Integrating Deep Learning and Statistical Methods for Cryptocurrency Price Forecasting: The case of Ripple

MSc Research Project
Master of Science in Fintech

John Guevara Rincón
Student ID: 23409517

School of Computing
National College of Ireland

Supervisor: Brian Byrne

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name:	John Fredy Guevara Rincón		
Student ID:	23409517		
Programme:	MSc in Fintech	**Year:**	2025
Module:	Msc Research Project		
Supervisor:	Brian Byrne		
Submission Due Date:	11th August 2025		
Project Title:	Integrating Deep Learning and Statistical Methods for Cryptocurrency Price Forecasting: The case of Ripple		

Word Count: 5931 **Page Count:** 17

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature: John Guevara Rincón

Date: 11th August 2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project	□

is lost or mislaid. It is not sufficient to keep a copy on computer.	

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Integrating Deep Learning and Statistical Methods for Cryptocurrency Price Forecasting: The case of Ripple

John Guevara Rincón
23409517

Abstract

The cryptocurrency market, and especially digital assets like XRP (Ripple), are becoming more popular and gaining greater acceptance in the market every day. One of its main characteristics is volatility and complex price dynamics, which pose significant challenges for predicting this market. This study presents and compares two distinct approaches to training neural networks: the autoregressive integrated moving average (ARIMA) statistical model and the long short-term memory (LSTM) deep learning architecture. By evaluating the strengths and limitations of each method, this study aims to provide a deeper understanding of how statistical and deep learning techniques can be applied to obtain findings that can guide model selection strategies and establish a methodology for future research with a solid methodological and practical foundation in financial markets.

Keywords: Ripple, Cryptocurrency, ARIMA, Long Short-Term Memory, Deep Learning.

1. Introduction

The financial landscape is undergoing an accelerated and fundamental transformation. In recent years, cryptocurrencies have been the financial system's greatest innovation in decades, becoming a significant force in conventional finance, challenging traditional monetary systems, and creating new paradigms in financial models. This new digital financial era is being driven primarily by blockchain technology, which is revolutionary not only in the various forms of money flow but also in questioning fundamental economic models such as money, trust, and financial intermediation.

There are thousands of cryptocurrencies in circulation today, each tailored to a specific market and need, where competition for relevance and adoption is fundamental to their growth. Ripple (XRP) offers added value. Unlike the mother of cryptocurrencies (Bitcoin), which is more focused on being a unit of change or Ethereum, which is geared towards smart contracts, Ripple aims to be a bridge between the old and the new financial worlds. Ripple's main objective (as a project) is to revolutionise cross border payments by having financial institutions worldwide adopt its cryptocurrency. This approach and project, aligned with Ripple's innovations and the implications of its consensus protocol and settlement speed, make it an important topic to analyse its value as a project and future as a cryptocurrency.

However, it is the characteristics of cryptocurrencies that, while opening up new fields of use, also present challenges for their application, investment, and analysis. The evident volatility of this market can produce significant fluctuations in relatively short periods, thus requiring more innovative forecasting tools that can adapt to these changes. Traditional forecasting models used to predict stocks or other traditional financial assets typically rely on historical data and relevant patterns, making it more complex to capture the characteristics of digital asset markets. This volatility not only reflects the fluctuation in the currency's value, but also the market's sensitivity, which is affected by regulatory or political announcements, technological developments, macroeconomic factors, and more, making it more complex to understand and analyse. Predicting the price of XRP goes beyond cryptocurrency volatility; its value is tied to government regulations, adoption by financial institutions, and, most importantly, the use of traditional payment systems such as SWIFT. These challenges create a prediction problem that, in addition to recognising changes in XRP supply and demand, also requires interpretability capabilities.

This research encompasses several models that can be used and have already been investigated in previous studies; however, few have focused specifically on Ripple despite its market position. Additionally, existing research employs relatively statistical methods or isolated machine learning techniques. However, this analysis focusses on a comparison of long short-term memory (LSTM) networks recognized for their ability to capture temporal dependencies and ARIMA models, in which statistical rigour and interpretability are more pronounced.

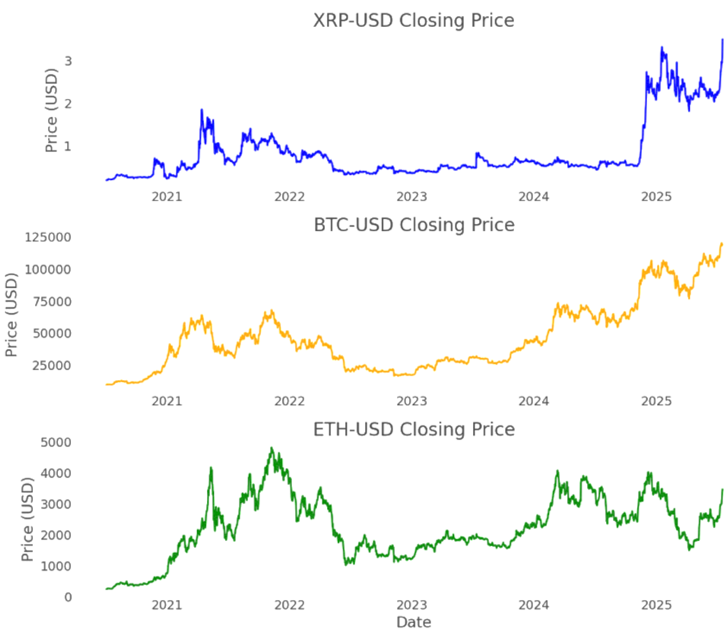

Figure 1 : Closing Price over 5 years (XRP, BTC, ETH)

Accurately forecasting the price of XRP goes beyond the academic interest of this analysis. The model can serve as a basis for risk management portfolio decisions in economies that are shifting towards digital. Similarly, financial institutions that are beginning to implement blockchain technology need more robust tools to evaluate metrics and manage exposure. Likewise, policymakers and legal frameworks for regulating crypto-assets must understand the benefits and dynamics of this technology, helping to design frameworks for these emerging assets. Therefore, having multiple tools for cryptocurrency evaluation helps improve understanding and regulatory capabilities, which in turn enhances their assessment, regulation, and risk management.

This research focusses on a 30-day forecast, a time frame that balances practical utility with technical feasibility short enough for timely trend-based decision making, yet long enough to mitigate daily noise and intrinsic volatility. At the same time, the study provides an understanding of the evidence regarding the effectiveness of different approaches to Ripple, one of the most institutionally orientated cryptocurrencies on the market. The objective of this research is to use the methodology and analysis, including a comparison of research models, contribute to the theoretical understanding of price prediction for this type of asset. The findings will reveal the dynamics of XRP price movements, while also contributing to the debate on how machine learning can be leveraged to address the volatility and complexity of digital asset markets.

2. Related Work

The study of cryptocurrency prediction has become a significant area of research at the intersection of finance, technology, and econometrics. More accurate models for handling volatility in assets such as cryptocurrencies, specifically XRP (Ripple), can integrate various features from traditional existing approaches. A review of the structure surrounding the evolution of cryptocurrency prediction models, comparative deep learning performance, and the integration of two different training models will enable the analysis and identification of key factors for establishing theoretical foundations in the comparison of deep learning models for cryptocurrency prediction.

2.1 Evolution of cryptocurrency price prediction models.

2.1.1 Comparison and correlation of deep learning models.

Classical time series models in digital asset markets represent an important facet of research and price prediction. (Zhang, 2023) presents an analysis of ARIMA's performance on several cryptocurrencies, including Bitcoin and Ripple. The study delves deeper, showing that while ARIMA exhibits competencies in capturing data and linear trends during stable periods, its performance does not train well when volatility is high a relevant characteristic in the cryptocurrency market. Likewise, Zhang proposes that Ripple exhibits relatively better ARIMA model accuracy compared to other cryptocurrencies such as Bitcoin or Ethereum, a result directly influenced by the volatility of the aforementioned cryptocurrencies.

A deeper insight into the limitations of the ARIMA model can be found in the study by (Derbentsev, et al., 2023) who propose a hybrid BART (Bayesian Additive Regression Trees) approach in which the main challenge is nonlinearity. The implementation within the CRISP-DM[1] methodology demonstrates superior performance compared to standard ARIMA models, particularly in capturing regime shifts or structural breaks in cryptocurrency time series. The BART model's ability to train on time series allows it to maintain interpretability, with the model's intermediate stage being purely statistical. Despite your use of statistical precision, the model requires excessive computational intensity, which casts doubt on its practical implementation in a real-time prediction system.

[1] Cross-Industry Standard Process for Data Mining - https://www.ibm.com/docs/sr/spss-modeler/saas?topic=dm-crisp-help-overview

2.1.2 Transition to machine learning

The application of machine learning models to cryptocurrency prediction has unlocked a wealth of significant information regarding the comparative performance of different neural network architectures. In their study, Pakizeh proposes LSTM and GRU models to predict the prices of four cryptocurrencies (Bitcoin, Ethereum, Litecoin, and Ripple) and incorporates variables such as technical indicators or time lags using Random Forest. (Pakizeh, et al., 2022)The architectural design of three-layer deep neural networks regularises and prevents overfitting, in which the author trained six years of daily cryptocurrency data. The results revealed that LSTM models achieved better predictive accuracy compared to models such as GRU.

Ensemble models open a new field in the study proposed by (Nayeem Pinky & Akula, 2023), offering hybrid approaches such as combinations of machine learning algorithms that mitigate the weaknesses of individual forecasting models. This analysis excels in balancing interpretability, which is an important factor in the study to be conducted with LSTM and ARIMA. (Johnson, 2024) The authors demonstrate that the methods, when combined, can achieve higher predictive accuracy; however, interpretability may diminish their usefulness in understanding market dynamics, which is an essential consideration in implementing prediction models.

2.2 Deep learning architectures in cryptocurrency prediction

2.2.1 LSTM and temporal dependencies

The applicability of long short-term memory (LSTM) networks in cryptocurrency prediction has gained momentum and importance due to the model's ability to capture long-term temporal dependencies. (Ranjan, et al., 2023) compares LSTM, GRU, and Bi-LSTM architectures for cryptocurrencies. The model results show that the Bi-LSTM model provides improvements compared to the LSTM model, as evidenced by its impact on prices. Particularly, the models failed to capture the impact of future events, which is attributed to the anticipatory nature of cryptocurrency markets. The study's reliance on price data alone, without accounting for external variables, introduces limitations in the model's applicability and in the research.

The theory behind LSTMs in time series forecasting is presented in greater detail by (Mahapatro, et al., 2024), who propose a hybrid CNN–LSTM architecture to enhance attention mechanisms. In this way, the integration of some technical indicators, such as RSI, MACD, or MFI, with price data is trained using CNN-based feature extraction layers before being processed by LSTM. The mechanism for weighting different time series offers interpretability benefits that are typically absent in deep learning models (Singh, 2024). Improvements of only a few percentage points in prediction accuracy show that independent LSTM models for XRP reflect the potential value of architectural innovations; however, the model's robustness under different market conditions remains uncertain.

2.2.2 Challenges and limitations of deep learning approaches

Machine learning and deep learning models offer many advantages in applicability and scope, however, they also face several practical challenges in cryptocurrency prediction (Ali, et al., 2023). Their analysis shows that the "black box" nature of neural networks raises concerns about limited interpretability for adoption in regulated financial systems. Incorporating attention layers and generating prediction confidence intervals are presented as possible solutions. At the same time, (Chen, 2024) the computational cost of implementing them is another impediment. This trade-off between interpretability and computational efficiency is a challenge not only in the core area of study but also in the inclusion of models such as ARIMA, which is a more interpretable model that will be used in this comparative study.

Deep learning models such as LSTMs exhibit the requirements proposed by (Bouteska, et al., 2024) where training datasets larger than those used by traditional methods achieve better performance, particularly for cryptocurrencies with limited historical data. (Martinez, 2024)The analysis of the relationship between training and prediction accuracy across different cryptocurrencies reveals that models trained on historical data become less accurate as the data ages. This finding has implications for the data collection strategy in the proposed study and its interpretability.

2.3 Integration of variables and market sentiment

2.3.1 Sentiment Analysis in Price Prediction

Incorporating other external market factors into the training of cryptocurrency prediction models represents significant methodological value. (Noorani & Mehrdoust, 2024) demonstrate the integration of sentiment analysis with neural network models for cryptocurrency price prediction. The multisectoral sentiment approach combines news articles or digital trends, demonstrating how it can improve predictions by up to 20% during periods of higher volatility. Despite this, there are data quality challenges in sentiment analysis, so the network can be manipulated by bots or fraudulent campaigns that introduce noise into the model training. (Gorton, 2024)

The dynamics of the impact of sentiment on the market are also reflected in the information value (IV) thresholds (Mahapatro, et al., 2024), which are used to calculate statistical sentiment features. This sentiment indicator shows how its predictive power varies across time horizons, exerting a stronger influence on short-term forecasts than on time-series models based on historical data. However, it was shown market sentiment in forecasting analysis and even more in cryptocurrency market, does not reflect a considerable improvement in the train of the model.

2.3.2 Macroeconomic factors and intermarket dependencies

Macroeconomic variables are increasingly influencing cryptocurrency prices and attracting greater attention every day due to the integration of new market trends and the digital era, including financial markets. (Zhang, 2023) demonstrates, through correlations between Federal Reserve policy announcements and cryptocurrency price movements, how Ripple in particular exhibits an interesting sensitivity to banking events. This discovery relates to XRP's positioning as a bridge between traditional finance and blockchain technology, which in turn demonstrates that predictive models should take into account conventional financial indicators as well as metrics specific to crypto markets. (Zhao, 2023)

Several layers of complexity are evident in the cryptocurrency ecosystem, (Pakizeh, et al., 2022) show how incorporating Bitcoin or Ethereum price movements as features in precision-enhancing models reflects the interconnected nature of cryptocurrency markets, which is also reflected in XRP prediction. The use of Granger causality tests to identify lagged leadership relationships reflects a methodological framework for feature selection in research.

2.3.3 Evaluation metrics and performance assessment

The metrics for evaluating prediction models in crypto markets require special considerations. (Derbentsev, et al., 2023) argue that traditional metrics such as RMSE or MAE may not fully capture model performance in highly volatile markets, and instead propose a combination of directional accuracy and profitability. Similarly, the author's approach to evaluating models dictates that prediction from both statistical and economic perspectives incorporates traditional evaluation methodologies. (Ranjan, et al., 2023) also find an overfitting problem in which deep learning models tend to show outstanding out-of-sample results but not in-sample. The cross-validation strategy emerges as a design approach for time series data with abrupt changes.

3. Research Methodology

This section examines the comparative performance of deep learning and statistical methods—specifically Long Short-Term Memory (LSTM) networks and Autoregressive Integrated Moving Average (ARIMA)—in predicting the movement of XRP (Ripple). This cryptocurrency stands out for its usability and computational efficiency compared to other digital assets. Due to the volatile nature of the cryptocurrency market, precise research is necessary for making investment decisions and, above all, for proper risk management. This methodology integrates historical cryptocurrency data, a comparison between Bitcoin, Ethereum, and Ripple, and deep learning-derived technical indicators—as in a more interpretable linear classification model—making it possible to compare the model under similar data conditions.

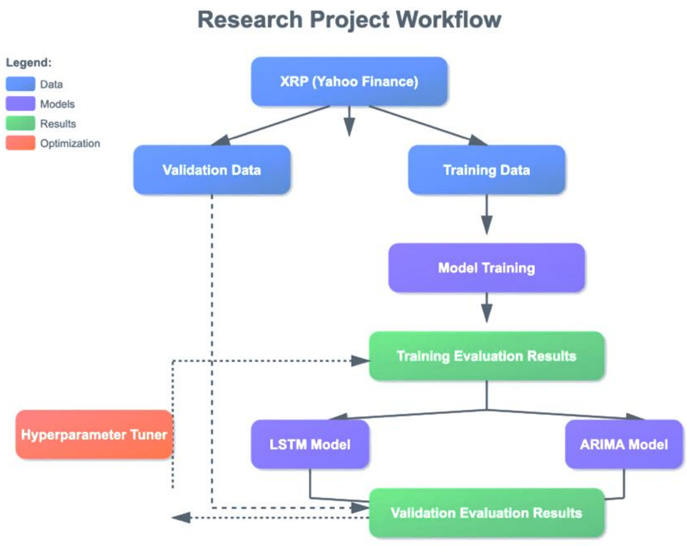

Figure 2: Research Project Workflow

Data preprocessing involves several stages in which missing values were identified, in order to prevent distortions and overfitting of the model. Likewise, the price series are normalised using a Min-Max scale to improve the numerical stability of the neural network learning. Similarly, time series segmentation is performed to generate training sequences, primarily for the LSTM model, which requires fixed-length windows for training. The calculation of metrics such as the Relative Strength Index (RSI). Moving Average Convergence Divergence (MACD) and moving averages for each cryptocurrency (10, 20, 50). These indicators are based on the predictive value documented in previous research, as well as on the predictive window used to capture market signals or trends.

3.1.2 Data acquisition and Preprocessing

Data quality ensures proper training and the interpretability of the models used, as well as the production of models with accurate and reliable forecasts. During this phase of the project, historical price data for the cryptocurrencies Ripple (XRP), Bitcoin (BTC), and Ethereum (ETH) were obtained using Python libraries such as *yfinance*, which provides direct access to verified market records via Yahoo Finance. This dataset covers the period from January 1st, 2018, to current date, and captures multiple market phases in which cryptocurrency volatility has always been present, showing growth spurts, corrections or consolidation periods.

This phase begins with an analysis and comparison of the behaviour of the three most representative cryptocurrencies in the digital market. From this, daily records were selected to balance short-term fluctuations and long-term patterns, resulting in a more robust and computationally efficient dataset. The observations in the dataset include the following features: Date, Opening Price, High Price, Low Price, Closing Price, Volume, and company name, all analysed in U.S. dollars (USD).

Table 1: Dataset Description

Feature	**Description**	**Data Type**
Date	The calendar date corresponding to the market data record	Nominal
Open	The opening price of XRP in USD at the start of the trading day.	Numeric
High	The highest price reached by XRP during the trading day.	Numeric
Low	The lowest price reached by XRP during the trading day.	Numeric
Close	The closing price of XRP in USD at the end of the trading day.	Numeric
Volume	The total number of XRP units traded during the trading day.	Numeric
Company Name	The name of the company who own the project	Object

The LSTM model's preprocessing focused on restructuring the training data. Closing prices were normalised to the range (0,1) using Min-Max, thereby improving training stability. The series normalisation was restructured into input-output pairs using a 30-day observation window, where each training sample showing the closing price was considered the input. The output price was the closing price of the following day. This approach helped the model learn temporal dependencies from three months of historical data.

For ARIMA, the primary step in preprocessing the data is to ensure stationarity; to this end, the Augmented Dickey-Fuller (ADF) test was applied to the closing price series. Due to the asset type, cryptocurrencies typically exhibit non-stationarity, making it necessary to calculate first-order differences to remove trends. Similarly, the optimal order of differencing was applied to ensure stationarity without loss of information, eliminating the need for logarithmic transformations since the variance remained stable for the model.

4. Design Specification

Forecasting models, with appropriate training data selection, are a fundamental step in accurately modelling the precision and interpretability of cryptocurrency projections. In this research, the focus is on deep learning methods and statistical models, using a statistical time series model such as LSTM (Long Short-Term Memory) as well as the ARIMA (Autoregressive Integrated Moving Average) model. These two approaches allow for a comparative evaluation of how different methods respond to the unique characteristics of the data, especially in the cryptocurrency market, in this case XRP (Ripple).

The selection of LSTM is due to its ability to capture nonlinear dependencies and long-term temporal relationships in sequential data, two important characteristics in financial markets, especially in markets with volatility patterns such as the digital asset or cryptocurrency market. On the other hand, ARIMA is a robust and interpretable statistical model whose ability to capture linear trends in stationary data is especially useful in changing markets. By combining these two types of models, the research charts a course for exploring a more complex, computationally intensive LSTM training model, in which the strength of the training is well complemented by the mathematical foundations of the ARIMA model.

4.1.1 Long Short-Term Memory (LSTM)

The Long Short-Term Memory (LSTM) network is a special deep learning model designed to address the vanishing and exploding gradient problems found in traditional Recurrent Neural Networks (RNNs). This model has the ability to control the flow of information through the internal state of network cells. This architecture helps the model retain more relevant information across time series, discarding irrelevant patterns.[2]

[2] "On the vanishing and exploding gradient problem in Gated Recurrent Units" - Alexander Rehmer: https://www.sciencedirect.com/science/article/pii/S2405896320317481

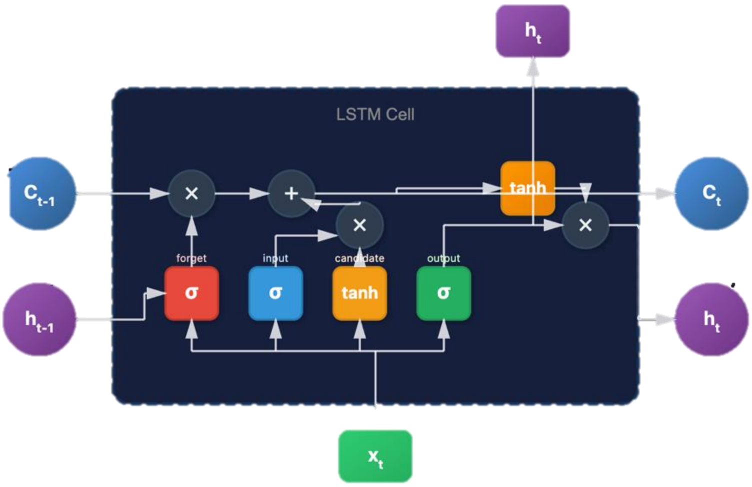

Figure 3: LSTM Architecture

In this study, a three-layer LSTM model is applied, with 128, 64, and 64 units stacked, each incorporating a dropout layer to mitigate overfitting—a common issue with non-stationary data, especially in highly volatile markets such as the cryptocurrency market. The output layer is trained with a single dense neurone using a linear activation function to produce continuous projections. Similarly, an optimisation model (Adam) was applied with a learning rate of 0.001 and a loss function interpreted as mean squared error (MSE). This training requires computational power, so a maximum of 100 epochs was set, with early stopping to prevent overfitting and validation data loss.

4.1.2 Autoregressive Integrated Moving Average (ARIMA)

The ARIMA model offers a broad window into the analysis of linear models, used in time series forecasting due to their interpretability and precise theoretical foundation. The three main components of this model are: the autoregressive (AR) part, which models the relationship between observations and lagged values; the integrated (I) part, which applies differencing to achieve stationarity; and, finally, the moving average (MA) part, which models the relationship between the data and past forecast errors.

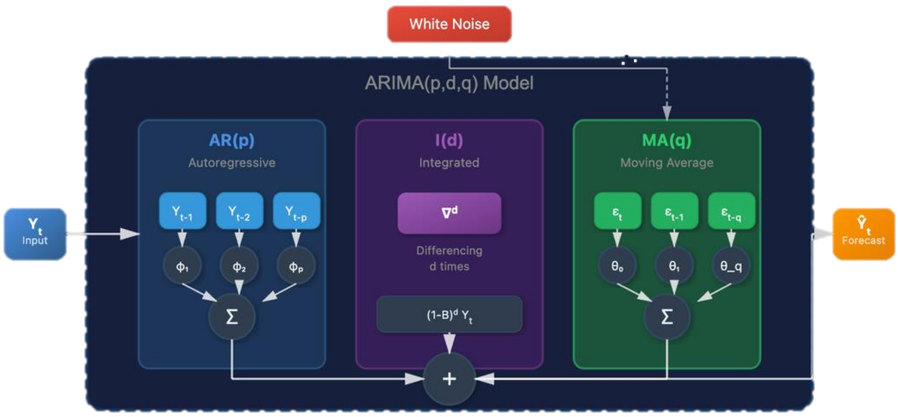

Figure 4: ARIMA Architecture

The choice of model is based on the plots of the autocorrelation function (ACF) and the partial autocorrelation function (PACF). Similarly, several combinations of p, d, and q were tested, where ‘p denotes the number of autoregressive terms, d is the number of non-seasonal differences required for stationarity, and q is the number of lagged forecast errors in the prediction’[3]. These were selected in a configuration to improve performance by balancing goodness of fit and parsimony. Similarly, residual diagnostic tests were performed to ensure that the residuals were not autocorrelated, demonstrating that the ARIMA model is suitable for the linear structure of cryptocurrency forecasting, especially for XRP prices.

4.1.3 Model Evaluation

To evaluate the performance of the predictions for both the LSTM and the ARIMA model, various statistical evaluation metrics and diagnostic analyses were employed. The purpose of the evaluation is to quantify the predictive accuracy of the models while also enabling a critical comparison of the two approaches.

The selected performance metrics were: Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and R-squared (R^2). The perspectives provided by these metrics complement the model’s performance very well, and they offer a different emphasis on the magnitude and nature of the prediction errors. The evaluations were conducted in conjunction with tests to ensure the validity and interpretability of the results.

- **Mean Absolute Error (MAE):** It calculates the average of the absolute errors between the predicted and actual values. Low absolute error values indicate better performance. It is easy to interpret and less sensitive to deviations.

- **Mean Squared Error (MSE):** The squared error shows the differences between the actual values and the predicted values. Values close to 0 indicate better model performance.

- **Root Mean Squared Error (RMSE):** This measurement metric is similar to MSE; however, the error calculation is more closely scrutinised, making it particularly relevant in volatile markets such as the cryptocurrency market.

- **R-squared (R^2):** This metric helps determine the model’s dependence on variables. Here, values close to 1 indicate a model in proper, well-fitted conditions.

In addition to the metrics used for model analysis, visual diagnostics were implemented in the qualitative evaluation of the model’s accuracy. The analysis of predicted price charts, contrasted with actual prices, helps to determine whether the models can capture trends, seasonality, or volatility in the XRP market. Regarding the ARIMA model, residual analyses are conducted to confirm the absence of significant autocorrelation, thereby validating the model’s suitability for capturing linear components. On the other hand, LSTM metrics for overfitting prevention help mitigate validation loss during training.

[3] Introduction to Arima - https://people.duke.edu/~rnau/411arim.htm

5. Implementation

The main reason for creating this project is to develop a predictive model for XRP (Ripple) prices by combining deep learning techniques and statistical methods, with a focus on time series. The choice of models is based on LSTM's time-series modelling capabilities and accuracy, as well as ARIMA's ability to capture nonlinear patterns and volatility, which is important for studying and predicting the cryptocurrency market. Beyond predicting the price of XRP, the primary objective is to showcase market behaviour and analyse, through deep learning, how the trend of a market as volatile as cryptocurrencies can be projected.

This project uses historical price data for Ripple, Bitcoin, and Ethereum as the basis for comparison, evaluation, and model training. The implementation phase shows the transition from the methodology to the execution of the research. Here, through a direct comparison of the main cryptocurrencies and two different training models, we aim to forecast XRP prices and compare them.

5.1.1. Data analysis

The dataset selected for this research contains historical market daily price data for XRP (Ripple), BTC (Bitcoin), and ETH (Ethereum) against the U.S. dollar. The data were collected from the yfinance library, taking data from January 1st, 2018, to the current date. The model's dataset includes opening, high, low, and closing prices, as well as trading volume for each day of the analysed period, totalling approximately 2,270 observations. Through exploratory data analysis (EDA), descriptive statistical calculations are performed, revealing trends and volatilities; similarly, correlations with the other cryptocurrencies studied are evaluated, enabling the identification of patterns, variability levels, and relationships.

On the other hand, the integrity and consistency of the series were verified, allowing us to discard any missing or atypical values that could disrupt the model's training. Following this, the dataset is split into three subsets: training (70%), validation (15%), and test (15%), thereby helping to prevent future information leakage. In the analysis, moving averages and daily logarithmic returns were taken into account, helping to detect cyclical movements or structural changes in the model.

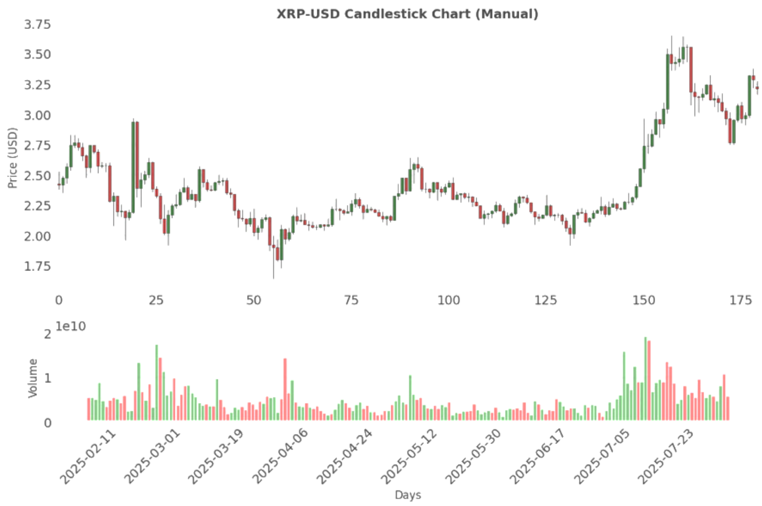

Figure 5: Candlestick chart - XRP-USD Price

In the below graph, the performance of the percentage changes in each of the analysed cryptocurrencies is reviewed over a five-year period. This highlights how the cryptocurrency market is highly volatile, while also comparing the three cryptocurrencies, notably showing that XRP has experienced a significant and distinctive increase over the past year. This served as the basis for training the model; due to the substantial rise in XRP's value, training over very short periods would introduce more errors into the data sample or make it more prone to overfitting.

Figure 6: Cryptocurrency Performance Comparison (% Change)

The implementation was based on two complementary approaches: a neural network model LSTM and a statistical model ARIMA. Time series forecasting was applied to both, preserving the temporal sequence of the data. Each model was trained individually, optimising hyperparameters for each one. In the LSTM model, the window size, number of layers, and regularisation parameters were adjusted; in contrast, the ARIMA model's optimal (p, d, q) order was selected based on an ACF/PACF[4] analysis. These types of evaluation were performed on the test set using metrics such as absolute error, squared error, and R-squared, as well as graphical representations comparing actual and estimated values.

5.1.2 LSTM Model

The LSTM training is aimed at improving the time series of XRP closing prices; primarily, the ranges given as [0, 1][0, 1][0, 1] are normalised using Min-Max scaling. Defining the forecast objective as the basis for the projection is essential, thus, a 30-day window is set as the input horizon, providing the three-dimensional temporal sequences needed to train recurrent neural networks. A three-layer architecture was implemented, with 50 units each, regularised with 20% dropout to reduce overfitting. A dense output layer is configured to estimate normalised price values. Similarly, the Adam technique and the squared error are used to mitigate data loss.

The training and validation datasets are kept in temporal order. During training, techniques to prevent overfitting are applied, namely early stopping and ReduceLROnPlateau callbacks, to mitigate stagnation and validation loss. The process was run with 50 epochs and a batch size

[4] ACF/PACF (Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) - Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF)

of 42. Additionally, different metrics were used to validate the training and assess the accuracy of pattern capture in the price dynamics.

5.1.3 ARIMA Model

In the ARIMA model, closing prices were transformed using logarithms to stabilise variance, and stationarity was adjusted via the Augmented Dickey-Fuller (ADF) test by applying differencing until p-values below 0.05 were obtained, indicating that the model is non-stationary. The parameters p and q were determined using PACF autocorrelations, identifying that the best configuration for the ARIMA model was (1,0,1).

The final adjustments were made through statistical significance evaluations of coefficients and residual diagnostic tests; these were used for validation, which also employed a walk-forward scheme, cumulatively retraining values. The performance, like that of the previous model, was evaluated using metrics such as RMSE, MAE, MAPE, and R^2.

6. Evaluation

The performance evaluation of the statistical and deep learning models implemented in this research examines two predictive approaches with recurrent neural network architectures. For this evaluation, time series forecasting metrics such as: Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), Coefficient of Determination (R^2) and Directional Accuracy. Studies use these metrics to assess the effectiveness and efficiency of a model. An R^2 value close to one indicates a perfect correlation between the observed and predicted values.

5.2 Quantitative Evaluation of the Models

Figure 7: LSTM vs ARIMA Performance Comparison

5.3 Performance of Deep Learning Models

Graph 7 shows the results obtained for each model, with ARIMA achieving an R^2 value of 0.9912, thus demonstrating an almost perfect correlation between observed and estimated values. On the other hand, the model's performance in absolute metrics was superior to that of the LSTM model, reaching a MAE of 0.0471, RMSE of 0.0910, and MAPE of 2.91%. The LSTM model, on the other hand, although its performance was lower than that of the ARIMA model, achieved an R^2 value of 0.9234, indicating that the model has a viable correlation and that the trained values are 92% close to the estimated ones. In other metrics, the model achieved MAE = 0.1735, RMSE = 0.2643, and MAPE = 8.76%, demonstrating stable performance given the inherent volatility of the cryptocurrency market. In terms of directional accuracy, both models showed similar results at 47%, indicating that despite model training, predicting the direction of movement remains a challenge in the volatile cryptocurrency market.

Figure 8: LSTM Model – XRP Price Prediction

Long Short-Term Memory model, the deep learning model analysed, successfully captures the main movements and exhibits slight deviations at inflection points. However, in the analysed series from 2018 to 2025, the training and validation datasets show that the XRP price exhibits high volatility, including notable peaks in 2018 and again in 2025. Although the model copes with and follows the general trend and fluctuations, it is clear that the training is consistent in short-term scenarios.

The model exhibits certain limitations at inflection points or during periods of heightened volatility, suggesting that the training has some difficulty anticipating abrupt shifts in the trend. This, in turn, is reflected in the quantitative data, where the mean absolute percentage error (MAPE) is 8.76 and the directional accuracy is 47%, showing that although the model and training are correct, approximating an exact numerical value is not as effective.

Figure 9: ARIMA Model – XRP Price Prediction

In Figure 9, the XRP price projection using an ARIMA model is shown; the model closely tracks both the trend and the magnitude of price movements, demonstrating how the adjustments made can help train it in volatile conditions. Likewise, the metrics obtained show that the ARIMA model exhibits a mean absolute percentage error (MAPE) of 2.91%, which is essential for minimising prediction error. This is corroborated by the R^2 value, demonstrating the high correlation between the trained data.

On the other hand, although ARIMA exhibits high numerical accuracy in the analysed metrics, its directional accuracy remains relatively low, indicating—just as the LSTM model does—that the cryptocurrency market still presents challenges in managing and predicting volatility. ARIMA is a clear model for short-term projections with absolute values and high predictive power; however, hybrid approaches or analysis incorporating additional variables may perform better in capturing market trends.

6.2 Discussion

The comparative analysis of deep learning models and statistical methods reveals substantial differences in predictive performance, driven by factors such as the nature of the time series data and the architecture of each model. Both the ARIMA and LSTM models performed considerably well; in quantitative terms, ARIMA shows an R^2 value of 0.9912 and a MAPE of 2.91%, indicating an almost perfect correlation with the actual series and a low relative error.

Thus, it can be seen that during the study period, the XRP closing price exhibits linear components and autocorrelation patterns that the ARIMA model can effectively capture. On the other hand, the LSTM model shows an R^2 value of 0.9234 and a MAPE of 8.76%, indicating relatively good performance, but its robustness is lower than that of the ARIMA model. It is evident that the LSTM model's architecture helps it model complex sequences and adapt to changes in temporal dynamics; however, its effectiveness is limited in volatile markets such as cryptocurrencies, where linear patterns dominate.

Table 2: Model Performance Comparison

MODEL PERFORMANCE COMPARISON

Metric	LSTM	ARIMA	Better Model
RMSE	$0.2643	$0.0910	ARIMA
MAE	$0.1735	$0.0471	ARIMA
MAPE (%)	8.76	2.91	ARIMA
R² Score	0.9234	0.9912	ARIMA
Directional Accuracy (%)	47.00	47.21	ARIMA

DETAILED ANALYSIS

ARIMA Performance Advantage:

- RMSE: ARIMA is 65.6% better
- MAE: ARIMA is 72.9% better
- MAPE: ARIMA is 66.8% better
- R² Score: ARIMA is 7.3% better
- Dir. Accuracy: ARIMA is 0.4% better

Directional accuracy is a relevant metric in this study; although metrics such as the R^2 value indicate effective performance and training, directional accuracy was the metric that both models struggled to achieve robustness on, showing that predicting the direction of market trends or volatility remains a challenge for training models. This also highlights the need to compare with other metrics that can help mitigate impacts in volatile markets such as cryptocurrencies, which are influenced by external factors.

In this way, the results of this research show there is no single approach to forecasting financial time series, especially with cryptocurrencies markets. Models such as ARIMA prove more efficient in situations where linear or stationary patterns predominate, while the LSTM model exhibits greater flexibility and potential in training with multivariate or hybrid schemes, where exogenous variables can be analysed. Thus, this performance underscores the importance of adapting methodologies to the characteristics of the series, helping to improve prediction and the objectives under analysis.

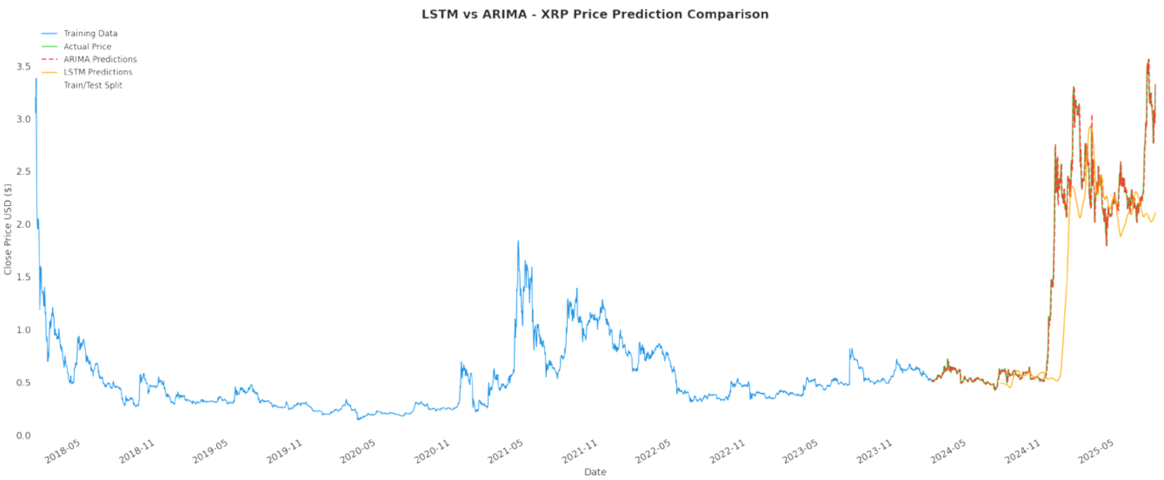

Figure 11: LSTM vs ARIMA – XRP Price Prediction Comparison

Finally, the results obtained show how deep learning models and statistical models can benefit from being combined. Hybrid models that integrate capabilities such as ARIMA to capture trends and linear components, along with deep learning models that handle long-term dependencies well, could help improve precision, such as the magnitude of accuracy. Similarly, analysing technical indicators, volume data, or even sentiment metrics can contribute to more robust model training in high-volatility scenarios.

7. Conclusion and Future Work

This research focused on analysing the behaviour of two training models Deep Learning and Statistical Model for forecasting the price of XRP (Ripple) by implementing the ARIMA and Long Short-Term Memory (LSTM) models. The results show a clear superior performance in the quantitative metrics for the ARIMA model (R^2 value 0.992), demonstrating a better correlation with the actual time series. This is not to detract from the LSTM model's performance, where its correlation and training score were also quite high. Both models achieved similar Directional Accuracy (47%), which highlights an opportunity to improve the training of such models in high-volatility market conditions.

This analysis also enabled the development of competencies in preparing time series data, fitting models, and conducting comparative analyses using statistical approaches. It was observed that the nature of the data plays an important role in model selection, with ARIMA proving more consistent for linear and stationary components, while LSTM demonstrates better performance on series from multiple information sources, adapting more effectively to the market. Additionally, direct comparisons of the two models reveal approaches that serve as a reference for selecting methodologies in similar projects.

However, enriching the training features and incorporating exogenous variables could serve as future work, helping to improve technical indicators and, on the other hand, including sentiment analysis or macroeconomic factors. Although the cryptocurrency market is uncertain. Models such as ARIMA and LSTM could help enhance the accuracy and stability of predictions if trained in conjunction with exogenous variables. In this way, not only will the accuracy of forecasts be strengthened, but the practical use of this type of analysis in the financial sector will also be increased.

In conclusion, this research project helps the field of cryptocurrency prediction by providing a rigorous comparative analysis of two types of neural network training, applied to one of the market's most prominent cryptocurrencies. The research, in turn, highlights the importance of data preparation and the appropriate model selection for project adaptation. Thus, this project serves as a foundation for proposing more precise and adaptive solutions in highly volatile financial markets.

8. References

Ahmed, S. R. M. a. K. A., 2023. *'Forecasting cryptocurrency prices using LSTM, GRU, and bi-directional LSTM: A deep learning approach',*, s.l.: Fractal and Fractional.
Ali, F., Suryakant, R. & Nimbore, S., 2023. *Ensemble Model Based on Deep Learning for Forecasting Crypto Asset Futures in Markets,* Karnataka: s.n.
Bouteska, A., Zoynul Abedin, M., Hajek, P. & Yuan, K., 2024. *Cryptocurrency price forecasting – A comparative analysis of ensemble learning and deep learning methods,* Tunis: ELSEVIER.
Chen, L. a. W. Y., 2024. Forecasting cryptocurrency returns using classical statistical and deep learning techniques. *International Journal of Information Management Data Insights,* Volume 4.
Derbentsev, V., Datsenko, N., Stepanenko, O. & Bezkorovainyi, V., 2023. *Forecasting Cryptocurrency Prices Time Series Using Machine Learning,* Kiev: s.n.
Gorton, G., 2024. The Digitalization of Money. In: s.l.:Cambridge Scholars Publishing, pp. 94-99.
Johnson, R. S. T. a. B. K., 2024. Machine learning-based cryptocurrency price prediction models: From LSTM to transformer. *Financial Innovation,* Volume 10, pp. 45-67.
Kumar, A. S. P. a. G. R., 2024. Review of deep learning models for crypto price prediction: Implementation and evaluation. *ArXiv Preprint.*
Li, H. Z. W. a. C. X., 2024. Deep learning for Bitcoin price direction prediction: Models and trading strategies empirically compared. *Financial Innovation,* Volume 10.
Mahapatro, S., Kumar Sahu, P. & Subudhi, A., 2024. *Navigating XRP Volatility: A Deep Learning Perspective on Technical Indicators,* Bhubaneswar: (IJACSA) International Journal of Advanced Computer Science and Applications.
Martinez, J. a. R. C., 2024. Forecasting cryptocurrency prices using recurrent neural network and long short-term memory. *nternational Journal of Advanced Computer Science and Applications,* Volume 15, pp. 234-245.
Nakamoto, S., 2009. *Bitcoin: A Peer-to-Peer Electronic Cash System,* s.l.: Bitcoin.org.
Nayeem Pinky, J. & Akula, R., 2023. *Enhancing Cryptocurrency Market Forecasting: Advanced Machine Learning Techniques and Industrial Engineering Contributions,* Florida: University of Central Florida,.
Noorani, I. & Mehrdoust, F., 2024. *Forecasting bitcoin price by a hybrid structure based on ARIMA, SVM and LSSVM models,* Rasht: University of Guilan.
Pakizeh, K., Malek, A., Karimzadeh, M. & Hamidi Razi, H., 2022. *Assessing machine learning performance in cryptocur- rency market price prediction,* Tehran: Journal of Mathematics and Modeling in Finance.
Ranjan, A. R., Nilakhe, V., Bhatkar, V. & Iinje, B., 2023. *XRP Price Prediction Using Advanced Deep Learning Models,* s.l.: International Journal of Research in Engineering and Science (IJRES).
Singh, A. a. P. N., 2024. Crypto price prediction using LSTM+XGBoost. *ArXiv Preprint.*
Zhang, Z., 2023. *Price Prediction of Bitcoin, Ethereum and XRP Based on the ARMA Model,* Bristol: BCP Business & Management.
Zhao, L. W. S. a. L. H., 2023. Attention! Transformer with sentiment on cryptocurrencies price prediction. *roceedings of the International Conference on Machine Learning and Applications,* pp. 456-467.