minimal initial bias. The R^2 at h=4 and h=12 shifted slightly, especially at h=12 with changes up to 0.13, due to a shorter training window. Results for h=1 stay the same. These results align with the efficient market hypothesis, showing increased uncertainty in longer horizons of non-stationary data. Volatility persistence from lagged RV features predicts the next period well, but this signal weakens as the horizon lengthens.

3.5 Time Frame Analysis

The best R^2 value is on the daily (1d) timeframe (0.5140 at h=1), showing daily granularity offers the best balance between signal and noise for Bitcoin volatility prediction. Similar patterns appear in intraday timeframes (2h, 4h, 8h): strong at h=1, moderate at h=2, negligible at h≥4. The 8h timeframe has a slightly higher R^2 at h=2 (0.2783) than 4h (0.2540), and 2h has the highest R^2 at h=1 (0.4888). These results account for look-ahead bias after correction. The linear kernel, which performs best, produces Figure 3, showing actual vs. predicted volatility (h=1) for daily data. It captures broad volatility movements well, including levels and fluctuations, but underestimates peaks and smooths extreme spikes.

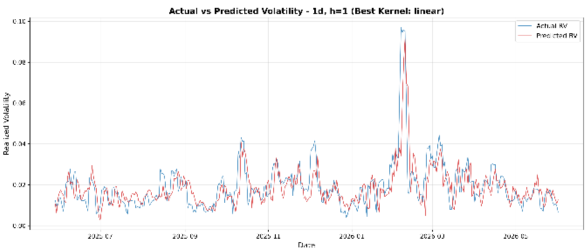

Figure 3. Actual and Predicted Realized Volatility

3.6 RMSE Comparison

Figure 4 presents the $RMSE$ for all combinations of time frame and horizon for the best-performing kernel (linear in all cases). The bar chart clearly shows the two effects of time frame and horizon on forecast accuracy. $RMSE$ increases monotonically with the coarseness of the time frame (the absolute magnitude of volatility is larger at lower frequencies) and with the length of the forecast horizon (prediction uncertainty increases).

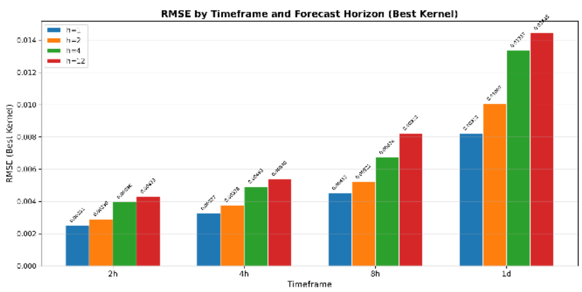

Figure 4. RMSE Across Timeframes and Horizons

3.7 Response Surface R^2

The R^2 heatmaps for all timeframe-horizon combinations are shown in Figure 5. The heatmap provides a more intuitive view of the performance landscape, where darker shading corresponds to higher explanatory power. The best performance is in the top-left corner ($h = 1$, all timeframes) and the worst in the lower right corner ($h = 12$, longer timeframes). Several combinations have performance below zero.

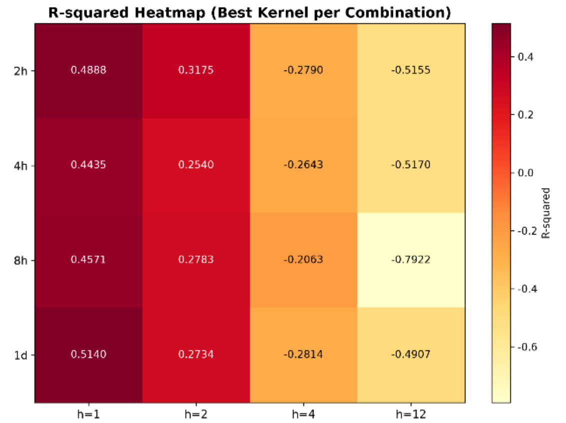

Figure 5. R^2 Across Timeframes and Horizons

3.8. Analysis of Prediction Accuracy

Figure 6 presents the plot of actual vs. predicted volatility for the daily time frame at $h = 1$ with a linear kernel. The points cluster around the diagonal line, especially in the low-to-moderate volatility range, confirming the model's ability to capture the central tendency of volatility dynamics. However, the scatter increases at higher volatility levels, indicating lower accuracy in extreme market conditions. This heteroscedastic behavior in prediction errors is characteristic of volatility forecasting models and suggests directions for improvement through heteroscedasticity-aware extensions.

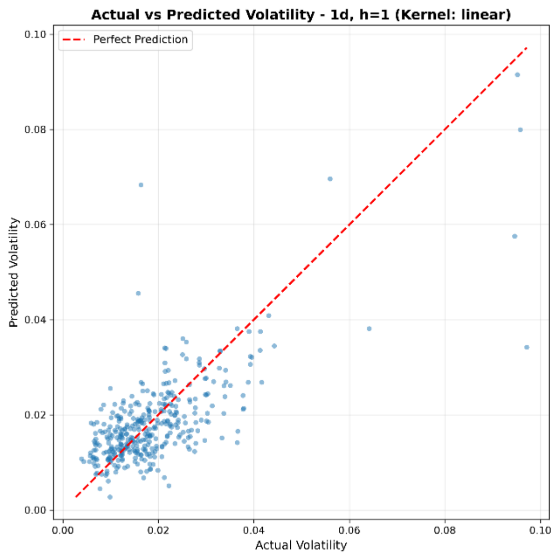

Figure 6. Actual Versus Predicted Volatility

P-ISSN : 2089-3353
E-ISSN : 2808-9162

3.9 Analyzing Feature Importance

Figure 7 displays permutation feature importance for the daily timeframe at h=1 for the RBF SVR model. This metric measures each feature's contribution by the performance drop when values are shuffled. Lagged realized volatility, especially RV_lag1 and RV_lag2, are the top predictors, supporting the volatility persistence hypothesis. Technical indicators like MACD signal and Bollinger Bandwidth have moderate effects, while RSI is less important. Lagged log return and volume also contribute, indicating price momentum and trading activity add predictive value beyond volatility. Figure 8 shows similar importance rankings for the linear SVR kernel, with RV_lag1 and RV_lag2 remaining dominant. The linear model depends more on recent volatility, with RV_lag1 having a larger share of importance than in the RBF kernel, implying less reliance on secondary feature interactions. Technical indicators and lagged returns contribute modestly and consistently across kernels, confirming that short-term volatility autoregression is the key predictor of Bitcoin realized volatility.

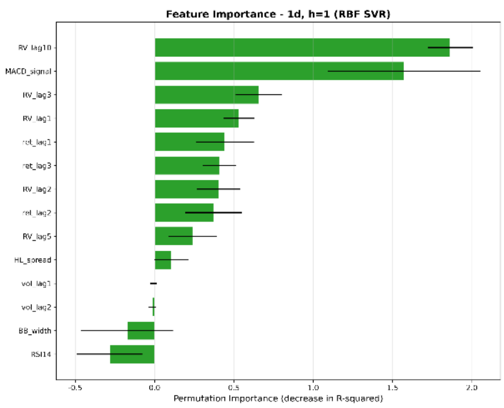

Figure 7. SVR Feature Importance Rankings (RBF)

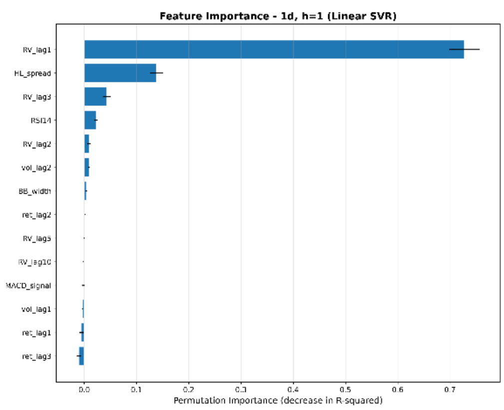

Figure 8. SVR Feature Importance Rankings (Linear)

4. Conclusion

Results show that the linear SVR kernel consistently outperformed the RBF kernel across all timeframes and forecast horizons, with the highest (R^2) of 0.5140 on the daily timeframe at (h=1). Forecast horizon strongly affected performance, with (h=1) and (h=2) maintaining predictive power, but near-zero or negative (R^2) at (h=4) and (h=12). The daily timeframe was the most reliable, while intraday offered finer resolution but more noise. Lagged realized volatility remained the key predictor, confirming volatility's dominance. These findings suggest daily linear SVR is useful for one-day-ahead volatility, with shorter intraday models providing signals for high-frequency uses. Complex nonlinear kernels are unnecessary for this forecasting task. Future research could combine linear SVR with regime-switching and use on-chain and sentiment data to extend prediction beyond (h=2), testing across other cryptocurrencies and market cycles.

References

[1] C. Fieberg, G. Liedtke, T. Poddig, T. Walker, and A. Zaremba, “A Trend Factor for the Cross Section of Cryptocurrency Returns,” *J. Financ. Quant. Anal.*, vol. 60, no. 7, pp. 3116–3153, Jul. 2025, doi: 10.1017/S0022109024000747.

[2] A. Shakourloo and A. Azimli, “Regime-switching in bitcoin volatility under global uncertainty: Markov-switching GARCH and hidden Markov Copula approaches,” *Res. Int. Bus. Financ.*, vol. 83, p. 103295, 2026, doi: 10.1016/j.ribaf.2026.103295.

[3] D. Ennagoura *et al.*, “A Comparative Evaluation of SARIMAX, LSTM, and Prophet Models for Cryptocurrency Price Trend Prediction,” *Eng. Technol. Appl. Sci. Res.*, vol. 16, no. 2, pp. 33146–33151, Apr. 2026, doi: 10.48084/etasr.15554.

[4] C. Y. Kaygın, M. Gün, O. N. Akarsu, H. Bağcı, and A. Yanık, “Algorithmic Stability in Turbulent Markets: Unveiling the Superiority of Shallow Learning over Deep Architectures in Cryptocurrency Forecasting,” *Mathematics*, vol. 14, no. 6, p. 989, Mar. 2026, doi: 10.3390/math14060989.

[5] K. Rakshitha, N. Naik, and S. Shetty, “Feature Selection of Technical Indicators Using Mutual Information and HSIC for Financial Time Series Analysis,” *Int. J. Math. Comput. Sci.*, vol. 21, no. 1, pp. 103–107, Jan. 2026, doi: 10.69793/ijmcs/01.2026/naik.

[6] S. V. Oprea and A. Bâra, “Regime-Aware Adaptive Forecasting Framework for Bitcoin Prices Using Probabilistic Generative Models,” *Comput. Econ.*, 2026, doi: 10.1007/s10614-026-11338-3.

[7] I. Maingo, T. Ravele, and C. Sigauke, “A Fusion of Statistical and Machine Learning Methods: GARCH-XGBoost for Improved Volatility Modelling of the JSE Top40 Index,” *Int. J. Financ. Stud.*, vol. 13, no. 3, p. 155, Aug. 2025, doi: 10.3390/ijfs13030155.

[8] K. Yan *et al.*, “Flexible Target Prediction for Quantitative Trading in the American Stock Market: A Hybrid Framework Integrating Ensemble Models, Fusion Models and Transfer Learning,” *Entropy*, vol. 28, no. 1, p. 84, Jan. 2026, doi: 10.3390/e28010084.

[9] E. Radmand, J. Pirgazi, and A. G. Sorkhi, “A Hybrid TLBO–XGBoost Model With Novel Labeling for Bitcoin Price Prediction,” *Int. J. Intell. Syst.*, vol. 2025, no. 1, Jan. 2025, doi: 10.1155/int/6674437.

[10] I. Ghosh, E. Alfaro-Cortés, M. Gámez, and N. García-Rubio, “When political leaders speak, market reacts: Unveiling the dynamic nexus of media chatter and crypto movements,” *Borsa Istanbul Rev.*, vol. 26, no. 3, p. 100786, Jan. 2026, doi: 10.1016/j.bir.2026.100786.

[11] Q. Zhao, H. Li, X. Liu, and Y. Wang, “A Hybrid Model of Multi-Head Attention Enhanced BiLSTM, ARIMA, and XGBoost for Stock Price Forecasting Based on Wavelet Denoising,” *Mathematics*, vol. 13, no. 16, p. 2622, Aug. 2025, doi: 10.3390/math13162622.

[12] T. Hall and K. Rasheed, “A Survey of Machine Learning Methods for Time Series Prediction,” *Appl. Sci.*, vol. 15, no. 11, p. 5957, May 2025, doi: 10.3390/app15115957.