A COMPARATIVE STUDY OF MACHINE LEARNING MODELS ON CRYPTOCURRENCY PRICES

Dušan ČATLOCH, Eva CHOVANCOVÁ
Department of Computers and Informatics, Faculty of Electrical Engineering and Informatics,
Technical University of Košice, Letná 9, 042 00 Košice, Slovak Republic, Tel: +421 55 602 2676
E-mails: dusan.catloch@tuke.sk, eva.chovancova@tuke.sk

ABSTRACT

The volatility of cryptocurrencies poses challenges for accurate price forecasting. This study compares traditional machine learning models (Linear Regression, Support Vector Regression), ensemble methods (Random Forest, Gradient Boosting), and deep learning architectures (LSTM, GRU) in predicting the daily prices of Bitcoin (BTC), Ethereum (ETH), Cardano (ADA), Ripple (XRP), and Polkadot (DOT). Using historical data from CoinGecko, a sliding window, and normalisation, we assess models by Root Mean Square Error (RMSE) and relative percentage error. Results show that LSTM and GRU achieve the best overall accuracy, while Linear Regression remains competitive for stable assets such as BTC, ADA, and DOT. Ensemble methods performed moderately, whereas SVR consistently underperformed. The findings underline the importance of matching prediction models to the characteristics of specific cryptocurrencies.

Keywords: *Cryptocurrency, Price prediction, Machine learning, Deep learning, Predictive models, Machine learning algorithms*

1. INTRODUCTION

Cryptocurrencies have become a major asset class, attracting both institutional and retail investors. Popular coins such as BTC, ETH, ADA, XRP, and DOT represent a significant share of the market, yet their extreme volatility makes price prediction difficult. External factors including sentiment, limited intrinsic value anchors, and macroeconomic influences further complicate forecasting.

Classical econometric approaches often fail to capture the nonlinear dynamics of cryptocurrency markets. Machine learning (ML) and deep learning (DL) methods provide flexible, data-driven alternatives capable of modelling complex temporal patterns. Their successful application in stock and forex markets motivates their use in cryptocurrency prediction.

This study compares a range of models. We compared Linear Regression, Support Vector Regression (SVR), Random Forest, Gradient Boosting Regression, LSTM, and GRU and also we applied them to daily closing prices of five major cryptocurrencies. Historical data from CoinGecko are processed with a sliding window approach, and models are evaluated on a fixed test set (January-October 2024) using RMSE and relative percentage error. Beyond accuracy, the comparison highlights balance between simplicity and complexity, offering guidance on model selection for practical forecasting.

2. SUBJECT

The rapid evolution of financial technologies and the increasing popularity of decentralised assets have sparked a surge of interest in cryptocurrencies. As their role in global markets expands, so does the need to understand and anticipate their behavior. This chapter delves into the field of cryptocurrency price forecasting, exploring the challenges and opportunities it presents. It outlines the motivation for using modern computational techniques to improve forecasting accuracy and lays the foundation for the models and methodologies analyzed in this study.

2.1. Introduction to cryptocurrency forecasting

Cryptocurrencies have become a significant part of the global financial ecosystem, attracting investors, researchers, and institutions alike. Their decentralised nature, combined with high market volatility, presents both unique opportunities and challenges. One of the critical areas of interest is accurate price forecasting, which can assist in investment strategies, risk management, and the development of automated trading systems.

Unlike traditional financial assets, cryptocurrencies are highly susceptible to external factors such as regulatory changes, market sentiment, technological developments, and social media influence. This volatility makes short and medium term forecasting particularly complex but also highly valuable.

As the market continues to mature, the need for reliable prediction models grows. Machine learning (ML) and deep learning (DL) techniques have shown considerable potential in capturing the nonlinear and temporal dynamics of cryptocurrency prices. This study aims to explore and compare various forecasting models, including both classical and deep learning methods, applied to multiple digital assets. By doing so, we aim to identify general trends in model performance and offer insights into the most effective strategies for cryptocurrency forecasting.

2.2. Research objectives

The primary objective of this study is to evaluate and compare the performance of different machine learning and deep learning models in forecasting cryptocurrency prices. The focus is placed on both traditional regression based models, such as Linear Regression, Support Vector Regression (SVR), Random Forest, Gradient Boosting Regression (GBR) and deep learning models, namely Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks.

By applying these models to real world price data, the study seeks to:

- Assess the prediction accuracy of classical versus deep learning models.
- Analyse model behaviour across different types of cryptocurrencies with varying market characteristics.
- Provide next day forecasts and performance metrics such as Root Mean Square Error (RMSE) and relative percentage error.

The cryptocurrencies selected for this research are BTC, ETH, XRP, ADA and DOT. These assets were chosen due to their market relevance, differing use cases, and diverse volatility profiles. Including a variety of cryptocurrencies allows the models to be tested under different market dynamics and enhances the generalisability of the results.

2.3. Related work

Cryptocurrency price forecasting has received increasing attention in recent years due to the high volatility of digital assets and their growing economic significance. Several studies have explored the use of machine learning and deep learning techniques to predict price movements.

In Mohil Maheshkumar Patel work [1], the authors compared multiple machine learning models, including Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) for predicting Litecoin and Monero prices. While the study provided valuable insights into model performance, it focused solely on Litecoin and Monero and did not explore other learning models which are capable of capturing temporal dependencies in time series data.

Another notable study is Patel Jay's work [2], where the authors applied MultiLayer Perceptron (MLP) and Long Short-Term Memory (LSTM) models for BTC, ETH, and Litecoin price forecasting. These results confirmed the effectiveness of models, but the study was limited to 3 cryptocurrencies and lacked a comparison with traditional models or other deep learning architectures.

Compared to these studies, our work makes the following contributions:

- We evaluate both classical machine learning models (e.g., Linear Regression, SVR, Random Forest, Gradient Boosting Regression) and deep learning models (LSTM, GRU) in a unified framework.
- We conduct experiments on five major cryptocurrencies (BTC, ETH, XRP, ADA, DOT), offering a broader perspective on model generalisability across different assets.
- We provide a comparison based not only on RMSE but also on relative percentage error and next day forecast, highlighting practical implications of each model performance.

This broader scope allows us to identify trends in model accuracy and behaviour across multiple cryptocurrencies, filling a gap left by prior studies which have primarily focused on single-asset predictions or a limited set of algorithms.

3. METHODS

This chapter outlines the methodological framework used to conduct the forecasting experiments. It includes a description of the dataset, the preprocessing techniques applied, and the implementation of various machine learning and deep learning models. The goal is to ensure a consistent and reproducible approach for comparing the predictive performance of each method.

3.1. Dataset description

The dataset used in this study was obtained from the CoinGecko API, a widely recognised source for cryptocurrency market data. For consistency and reliability, historical daily price data were collected for five major cryptocurrencies: BTC, ETH, XRP, ADA, DOT. These cryptocurrencies were selected due to their high market capitalisation, active trading volume, and relevance in both research and practical applications. BTC and ETH are widely considered foundational assets in the cryptocurrency ecosystem, often serving as benchmarks for the broader market. XRP is known for its unique consensus mechanism and application in crossborder payments, making it of particular interest in financial technology research. ADA and DOT, both representing third-generation blockchain platforms, offer innovative approaches to scalability, interoperability, and sustainability. Their inclusion enables the evaluation of prediction models across a diverse set of technological architectures and market dynamics.

The dataset spans from the earliest available records (typically 2013 for BTC and 2020 for newer assets like DOT) up to the first quarter of 2025. Each record includes the date and the corresponding daily closing price in USD. The temporal range allows the models to capture both long term trends and short term fluctuations.

Prior to training, the data underwent several preprocessing steps. Missing values, which were minimal, were removed to avoid introducing bias. All date fields were converted into a consistent datetime format. For machine learning algorithms, closing prices were normalised using Min-Max scaling to improve convergence and model performance. Additionally, the data was split into training and testing sets, with the final portion (typically the last 30-60 days) reserved for evaluating prediction accuracy.

3.2. Data preparation

Before feeding the data into machine learning models, several preprocessing steps were performed to ensure consistency and facilitate model training.

First, all price values were normalized using the MinMaxScaler technique to scale data into the range [0, 1]. This normalisation is essential for improving convergence during training, particularly for neural network models that are sensitive to feature scaling.

Next, the time series data were transformed into input-output sequences using a sliding window (also known as a rolling window) approach. In this step, fixed length sequences of consecutive daily prices were used as input features to predict the price on the following day. This method allows models to capture temporal dependencies and patterns in the price movements.

Finally, the dataset was split into training and testing subsets. The training set covers the period from the earliest available records up to December 31, 2023, while the testing set spans from January 1, 2024 to October 1, 2024. This corresponds approximately to an 85/15 ratio, ensuring that model performance is evaluated on a recent and unseen portion of the data. Such a temporal split is a standard approach in time series forecasting tasks, as it respects the chronological order of the observations.

3.2.1. Hyperparameter settings

At the beginning of the experiments, we initialized the deep learning models with standard baseline hyperparameters that are commonly used in financial time series forecasting tasks. For both LSTM and GRU architectures, the initial configuration consisted of 3 layers with 50 hidden units per layer, a dropout rate of 0.2, the Adam optimizer with a learning rate of 10^{-3}, a batch size of 32, and a maximum of 25 training epochs. These values were selected as a reasonable starting point to establish a baseline performance for comparison with traditional machine learning models.

The baseline settings were subsequently replaced by a dynamic hyperparameter tuning procedure in order to improve model accuracy and ensure fair comparison across all models. In this tuning phase, we systematically explored the following ranges: number of layers $\in \{2,3\}$, hidden units $\in \{32,64,96\}$, dropout $\in \{0.1,0.2,0.3\}$, batch size $\in \{32,64\}$, and learning rates $\in \{5\times10^{-4},10^{-3}\}$. Training was conducted for up to 100 epochs with early stopping (patience of 8 epochs) and learning rate scheduling (ReduceLROnPlateau, factor 0.5, patience of 4, minimum learning rate 10^{-5}). To avoid information leakage, the last 10% of the training set was reserved as a validation split for hyperparameter optimization, while the test period (2024-01-01 to 2024-10-01) was strictly held out for final evaluation only.

For the classical machine learning models, similar procedures were applied. SVR was tuned over the ranges $C \in \{10,100,300,1000\}$, $\gamma \in \{10^{-3},10^{-2},10^{-1}\}$, and $\varepsilon \in \{0.01,0.05,0.1\}$. Random Forest was tuned over $n\_estimators \in \{200,400,800\}$ and $max\_depth \in \{6,10,16,None\}$, while Gradient Boosting Regression explored $n\_estimators \in \{200,400,800\}$, $learning\_rate \in \{0.05,0.1,0.2\}$, and $max\_depth \in \{2,3,4\}$. Linear Regression was applied in its standard closed-form solution without additional parameters.

After the tuning phase, the best configurations were selected separately for each dataset based on the lowest validation RMSE, and the models were retrained on the full training set before being evaluated on the held-out test period. In practice, the deep learning models typically converged within 30–50 epochs, demonstrating both stability and efficiency of the optimization process. This two-step approach (baseline setup followed by dynamic tuning) ensured reproducibility, robustness, and a transparent methodology for fair performance comparison across all models.

3.2.2. Computational setup

All experiments were conducted on a local workstation equipped with an Intel Core i7-9750H CPU @ 2.60 GHz, 16 GB of RAM, and an NVIDIA GeForce RTX 2060 GPU with 6 GB GDDR6 memory (Windows 11 Home, version 24H2).

In terms of computational requirements, the classical machine learning models (Linear Regression, SVR, Random Forest, Gradient Boosting) were lightweight and executed within seconds to a few minutes on the CPU, depending on the hyperparameter configuration. Ensemble methods (RF, GBR) required longer training times than LR or SVR, but all remained computationally efficient.

The deep learning models (LSTM and GRU) were trained using the GPU-accelerated TensorFlow backend. A single training run with baseline hyperparameters (25 epochs, batch size 32) completed in approximately 2-3 minutes. During hyperparameter tuning with extended training (up to 100 epochs, early stopping, and validation monitoring), the models typically converged within 30-50 epochs, resulting in total runtimes of 10-15 minutes per configuration. Complete tuning across the search space for each cryptocurrency required approximately 2-3 hours, which remained manageable on the given hardware setup.

Overall, while the deep learning models demanded substantially more resources than the classical approaches, GPU acceleration ensured feasible training times and allowed systematic hyperparameter optimization without requiring specialized high-performance computing infrastructure.

3.3. Machine learning models

In this study, we implemented and compared several traditional machine learning models for time series forecasting of cryptocurrency prices. Each model was trained using the same rolling window input sequences and evaluated on the same test set for consistency.

3.3.1. Linear regression

Linear Regression (LR) is one of the simplest predictive models, which assumes a linear relationship between input variables and the target value. Despite its simplicity, LR often provides competitive performance on financial time series, especially when trends are stable or linear in nature [3].

3.3.2. Support vector regression (SVR)

Support Vector Regression extends the principles of Support Vector Machines (SVMs) to regression problems. SVR seeks to find a function that approximates the data within a specified error margin, and it is effective in high-dimensional spaces. Its kernel trick allows capturing nonlinear relationships between variables [4].

3.3.3. Random forest

Random Forest (RF) is an ensemble learning method based on decision trees. It constructs multiple decision trees during training and outputs the mean prediction of the individual trees. Random Forest is robust against overfitting and performs well in capturing nonlinear dependencies and feature interactions [5].

3.3.4. Gradient boosting regression

Gradient Boosting Regression (GBR) is an ensemble learning technique that builds predictive models in a stage wise manner. It sequentially adds decision trees to minimise the residual errors of the model using gradient descent. GBR is capable of capturing complex nonlinear relationships and often achieves high accuracy in financial forecasting tasks. In our study, a tree based implementation such as XGBoost or a similar library was employed [6].

3.4. Deep learning models

Deep learning models have gained significant attention in time series forecasting due to their ability to model complex, nonlinear temporal patterns. In this study, we employed two prominent architectures from the family of recurrent neural networks (RNNs): Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU). Both are designed to handle long range dependencies in sequential data, overcoming limitations of traditional RNNs such as the vanishing gradient problem.

3.4.1. Long short-term memory (LSTM)

LSTM networks utilise memory cells and gating mechanisms (input, forget, and output gates) to regulate the flow of information through the network. This architecture allows the model to selectively retain relevant information and discard less useful patterns across time steps, making it suitable for capturing the temporal dynamics of cryptocurrency prices [7].

3.4.2. Gated recurrent unit (GRU)

GRU is a streamlined alternative to LSTM that merges certain gates to reduce complexity while maintaining the ability to learn from long-term dependencies. Despite its simpler structure, GRU often delivers performance comparable to LSTM, especially in tasks involving noisy and highly volatile data such as financial time series [8].

The architectures of both models were designed to allow flexible tuning of hyperparameters such as the number of units, learning rate, batch size, and number of layers. These parameters were optimised dynamically during training using a systematic tuning strategy, ensuring that each model was adequately adapted to the characteristics of the dataset.

3.5. Evaluation metrics

To objectively assess the performance of the developed machine learning and deep learning models for cryptocurrency price prediction, we employed a combination of absolute and relative error based evaluation metrics. These are complemented by a practical forecasting test simulating a real world use case. By utilising multiple evaluation strategies, we ensure a comprehensive understanding of both numerical accuracy and practical utility of the models across varying market conditions.

3.5.1. Root Mean Squared Error (RMSE)

Root Mean Squared Error (RMSE) is one of the most widely adopted error metrics in regression problems, especially when predicting continuous values such as asset prices. RMSE provides a direct measure of the average magnitude of the prediction error. It is computed as follows:

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

where y_i is the actual value and $\hat{y}_i$ is the predicted value. This metric gives higher weight to larger errors by squaring the differences before averaging. As such, RMSE is particularly useful when large deviations from actual values are undesirable, a common scenario in financial domains, where even moderate forecasting errors can lead to significant financial consequences. A lower RMSE value indicates better model performance. However, since it is measured in the same units as the predicted variable, RMSE does not provide information on how significant the error is in relative terms [9].

3.5.2. Relative percentage error (RPE)

To address the limitations of absolute metrics like RMSE, we also employ a normalised evaluation measure, the Relative Percentage Error (RPE). This metric expresses the magnitude of error as a percentage of the actual value, providing a more interpretable and comparable measure across different cryptocurrencies with varying price ranges. The formula is defined as:

$$\text{RPE (\%)} = \frac{|\text{Actual} - \text{Predicted}|}{\text{Actual Value}} \times 100$$

By calculating this percentage for each prediction and averaging the result across the test dataset, we gain insight into the average relative deviation. This is especially valuable when comparing model performance on assets that differ in scale, for example BTC versus smaller cap cryptocurrencies. A lower RPE value indicates a model that maintains higher consistency and proportional accuracy, which is desirable for traders and analysts focusing on risk adjusted performance [10].

3.5.3. Next day price forecast

In addition to quantitative metrics, we perform a qualitative evaluation by simulating next day price predictions. This mirrors real world application, where a model would be used to forecast the closing price of a cryptocurrency for the following trading day based on recent historical data.

Although this next day prediction is not formally included as a numerical metric, it allows for visual comparison of predicted versus actual price trajectories, provides insight into the model‘s responsiveness to market fluctuations, and highlights potential issues such as time lag or instability in predictions.

By combining RMSE and Relative Percentage Error, we cover both the raw magnitude of prediction errors and their significance in context. The inclusion of next day forecasting further extends the evaluation beyond statistics, addressing practical feasibility and interpretability. This multifaceted approach provides a solid foundation for comparing different models and choosing the most effective one for cryptocurrency price forecasting in volatile environments.

4. RESULTS

This chapter presents the experimental results obtained from applying various forecasting models to the selected cryptocurrency datasets. The models are evaluated using standard performance metrics and visual comparisons, offering insights into their accuracy and reliability. The results highlight both the strengths and limitations of each approach in real world forecasting scenarios.

4.1. Prediction performance

The performance of all tested models was evaluated using Root Mean Squared Error (RMSE) and Relative Percentage Error across five selected cryptocurrencies: BTC, ETH, XRP, ADA, DOT. The results revealed that model performance varied depending on the asset, highlighting the importance of testing algorithms on diverse data sources.

For high market cap assets such as BTC and ETH, deep learning models, particularly LSTM and GRU, showed strong performance in terms of both RMSE and percentage error. In the case of BTC, LSTM achieved the lowest relative error (3.72%) and GRU closely followed (4.14%), outperforming classical models such as SVR, which produced the highest error (15.13%). Similar patterns were observed for ETH, where the relative error remained under 6% for most models, except for SVR, which again reached over 56%.

In contrast, for lower valued assets like XRP and ADA, classical models such as Linear Regression and Random Forest occasionally performed competitively or even outperformed deep learning models in RMSE. For example, on the XRP dataset, Random Forest and Gradient Boosting Regression achieved a low RMSE of 0.02 with a 3.62% error, matching Linear Regression and outperforming LSTM and GRU in raw accuracy. However, SVR again demonstrated instability, with over 9% error.

The results for DOT mirrored these findings: Linear Regression and Gradient Boosting Regression maintained a stable 4.2% error, while SVR produced the highest (27.3%). Deep learning models yielded around 10% error, suggesting that their advantage is less pronounced on smaller cap or more volatile assets.

Overall, LSTM and GRU consistently provided balanced results with low error rates across all datasets. SVR, while theoretically suitable for time series, underperformed in most scenarios due to its sensitivity to parameter tuning and volatility. Ensemble methods like Random Forest and Gradient Boosting Regression proved to be reliable, especially on assets with lower price ranges.

Deep learning models showed strong adaptability across assets, while simpler models remained effective for more stable markets. A summary of RMSE and Relative Percentage Error for each model and cryptocurrency is provided in Tables 1 and 2.

Table 1 Root Mean Squared Error (RMSE) of models across all cryptocurrencies

Model	BTC	ETH	ADA	DOT	XRP
LR	1754.86	101.13	0.02	0.30	0.02
SVR	9085.52	158.02	0.05	3.78	0.05
RF	3278.63	120.77	0.02	0.34	0.02
GBR	3466.31	113.72	0.02	0.34	0.02
LSTM	2233.34	138.04	0.03	0.52	0.03
GRU	2483.67	174.70	0.03	0.54	0.03

Table 2 Relative Percentage Error (%) of models across all cryptocurrencies

Model	BTC	ETH	ADA	DOT	XRP
LR	2.92%	3.34%	3.62%	4.46%	3.62%
SVR	15.13%	5.22%	9.04%	56.22%	9.04%
RF	5.46%	3.99%	3.62%	5.06%	3.62%
GBR	5.77%	3.76%	3.62%	5.06%	3.62%
LSTM	3.72%	4.56%	5.43%	7.73%	5.43%
GRU	4.14%	5.77%	5.43%	8.03%	5.43%

4.2. Forecast accuracy

To complement the numerical evaluation metrics presented earlier, we visualize the predicted versus actual prices for each cryptocurrency in Figures 1, 2, 3, 4, and 5. These plots provide a clear illustration of each model’s forecasting behaviour over time, capturing both short-term fluctuations and broader market trends.

Across most assets, traditional models such as Linear Regression and ensemble methods (Random Forest and Gradient Boosting Regression) followed the price curve with reasonable accuracy. However, they occasionally showed delays in responding to sharp price movements. Support Vector Regression (SVR) displayed greater variance and less stability, particularly on volatile assets like BTC and ETH.

Deep learning models, especially LSTM and GRU, consistently provided smoother and more realistic predictions, closely tracking the actual price trajectory. These models exhibited a better capacity to adapt to nonlinear temporal dependencies, which is especially evident in assets with complex patterns, such as DOT and XRP.

Additionally, the one day ahead forecasts, visible as the final extension of each model's curve, demonstrate the models' immediate predictive capabilities. LSTM and GRU again proved effective in generating predictions that remained close to the observed values, making them strong candidates for short-term trading decision support.

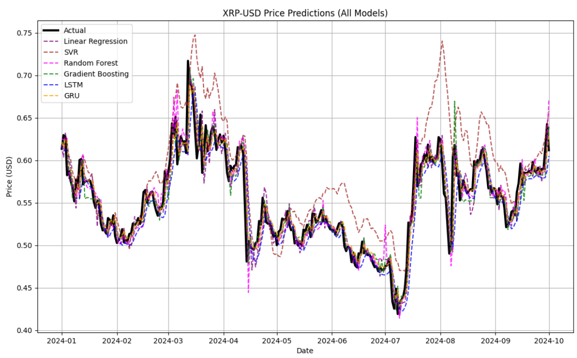

Fig. 1 XRP-USD Price Predictions (All Models)

Figure 1 shows the performance of all models in predicting XRP-USD prices.

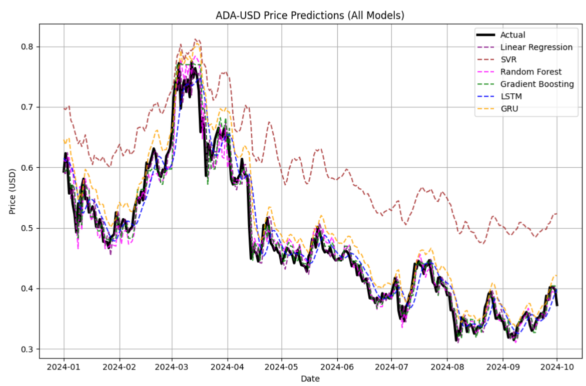

Fig. 2 ADA-USD Price Predictions (All Models)

Figure 2 shows the performance of all models in predicting ADA-USD prices.

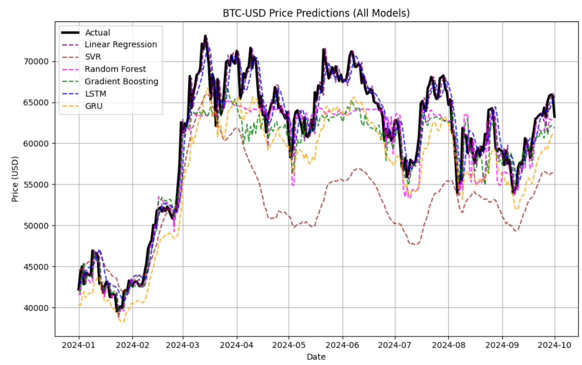

Fig. 3 BTC-USD Price Predictions (All Models)

Figure 3 shows the performance of all models in predicting BTC-USD prices.

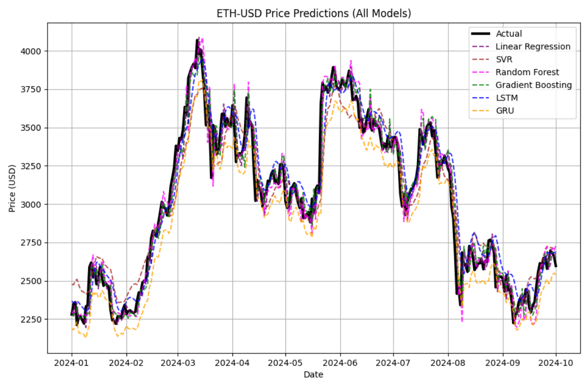

Fig. 4 ETH-USD Price Predictions (All Models)

Figure 4 shows the performance of all models in predicting ETH-USD prices.

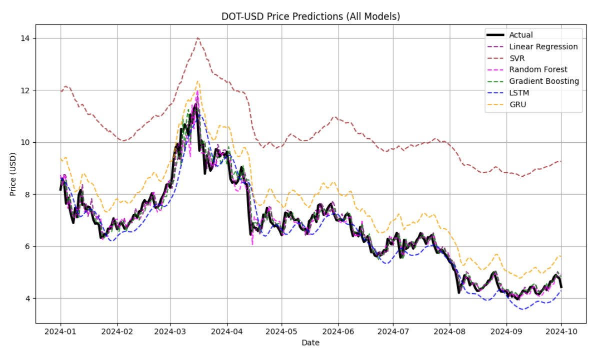

Fig. 5 DOT-USD Price Predictions (All Models)

Figure 5 shows the performance of all models in predicting DOT-USD prices.

Another practical evaluation of model usefulness can be observed through the next day price forecasts. Table 3 presents the predicted prices for each cryptocurrency on 2024-10-02, as generated by all six models. This approach allows us to assess each model's short-term forecasting capability in a realistic setting, which is highly relevant for time sensitive financial decision-making. Comparing these predictions with actual market values offers additional insight into the precision and potential practical deployment of each model.

Table 3 Next-day price predictions (USD) for each model and cryptocurrency

Model	BTC	ETH	ADA	DOT	XRP
LR	62681.37	2553.71	0.37	4.47	0.61
SVR	56234.61	2667.92	0.52	9.21	0.66
RF	61227.02	2636.11	0.38	4.46	0.62
GBR	60598.30	2597.29	0.37	4.56	0.62
LSTM	64167.47	2669.45	0.42	4.41	0.63
GRU	65941.16	2518.73	0.38	4.05	0.59

To better interpret prediction accuracy, Table 4 shows the real prices and absolute differences (Δ) between each model's prediction and the actual value.

Table 4 Real next day prices (R) and absolute difference (Δ) from prediction (USD)

	BTC	ETH	ADA	DOT	XRP
R	60872.67	2451.63	0.35	4.16	0.60
ΔLR	+1808.70	+102.08	+0.02	+0.31	+0.01
ΔSVR	-4638.06	+216.29	+0.17	+5.05	+0.06
ΔRF	+1354.35	+184.48	+0.03	+0.30	+0.02
ΔGBR	-274.37	+145.66	+0.02	+0.40	+0.02
ΔLSTM	+1294.80	+217.82	+0.07	+0.25	+0.03
ΔGRU	+5068.49	+67.10	+0.03	-0.11	-0.01

For 2024-10-02, most models slightly overestimated the price of XRP, ADA, ETH, and BTC, reflecting a general tendency toward optimistic forecasts. This pattern suggests that while the models were relatively well calibrated, they may not have fully captured short term corrections or negative momentum. Among the assets, DOT exhibited the highest prediction variability across models, with SVR and GBR overshooting the actual price by a considerable margin. In contrast, the GRU model was the only one to predict a lower price than the actual value, highlighting its potential to capture subtle downward trends that other models may overlook. These results emphasise the importance of model diversity and asset specific behaviour when selecting forecasting approaches.

5. DISCUSSION/CONCLUSIONS

This final chapter discusses the implications of the experimental findings, highlighting key insights, limitations, and directions for future research. By reflecting on the comparative performance of the applied models, we provide guidance for both academic exploration and practical applications in cryptocurrency forecasting.

5.1. Interpretation of results

The results of the conducted experiments reveal key differences between traditional machine learning models and deep learning architectures in the context of cryptocurrency price prediction.

Traditional models such as Linear Regression, Random Forest, and Gradient Boosting often achieved comparable or even better performance in terms of RMSE and percentage error, especially in the case of relatively stable cryptocurrencies like XRP and ADA. Their lower computational complexity and faster training times make them suitable for real time systems or scenarios where resources are limited.

On the other hand, deep learning models like LSTM and GRU demonstrated greater adaptability to more volatile assets such as BTC and ETH. Their ability to capture complex sequential dependencies provided improved predictive accuracy in more turbulent market conditions. However, these models require more computational resources, longer training times, and careful hyperparameter tuning, which may not always be feasible in time constrained applications.

The comparison also highlights the importance of selecting models not only based on predictive accuracy, but also considering the trade-offs between complexity, generalisability, and interpretability. While simpler models may be sufficient for certain assets or conditions, deep learning models tend to generalise better across unpredictable patterns, making them valuable for assets with high volatility or irregular behaviour.

5.2. Limitations

Despite the promising results, this study is subject to several limitations that should be acknowledged.

First, the testing period is limited to the year 2024 up to October 2nd. While this range includes various market movements, it may not fully capture long term trends or extreme events such as bull or bear market cycles. A longer historical testing window could provide a more comprehensive evaluation of model robustness.

Second, the models do not account for external market moving factors such as breaking news, geopolitical events, or regulatory changes, which can cause abrupt and significant price fluctuations. Since the models are purely data driven and rely solely on historical prices, their predictive accuracy may be limited during periods of sudden market disruptions.

Additionally, no external features (e.g., trading volume, sentiment scores, macroeconomic indicators) were included, which might further improve prediction accuracy. Future work could address these limitations by incorporating a broader dataset and exploring the effects of external signals on model performance.

5.3. Future work

There are several directions for future work that could build upon the findings of this study.

First, exploring hybrid modelling approaches that combine traditional and deep learning methods could yield improved accuracy and robustness. For example, combining tree based models for feature selection with recurrent neural networks for sequence prediction may take advantage of the strengths of both paradigms.

Second, extending the analysis to a larger and more diverse set of cryptocurrencies, including smaller cap or emerging tokens, could help assess model generalisability across different market structures. In addition, evaluating the models over longer historical periods and during varying market conditions (e.g., crashes, rallies) would provide deeper insight into model stability and reliability.

Finally, incorporating additional data sources such as sentiment analysis, social media trends, or macroeconomic indicators could enhance the context awareness of the predictions and improve real world applicability in trading systems or financial decision making.

6. CONCLUSIONS

This study evaluated the forecasting performance of multiple machine learning models, including traditional regressors and deep learning architectures, on five major cryptocurrencies: BTC, ETH, XRP, ADA, and DOT. Through comprehensive analysis using RMSE and relative percentage error metrics, we found that deep learning models, particularly GRU and LSTM, generally offered higher accuracy, especially in capturing complex temporal dependencies. However, simpler models such as Linear Regression and Gradient Boosting remained competitive in more stable market conditions.

The results underscore the importance of model selection based on asset specific behaviour and trading goals. While advanced models can offer improved accuracy, they also demand greater computational resources and data preprocessing, which may not be ideal in all practical scenarios.

From a research perspective, these findings encourage continued exploration of tailored modelling strategies, including hybrid or ensemble methods. Practically, financial analysts and algorithmic traders should consider the volatility profile and liquidity of an asset when selecting predictive models for deployment.

Ultimately, there is no one size fits all solution. Model performance is context dependent, and the most effective forecasting approach is one that balances precision, complexity, and interpretability.

ACKNOWLEDGEMENT

Funded by the EU NextGenerationEU through the Recovery and Resilience Plan for Slovakia under the project No. 17R05-04-V01-00004, Competence and Training Center for Cyber and Information Security TUKE.

REFERENCES

[1] PATEL, M. M. – TANWAR, S. – GUPTA, R. – KUMAR, N.: *A Deep Learning-based Cryptocurrency Price Prediction Scheme for Financial Institutions*, Journal of Information Security and Applications , No. (2020) `https://www.sciencedirect.com/science/article/pii/S2214212620307535`

[2] PATEL, J. – VASU, K. – PUSHPENDRA, P. – SUDEEP, T. – KUMAR, N. – MAMOUN, A.: *Stochastic Neural Networks for Cryptocurrency Price Prediction*, IEEE Access , No. (2020) `https://ieeexplore.ieee.org/document/9079491`

[3] SU, X. – YAN, X. – TSAI, C.-L.: *Linear regression*, WIREs Computational Statistics **4**, No. 3 (2012) 275–294 `https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wics.1198`

[4] AWAD, M. – KHANNA, R.: *Support Vector Regression*, Efficient Learning Machines: Theories, Concepts, and Applications for Engineers and System Designers , No. (2015) 67–80 `https://doi.org/10.1007/978-1-4302-5990-9_4`

[5] AAKASH, P. – RAKESH, K. – PATEL, V.: *A Review on Random Forest: An Ensemble Classifier*, International Conference on Intelligent Data Communication Technologies and Internet of Things (ICICI) 2018 , No. (2019) 758–763 `https://doi.org/10.1007/978-3-030-03146-6_86`

[6] OTCHERE, D. A. – GANAT, T. O. A. – OJERO, J. O. – TACKIE-OTOO, B. N. – TAKI, M. Y.: *Application of Gradient Boosting Regression Model for the Evaluation of Feature Selection Techniques in Improving Reservoir Characterisation Predictions*, Journal of Petroleum Science and Engineering **208**, No. (2022) 109244 `https://www.sciencedirect.com/science/article/pii/S0920410521008998`

[7] DIPIETRO, R. – HAGER, G. D.: *Chapter 21 - Deep Learning: RNNs and LSTM*, Handbook of Medical Image Computing and Computer Assisted Intervention. Academic Press, The Elsevier and MICCAI Society Book Series , No. (2020) 503–519 `https://www.sciencedirect.com/science/article/pii/B9780128161760000260`

[8] DEY, R. – SALEM, F. M.: *Gate-variants of Gated Recurrent Unit (GRU) Neural Networks*, Proceedings of the 2017 IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS) , No. (2017) 1597–1600 `https://doi.org/10.1109/MWSCAS.2017.8053243`

[9] Coralogix: *Root Mean Square Error (RMSE): The Cornerstone for Evaluating Regression Models*, Coralogix AI Blog , No. (2023) `https://coralogix.com/ai-blog/root-mean-square-error-rmse-the-cornerstone\-for-evaluating-regression-models/`

[10] Lumen Learning: *Absolute and Relative Error | Calculus II*, courses.lumenlearning.com , No. (n.d.) `https://courses.lumenlearning.com/calculus2/chapter/absolute-and-relative-error/`

Received June 23, 2025, accepted October 6, 2025

BIOGRAPHIES

Dušan Čatloch graduated with a Masteros degree in 2024 at the Department of Computers and Informatics, Technical University of Košice, Slovakia. Currently, he is pursuing his Ph.D. at the same department. The main topic of his research is machine learning.

Eva Chovancová graduated (Ing.) at the Department of Computers and Informatics at the Faculty of Electrical Engineering and Informatics of the Technical University of Košice in 2009. She defended her Ph.D. t hesis in the field of computers and computer systems in 2012; her thesis title was “Specialized Processor for Computing Acceleration in the Field of Computer Vision”. Since 2012 she has been working as Assistant Professor at the Department of Computers and Informatics. Her scientific research is focused on the multicore computer architectures