

Evaluation of the Performance of Deep Learning Models in Cryptocurrency Price Prediction: A Case Study of Bitcoin, Dogecoin, Ethereum, and Ripple

Reza Taleblou¹ | Parisa Mohajeri^{2*}

Article Type:
Research Article

Reza Taleblou

Associate Professor of Economics,
Department of Theoretical Economics,
Allameh Tabataba'i University, Tehran,
Iran.
E-mail: talebloo.r@gmail.com

Parisa Mohajeri

Corresponding Author, Associate
Professor of Economics, Department of
Theoretical Economics, Allameh
Tabataba'i University, Tehran, Iran.
E-mail: parisa_m2369@yahoo.com

Spring & Summer (2025) 2(1): 7-19

Received: 23 February 2025

Revised: 15 March 2025

Accepted: 7 April 2025

Available Online: 14 April 2025

ABSTRACT

Cryptocurrencies, as one of the emerging asset classes, have gained significant popularity in recent years. Accurate forecasting of cryptocurrencies' prices has become highly attractive for both researchers and investors due to their volatile and non-linear price behavior. However, predicting the cryptocurrencies' prices accurately remains challenging due to their substantial fluctuations and complex dynamics. Research findings indicated that the methods of deep learning and neural networks outperform traditional econometric approaches in forecasting financial and economic time series. Among the techniques of neural network and deep learning, various types of Recurrent Neural Network (RNN) models have been proven to be effective. This study employed three Recurrent Neural Network architectures—RNN, Long Short-Term Memory (LSTM), and Gated Recurrent Unit (GRU)—to predict the logarithm of the prices of four major cryptocurrencies of Bitcoin (BTC), Dogecoin (DOGE), Ethereum (ETH), and Ripple (XRP). Daily time-series data from January 17, 2018, to December 18, 2024, were utilized for this purpose. The data were collected using the `cryptocmd` python package. The experimental results which were assessed using four metrics—Root Mean Square Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and Symmetric Mean Absolute Percentage Error (SMAPE)—revealed two key findings: First, as the forecasting horizon increases, the required input size for achieving the best predictions increases for all models. Second, the LSTM model demonstrates a superior performance in predicting the prices of major cryptocurrencies for 1-day and 30-day horizons, whereas the GRU model exhibits the lowest prediction error for a 7-day horizon. These findings provided valuable insights for estimating the mean equation, which is instrumental in forecasting the expected returns of cryptocurrency assets for risk management purposes.

KEYWORDS

Artificial Intelligence, Cryptocurrencies, Deep Learning, Gated Recurrent Units, Long Short-Term Memory, Recurrent Neural Networks.

Cite this article: Taleblou, R., & Mohajeri, P. (2025). Evaluation of the Performance of Deep Learning Models in Cryptocurrency Price Prediction: A Case Study of Bitcoin, Dogecoin, Ethereum, and Ripple. *Journal of Knowledge Economy Studies (JKES)*, 2(1), 7-19.

DOI: <http://doi.org/10.22034/kes.2025.2054189.1046>



Authors retain the copyright and full publishing rights.

Online ISSN: 3060-7329

Publisher: Hazrat-e Masoumeh University.

Introduction

Cryptocurrencies are digital assets that have gained significant attention in recent years. By nature, they are decentralized, meaning their issuance and value are not controlled by any central authority or institution. Consequently, their prices are highly volatile, making the prediction of future prices a challenging and complex task. Accurate forecasting of cryptocurrencies' price is crucial for informed decision-making and reducing the risks associated with these assets for investors, traders, and financial analysts (Kubat, 2015). Various methods exist for predicting the cryptocurrencies' prices, including time-series analysis, sentiment analysis, and neural networks and deep learning algorithms. However, due to the dynamic and complex nature of cryptocurrencies, these methods face limitations in terms of their predictive accuracy.

Time series forecasting plays a pivotal role in the cryptocurrency market, where accurate predictions of asset prices and volatility can significantly influence investment strategies, risk mitigation, and liquidity management. Traditional statistical and econometric models, such as ARIMA and GARCH, have historically been employed for financial forecasting. However, their reliance on assumptions like linearity, stationarity, and normality—rarely satisfied in cryptocurrency markets—limits their effectiveness. Cryptocurrencies exhibit extreme volatility, non-linear price dynamics, and susceptibility to exogenous shocks rendering conventional methods inadequate.

In recent years, machine learning—particularly deep learning architectures—has emerged as a transformative tool for cryptocurrency forecasting. Advanced neural network techniques, such as Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs), have gained significant traction due to their ability to capture complex temporal dependencies, abrupt trend reversals, and multi-scale patterns inherent in crypto data. These models, as part of the broader family of Artificial Neural Networks (ANNs), dynamically adapt to market anomalies like pump-and-dump cycles or flash crashes. By overcoming the limitations of traditional statistical methods, deep learning enables robust predictions across diverse horizons—from short-term intraday fluctuations to long-term trend analysis—positioning it as a cornerstone of modern cryptocurrency analytics. This growing trend is supported by extensive international research demonstrating the efficacy of LSTM and GRU in modeling long-term dependencies and nonlinear relationships, particularly in volatile asset markets like cryptocurrencies.

Despite extensive international research employing various neural networks and deep learning models to predict economic variables, there is a scarcity of published studies in this domain in an Iranian context. Moreover, apart from the prediction of Bitcoin price in the study by Sayadi Nezhad et al. (2023), the prediction of other major cryptocurrencies has not been addressed in prior Iranian research studies. In light of the aforementioned considerations, the primary objective of this paper is to predict the logarithm of the prices of four major cryptocurrencies using three models—RNN, LSTM, and GRU—and to identify the superior model with the lowest prediction error, utilizing daily price data

from January 17, 2018, to December 18, 2024. This study is innovative in several respects: First, in addition to using the widely used RNN and LSTM methods, the GRU method is also employed. Second, the prices of four key cryptocurrencies are predicted, whereas previous studies have predominantly focused on Bitcoin. Third, approximately 15% of the last observations from the end of the period (around 360 observations) are reserved for testing the prediction accuracy, followed by the calculation of prediction errors for 1-day, 7-day, and 30-day horizons. In this paper, the data under examination were analyzed within the framework of univariate time series analysis, addressing the key limitations inherent in Recurrent Neural Network (RNN) family architectures. Additionally, lagged variables and other deterministic variables, such as days and months of the year, were also incorporated as explanatory variables in the analysis.

Literature Review

The prediction of financial time series primarily focuses on forecasting the asset prices (Tsay, 2005). Despite the existence of various methodologies, recent research has concentrated on using deep learning models to predict future movements in the assets (Fischer & Krauss, 2018). This encompasses a wide range of topics, including stock price prediction, index forecasting, foreign exchange rates, commodity prices (e.g., oil and gold), bond prices, and cryptocurrency prices (Sezer et al., 2020).

The value of cryptocurrencies theoretically reflects their desirability as a medium of exchange, which has significantly increased over the past decade. Given the growing importance of cryptocurrencies for financial systems, early studies focused on their volatility, confirming that these fluctuations are substantial (Klein et al., 2018) and extremely challenging to predict (Fang et al., 2020; Walther et al., 2019). Furthermore, some empirical studies have highlighted several key characteristics, such as the "fat-tailed distribution of cryptocurrency returns", the "weakening of absolute and relative autocorrelation of returns at different rates", the "strong presence of leverage effects and volatility clustering", "long-term dependencies in volatility and returns", and "compliance with power-law dynamics in prices and volatility" (Zhang et al., 2018). These features make cryptocurrency price prediction difficult and investing in them riskier than traditional financial assets. Additionally, empirical findings indicate that Bitcoin and other cryptocurrencies exhibit cyclical patterns in recent years (Dong et al., 2022; Kyriazis et al., 2020).

Most studies on cryptocurrency price prediction have relied on traditional statistical methods. For instance, Catania et al. (2019) utilized a vector autoregression model to forecast Bitcoin, Ripple, Litecoin, and Ethereum prices, while Conrad et al. (2018) employed the GARCH-MIDAS model to study the impact of mid- and long-term S&P 500 volatility on Bitcoin. Similarly, Walther et al. (2019) used the GARCH-MIDAS model to predict volatility for five major cryptocurrencies.

In recent years, researchers' attention has shifted toward deep learning models for cryptocurrency price prediction. A notable contribution in this area is the work of

Kristjanpoller and Minutolo (2018), who introduced the MLP-GARCH model for predicting the volatility of Bitcoin price. Artificial neural networks represent a collection of advanced techniques and computational methods in neural networks and deep learning. These networks consist of structures made up of interconnected simple processing units capable of performing highly parallel computations to process data efficiently (Goodfellow et al., 2016). Nakano et al. (2018) also applied Multi-Layer Perceptron (MLP) models to predict Bitcoin returns. Empirical evidence suggests that this approach outperforms ARIMA, Prophet, and Random Forest models in forecasting directional movements (Ibrahim et al., 2021).

Artificial neural networks offer numerous advantages over other methods, such as their ability to use nonlinear functions, perform parallel computations, learn from data, and incorporate qualitative information. Using nonlinear functions improves the evaluation of input data and reduces the time required for assessment due to their parallel computation capabilities (Limsombunchai et al., 2004).

Recently, recurrent neural networks (RNNs) like LSTM and GRU have been employed to automatically extract short-term patterns from high-frequency cryptocurrency time series. Notable examples include studies by Patel et al. (2020) and Parekh et al. (2022), both of which utilized LSTM and GRU models. These advanced deep learning-based neural networks have effectively learned complex sequential patterns in data, explaining why LSTM has outperformed the traditional neural network models like MLP in recent studies (Chen et al., 2021; Lahmiri & Bekiros, 2019; Li & Dai, 2020). Moreover, according to Zhang et al. (2021), the GRU model demonstrated a superior predictive performance for four major cryptocurrencies compared to both traditional neural network methods and LSTM-based models. Seabe, et al. (2023) used LSTM, GRU, and Bi-Directional LSTM for forecasting the cryptocurrency prices. Hamayel et al. (2021) used GRU, LSTM and bi-LSTM machine learning algorithms for predicting the cryptocurrency prices.

Methodology

Research Method

1. Simple Recurrent Neural Networks (RNN)

Recurrent Neural Networks (RNNs) have been one of the cornerstones of deep learning research for several decades. The concept of RNN was first introduced in the 1980s by David Rumelhart, Geoffrey Hinton, and Ronald Williams, who proposed the use of recurrent connections to model temporal relationships in data. The RNN models are suitable for multivariate time-series data and are capable of capturing temporal dependencies across different time periods (Che et al., 2018; Graves, 2012). In an RNN, input values enter the network, and output values are generated based on these inputs and network parameters.

The backpropagation algorithm can be applied to train a recurrent network by unfolding it over time and constraining certain connections to maintain consistent weights across all time steps (Xu et al., 2018). Thus, RNNs can encode previous

information into the current hidden layer during the learning process, enabling effective learning of time-series data (Bengio et al., 1994).

In simple recurrent neural networks (RNNs), the output y_t is calculated as follows:

$$y_t = f(W_y h_t) \quad (1)$$

$$h_t = \sigma(W_h h_{t-1} + W_x x_t) \quad (2)$$

Here, W_y , W_h , and W_x represent matrices for the hidden layer output h_t , the activity of the previous hidden layer h_{t-1} , and the input x_t , respectively. The recurrence in Equation (2) links the current hidden layer activity h_t with the previous hidden layer activity h_{t-1} . This dependency is nonlinear due to using sigmoid function $\sigma(\cdot)$ (Yao et al., 2015).

2. Long Short-Term Memory (LSTM)

Theoretically, RNNs should possess the ability to learn sequences of any complexity. However, empirical observations indicate that these networks struggle to retain relevant information from past inputs over long durations. This inability to store and preserve past information leads to instability when generating sequences (Hansson & Rostami, 2019). Addressing this instability requires longer-term memory capabilities, as a network with extended memory can still learn from inputs despite long intervals between significant events. This need led to the introduction of the Long Short-Term Memory (LSTM) model, a subtype of RNN with enhanced storage and access capabilities compared to traditional versions. Consequently, LSTMs can detect important features in input sequences early on and propagate them along a long path while maintaining long-term dependencies throughout the process (Chen et al., 2017).

Hochreiter and Schmidhuber (1997) were the first researchers who introduced the LSTM model, which was later refined in Graves' research study (2012). Designed as a primary alternative to RNNs, LSTMs address issues such as vanishing gradients and facilitating more efficient computations. In many recent studies, LSTMs have demonstrated a superior performance and are relatively easier to train. As a result, LSTMs have become baseline architectures for processing sequential data with temporal information (Jozefowicz et al., 2015).

An LSTM unit consists of a cell and three gates: the input gate, forget gate, and output gate (Cho et al., 2014). The cell serves as the memory component of the LSTM, used to store values over arbitrary time intervals. In LSTMs, a linear dependency is introduced between the memory cells c_t and their previous state c_{t-1} . Additionally, LSTMs include input and output gates, which apply nonlinear functions to the input and output of the LSTM, respectively. The relationships within an LSTM are expressed as follows:

$$i_t = \sigma(W_{xi}x_t + W_{hi}h_{t-1} + W_{ci}c_{t-1}) \quad (3)$$

$$f_t = \sigma(W_{xf}x_t + W_{hf}h_{t-1} + W_{cf}c_{t-1}) \quad (4)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot \tanh(W_{xc}x_t + W_{hc}h_{t-1}) \quad (5)$$

$$o_t = \sigma(W_{xo}x_t + W_{ho}h_{t-1} + W_{co}c_t) \quad (6)$$

$$h_t = o_t \odot \tanh(c_t) \quad (7)$$

Here i_t , f_t , and o_t represent the input gate, forget gate, and output gate, respectively (Yao et al., 2015).

3. Gated Recurrent Unit (GRU)

The Gated Recurrent Unit (GRU), introduced by [Cho et al. \(2014\)](#), shares many similarities with LSTM but differs in how the gates are applied. Unlike LSTM, which uses three gates, GRU employs only two gates: reset gate and update gate. The reset gate determines how to combine new input with previous memory and controls the degree to which the prior memory is ignored. Smaller values for the reset gate indicate that the previous memory is largely disregarded. The update gate specifies the extent to which the prior memory is retained ([Song, 2018](#)). A key distinction between GRU and LSTM is that GRU does not have a separate memory cell; instead, its operations are simplified as follows:

$$h_t = (1 - z_t)h_{t-1} + z_t\tilde{h}_t \quad (8)$$

$$z_t = \sigma(W_z x_t + U_z h_{t-1}) \quad (9)$$

$$\tilde{h}_t = \tanh(W_h x_t + U(r_t \odot h_{t-1})) \quad (10)$$

$$r_t = \sigma(W_r x_t + U_r h_{t-1}) \quad (11)$$

Here, h_t represents the output of the GRU at time t , while z_t and r_t denote the update and reset gates, respectively. W_z , W_h , W_r , U_z , and U_r are weight matrices within the GRU architecture (Yao et al., 2015). The GRU's streamlined structure reduces computational complexity while maintaining strong performance in capturing temporal dependencies.

Time-Series Cross-Validation

Time-series cross-validation is a method used to assess how well a model performs on historical data. Unlike standard cross-validation, which randomly partitions the dataset, time-series cross-validation preserves the chronological order of the data by sliding a window over past observations and predicting the subsequent periods. This approach allows for a more accurate estimation of a model's predictive capabilities across multiple forecast horizons. When only a single window is used, it resembles a standard train-test split, where the test set consists of the most recent observations, and the training set includes all the prior data. Given the study period spans from January 17, 2018, to December 18, 2024, a total of 2,524 daily observations (approximately 7 years) are available. The first 6 years (85% of the data) are allocated for training, while the final year (15% of the data) is reserved for testing.

To define the forecast horizon based on data frequency, the present study utilizes a daily data frequency. For a one-week forecast, $h=7$ days, and for a one-month forecast, $h=30$ days. Another parameter, the step size for advancing the data window, is set to 7 days (one week) or 30 days (one month) depending on the specific forecast horizon. Additionally, the number of prediction windows is fixed at 360 to ensure carrying out a comprehensive evaluation across the dataset.

Model Validation

To evaluate the predictive performance of the models, this study employs four metrics of Root Mean Square Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and Symmetric Mean Absolute Percentage Error (SMAPE). These metrics quantify the deviation between predicted and actual values, with lower

values indicating a higher predictive accuracy. The mathematical formulations of these metrics are provided below:

$$RMSE = \sqrt{\frac{1}{H} \sum_{\tau=t+1}^{t+H} (y_{\tau} - \hat{y}_{\tau})^2} \quad (12)$$

$$MAE(y_{\tau}, \hat{y}_{\tau}) = \frac{1}{H} \sum_{\tau=t+1}^{t+H} |y_{\tau} - \hat{y}_{\tau}| \quad (13)$$

$$MAPE(y_{\tau}, \hat{y}_{\tau}) = \frac{1}{H} \sum_{\tau=t+1}^{t+H} \frac{|y_{\tau} - \hat{y}_{\tau}|}{y_{\tau}} \quad (14)$$

$$SMAPE(y_{\tau}, \hat{y}_{\tau}) = \frac{1}{H} \sum_{\tau=t+1}^{t+H} \frac{|y_{\tau} - \hat{y}_{\tau}|}{|y_{\tau}| + |\hat{y}_{\tau}|} \quad (15)$$

Statistical Bases and Descriptive Statistics

In this study, data from four cryptocurrencies with the highest market capitalization during the period from January 17, 2018, to December 18, 2024, were utilized. Table 1 presents the descriptive statistics for these four cryptocurrencies, while Figure 1 illustrates the trend of changes in their logarithmic price data. 85% of the initial portion of the data was used for training, and 15% of the remaining data was allocated for conducting cross-validation.

Table 1.

Data Descriptive Statistics

	Count	mean	min	25%	50%	75%	max	std
BTC	2524	9.883	8.082	9.095	10.007	10.663	11.573	0.885
DOGE	2524	-3.767	-6.478	-5.854	-2.780	-2.218	-0.379	1.867
ETH	2524	6.794	4.434	5.544	7.290	7.793	8.479	1.164
XRP	2524	-0.767	-1.969	-1.142	-0.745	-0.481	0.998	0.508

(Source: Researcher's Findings)

Figure 1.

The Logarithmic Price Trends of Four Major Cryptocurrencies



(Source: Researcher's Findings)

As evident in the data, both trend-based and cyclical behaviors, alongside stochastic components, are observable in the charts. Subsequently, efforts were made to estimate the models aimed at decomposing these elements within the logarithmic prices of four major cryptocurrencies.

Findings

Initially, three key time-series forecasting methods based on the Recurrent Neural

Network (RNN) approach were selected for conducting a performance comparison. Each model was cross-validated using various hyperparameters (Table 2), and their predictive capabilities were evaluated across three forecast horizons: 1 day, 7 days (one week), and 30 days (one month). As shown in Table 2, the required input size for achieving optimal predictions increases on average across all models as the forecast horizon extends.

Table 2.
Hyperparameters Tuning of three Models at Different Horizons

Time Horizons	Model	Input Size	Batch Size	Encoder Hidden Size	encoder_n_layers	decoder_hidden_size	decoder_layers	learning_rate	max_steps	encoder_dropout
1 day	RNN	30	5	54	2	128	2	0.010	500	0.20
	LSTM	28	5	54	2	128	2	0.015	500	0.15
	GRU	28	5	32	2	64	2	0.015	500	0.15
7 days	RNN	30	5	54	2	128	2	0.010	500	0.20
	LSTM	28	5	54	2	128	2	0.015	500	0.15
	GRU	28	5	32	2	64	2	0.015	500	0.15
30 days	RNN	60	8	54	2	128	2	0.010	500	0.20
	LSTM	60	5	54	2	128	2	0.015	500	0.15
	GRU	50	5	32	2	64	2	0.015	500	0.15

(Source: Researcher's Findings)

Table 3 presents the prediction errors of the three deep learning and neural networks models for different forecast horizons using 360 daily observations (approximately one year). The findings indicated the following:

First: In a short-term forecasting (1-day horizon), the LSTM model performs well in predicting the prices of three cryptocurrencies—DOGE, ETH, and XRP. For BTC price prediction, the performance of LSTM and GRU models is nearly identical, while the RNN model does not demonstrate a superior predictive power for any of the cryptocurrencies.

Second: In a medium-term forecasting (7-day horizon), the GRU model outperforms the others in predicting the prices of all four cryptocurrencies. It is worth noting that the prediction errors, assessed using RMSE, MAE, MAPE, and SMAPE metrics, increase compared to the 1-day horizon, reflecting reduced accuracy with longer forecast horizons.

Third: In a long-term forecasting (30-day horizon): the LSTM model exhibits the best performance in predicting the prices of all four cryptocurrencies. Similar to previous findings, the prediction errors for the 30-day horizon are higher than those for the 7-day horizon, further emphasizing the decline in accuracy as the forecast horizon expands.

Table 3.

Forecasting the Error of the Cryptocurrencies' Prices for Different Deep Learning Models


Horizon = 1 day								
crypto	metric	RNN	LSTM	GRU	crypto	RNN	LSTM	GRU
BTC	rmse	0.145	0.110	0.117	ETH	0.217	0.056	0.114
	mae	0.131	0.101	0.099		0.193	0.047	0.108
	mape	0.012	0.009	0.009		0.024	0.006	0.013
	smape	0.006	0.005	0.004		0.012	0.003	0.007
DOGE	rmse	0.309	0.253	0.298	XRP	0.187	0.150	0.21
	mae	0.287	0.229	0.241		0.103	0.064	0.099
	mape	0.151	0.146	0.166		0.172	0.121	0.192
	smape	0.079	0.064	0.069		0.113	0.077	0.122
Horizon = 7 day								
crypto	metric	RNN	LSTM	GRU	crypto	RNN	LSTM	GRU
BTC	rmse	0.213	0.186	0.067	ETH	0.255	0.267	0.071
	mae	0.189	0.139	0.052		0.213	0.249	0.052
	mape	0.017	0.013	0.005		0.027	0.031	0.006
	smape	0.008	0.006	0.002		0.013	0.015	0.003
DOGE	rmse	0.433	0.524	0.184	XRP	0.193	0.264	0.182
	mae	0.388	0.371	0.129		0.127	0.154	0.087
	mape	0.196	0.270	0.09		0.243	0.302	0.197
	smape	0.108	0.098	0.040		0.137	0.168	0.103
Horizon = 30 day								
crypto	metric	RNN	LSTM	GRU	crypto	RNN	LSTM	GRU
BTC	rmse	0.558	0.191	0.237	ETH	0.425	0.145	0.238
	mae	0.544	0.143	0.211		0.402	0.111	0.207
	mape	0.049	0.013	0.019		0.05	0.014	0.026
	smape	0.025	0.007	0.010		0.026	0.007	0.013
DOGE	rmse	1.277	0.487	0.505	XRP	0.272	0.241	0.257
	mae	1.243	0.360	0.43		0.14	0.121	0.126
	mape	0.69	0.227	0.255		0.372	0.306	0.336
	smape	0.242	0.088	0.103		0.124	0.115	0.116

(Source: Researcher's Findings)

Finally, to provide a more detailed assessment of the predictive power of the deep learning and neural network models (RNN, GRU, and LSTM) for the studied cryptocurrencies across 1-, 7-, and 30-day horizons, Figure 2 has been plotted in four separate panels.

Figure 2.

Prediction of the Logarithmic Prices of Four Cryptocurrencies Using Deep Learning and Neural Network Methods for Different Horizons (1, 7, and 30 Days)



(Source: Researcher's Findings)

Discussion and Conclusion

In this study, the performance of different deep learning and neural networks approaches for forecasting the logarithmic prices of four major cryptocurrencies was evaluated. The dataset consisted of daily observations spanning from January 17, 2018, to December 18, 2024. Three models from the RNN family—RNN, LSTM, and GRU—were employed, as they represent some of the most important deep learning and neural networks and AI techniques for time-series forecasting. For the analysis, 85% of the total data

(approximately the first 6 years) was allocated for training, while the remaining 15% (approximately the final year) was reserved for testing. Initially, three prominent RNN-based forecasting methods were selected for conducting the performance comparison. These models were then tested across the forecast horizons of 1 day, 7 days (one week), and 30 days (one month), with each being cross-validated under varying hyperparameters. The key findings regarding the model specifications and the comparative performance of three nonlinear deep learning and neural network models (RNN, LSTM, and GRU) are summarized below:

- As the forecast horizon increases, the average input size required for achieving the best predictions rises across all the models.
- In short-term forecasting (1-day horizon), the LSTM model demonstrates a superior performance for DOGE, ETH, and XRP price predictions. For BTC, both LSTM and GRU exhibit a similar performance, while the RNN model lags behind.
- In medium-term forecasting (7-day horizon), the GRU model outperforms the others for all four cryptocurrencies. The prediction errors increase compared to the 1-day horizon, indicating a reduced accuracy with extended horizons.
- In long-term forecasting (30-day horizon), the LSTM model shows the best overall performance for all four cryptocurrencies. Consistent with earlier results, the prediction errors are highest for the 30-day horizon, highlighting the general decline in accuracy as the forecast horizon lengthens.
- The choice of an appropriate deep learning and neural networks method depends on the forecast horizon and data characteristics.

The findings suggested that selecting the right deep learning and neural networks approach is contingent upon the specific forecast horizon and data attributes. Given the superior accuracy of the LSTM model for 1- and 30-day horizons and the GRU model's dominance in the 7-day horizon, it is recommended that investors and financial market participants pay attention to cross-market signaling and profitable strategies. By understanding the inter-market dynamics and leveraging these predictive insights, stakeholders can better manage risks and achieve desirable returns.

Future research could enhance the forecasting accuracy and robustness by exploring hybrid models that integrate deep learning architectures with external factors, such as macroeconomic indicators, social media sentiment, or regulatory changes. Additionally, leveraging advanced architectures like Graph Neural Networks (GNNs) —capable of capturing interdependencies among cryptocurrencies or market entities—could address complex spatiotemporal patterns. These approaches, combined with real-time data streams and interpretable AI techniques, may improve the adaptability to abrupt market shifts and enhance the scalability across diverse forecasting horizons.

These findings provide valuable insights for estimating the mean equation, which is instrumental in forecasting the expected returns of cryptocurrency assets for risk management purposes.

REFERENCES

- Bengio, Y., Simard, P., & Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Trans. Neural Networks*, 5(2), 157–166.
- Catania, L., Grassi, S., & Ravazzolo, F. (2019). Forecasting cryptocurrencies under model and parameter instability. *International Journal of Forecasting*, 35(2), 485–501.
- Che, Z., Purushotham, S., Cho, K., Sontag, D., & Liu, Y. (2018). Recurrent neural networks for multivariate time series with missing values. *Scientific Reports*, 8(1), 6085, 1–12.
- Chen, W., Xu, H., Jia, L., & Gao, Y. (2021). Machine learning model for bitcoin exchange rate prediction using economic and technology determinants. *International Journal of Forecasting*, 37(1), 28–43.
- Chen, X., Wei, L., & Xu, J. (2017). House price prediction using LSTM. *arXiv preprint arXiv:1709.08432*.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Conrad, C., Custovic, A., & Ghysels, E. (2018). Long-and short-term cryptocurrency volatility components: A GARCH-MIDAS analysis. *Journal of Risk and Financial Management*, 11(2), 23.
- Dong, F., Xu, Z., & Zhang, Y. (2022). Bubbly Bitcoin. *Economic Theory*, 74(3), 9731015.
- Fang, T., Su, Z., & Yin, L. (2020). Economic fundamentals or investor perceptions? The role of uncertainty in predicting long-term cryptocurrency volatility. *International Review of Financial Analysis*, 71, 101566.
- Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European journal of operational research*, 270 (2), 654–669.
- Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning* (Vol. 1, No. 2). Cambridge: MIT press.
- Graves, A. (2012). Sequence transduction with recurrent neural networks. *ArXiv preprint arXiv: 1211.3711*.
- Hamayel, M. J., & Owda, A. Y. (2021). A novel cryptocurrency price prediction model using GRU, LSTM and bi-LSTM machine learning algorithms. *AI*, 2(4), 477–496.
- Hansson, F., & Rostami, J. (2019). *Time series forecasting of house prices: an evaluation of a support vector machine and a recurrent neural network with LSTM cells*. Uppsala University Thesis. May 2019.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780.
- Ibrahim, A., Kashaf, R., & Corrigan, L. (2021). Predicting market movement direction for Bitcoin: A comparison of time series modeling methods. *Computers and Electrical Engineering*, 89, 106905.
- Jozefowicz, R., Zaremba, W., & Sutskever, I. (2015, June). An empirical exploration of recurrent network architectures. In *International conference on machine learning* (pp. 2342–2350). PMLR.
- Klein, T., Thu, H. P., & Walther, T. (2018). Bitcoin is not the new gold—a comparison of volatility, correlation, and portfolio performance. *International Review of Financial Analysis*, 59, 105–116.

- Kristjanpoller, W., & Minutolo, M. C. (2018). A hybrid volatility forecasting framework integrating GARCH, artificial neural network, technical analysis and principal components analysis. *Expert Systems with Applications*, 109, 1–11.
- Kubat, M. (2015). Virtual currency Bitcoin in the scope of money definition and store of value. *Procedia Economics and Finance*, 30, 409–416. [https://doi.org/10.1016/S2212-5671\(15\)01308-8](https://doi.org/10.1016/S2212-5671(15)01308-8).
- Kyriazis, N., Papadamou, S., & Corbet, S. (2020). A systematic review of the bubble dynamics of cryptocurrency prices. *Research in International Business and Finance*, 54, 101254.
- Lahmire, S., & Bekiros, S. (2019). Cryptocurrency forecasting with deep learning chaotic neural networks. *Chaos, Solitons & Fractals*, 118, 35–40.
- Li, Y., & Dai, W. (2020). Bitcoin price forecasting method based on CNN- LSTM hybrid neural network model. *The Journal of Engineering*, 2020(13), 344–347.
- Limsombunchai, V., Gan Ch., & Lee, M. (2004). House price prediction: hedonic price model vs. ANN. *American Journal of Applied Sciences*, 1 (3), 193-201.
- Nakano, M., Takahashi, A., & Takahashi, S. (2018). Bitcoin technical trading with artificial neural network. *Physica A: Statistical Mechanics and its Applications*, 510, 587609.
- Parekh, R., Patel, N. P., Thakkar, N., Gupta, R., Tanwar, S., Sharma, G., ... & Sharma, R. (2022). DL-Guess: Deep learning and sentiment analysis-based cryptocurrency price prediction. *IEEE Access*, 10, 35398-35409.
- Patel, M. M., Tanwar, S., Gupta, R., & Kumar, N. (2020) A deep learning-based cryptocurrency price prediction scheme for financial institutions. *Journal of Information Security and Applications*, 55, 102583.
- Sayadi Nezhad, S., Esmailzadeh Maghari, A., & Rostami, M.R. (2023). Presenting the forecasting model of Bitcoin return using the hybrid method of deep learning-signal decomposition algorithm (CEEMD-DL). *Journal of Financial Economics*, 17(1), 217-238 (In Persian).
- Seabe, P. L., Moutsinga, C., & Pindza, E. (2023). Forecasting cryptocurrency prices using LSTM, GRU, and Bi-Directional LSTM: A deep learning approach. *Fractal Fractional*, 7(2), 203, 7020203.
- Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review: 2005– 2019, *Applied soft computing*, 90, 106181.
- Song, Y. (2018). *Stock trend prediction: Based on machine learning methods* (Doctoral dissertation, UCLA).
- Tsay R. S. (2005). *Analysis of financial time series*, John Wiley & Sons.
- Walther, T., Klein, T., & Bouri, E. (2019). Exogenous drivers of Bitcoin and cryptocurrency volatility—a mixed data sampling approach to forecasting. *Journal of International Financial Markets Institutions and Money*, 63, 101133.
- Xu, L., Li, C., Xie, X., & Zhang, G. (2018). Long short-term memory network based hybrid model for short-term electrical load forecasting. *Information*, 9(7), 165.
- Zhang, W., Wang, P., Li, X., & Shen, D. (2018). Some stylized facts of the cryptocurrency market. *Applied Economics*, 50(55), 5950–5965.
- Zhang, Z., Dai, H. N., Zhou, J., Mondal, S. K., García, M. M., & Wang, H. (2021). Forecasting cryptocurrency price using convolutional neural networks with weighted and attentive memory channels. *Expert Systems with Applications*, 183, 115378.