


**COPY RIGHT**



**ELSEVIER**  
**SSRN**

**2023IJIEMR**. Personal use of this material is permitted. Permission from IJIEMR must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. No Reprint should be done to this paper, all copy right is authenticated to Paper Authors IJIEMR Transactions, online available on 19th May 2023.

Link : <https://ijiemr.org/downloads/Volume-12/Issue-05>

**10.48047/IJIEMR/V12/ISSUE05/40**

Title **CRYPTOCURRENCY PRICE PREDICTION USING NEURAL NETWORKS, DEEP LEARNING AND MACHINE LEARNING**

Volume12, Issue 05, Pages: 408-417

Paper Authors

**Mrs. Sunitha Surarapu, P.Vasavi, V.Sravani, V.Sree Koumudi, E.Laxmi Narasimha, B.Shivudu**



USE THIS BARCODE TO ACCESS YOUR ONLINE PAPER

To Secure Your Paper As Per **UGC Guidelines** We Are Providing A Electronic Bar Code

## CRYPTOCURRENCY PRICE PREDICTION USING NEURAL NETWORKS, DEEP LEARNING AND MACHINE LEARNING

**Mrs. Sunitha Surarapu**, Assistant professor, Department of CSE, Sreyas Institute of Engineering and Technology

[sunithasurarapu@sreyas.ac.in](mailto:sunithasurarapu@sreyas.ac.in)

**P.Vasavi, V.Sravani, V.Sree Koumudi, E.Laxmi Narasimha, B.Shivudu**, Sreyas Institute of Engineering and Technology

**ABSTRACT:** The goal of this paper is to ascertain with what accuracy the direction of Bitcoin price in USD can be predicted. The price data is sourced from the Bitcoin Price Index. The task is achieved with varying degrees of success through the implementation of a Bayesian optimized recurrent neural network (RNN) and a Long Short Term Memory (LSTM) network. The LSTM achieves the highest classification accuracy of 52% and a RMSE of 8%. The popular ARIMA model for time series forecasting is implemented as a comparison to the deep learning models. As expected, the non-linear deep learning methods outperform the ARIMA forecast which performs poorly. Finally, both deep learning models are benchmarked on both a GPU and a CPU with the training time on the GPU outperforming the CPU implementation by 67.7%.

**Keywords** – *recurrent neural network (RNN) and a Long Short Term Memory (LSTM) network.*

### 1. INTRODUCTION

Bitcoin is the worlds' most valuable cryptocurrency and is traded on over 40 exchanges worldwide accepting over 30 different currencies. It has a current market capitalization of 9 billion USD according to <https://www.blockchain.info/> and sees over 250,000

transactions taking place per day. As a currency, Bitcoin offers a novel opportunity for price prediction due its relatively young age and resulting volatility, which is far greater than that of fiat currencies It is also unique in relation to traditional fiat currencies in terms of its open nature; no complete data exists regarding cash transactions or money in circulation for fiat currencies. Prediction of mature financial markets such as the stock market has been researched at length. Bitcoin presents an interesting parallel to this as it is a time series prediction problem in a market still in its transient stage. Traditional time series prediction methods such as Holt-Winters exponential smoothing models rely on linear assumptions and require data that can be broken down into trend, seasonal and noise to be effective . This type of methodology is more suitable for a task such as forecasting sales where seasonal effects are present. Due to the lack of seasonality in the Bitcoin market and its high volatility, these methods are not very effective for this task. Given the complexity of the task, deep learning makes for an interesting technological solution based on its performance in similar areas. The recurrent neural network (RNN) and the long short term memory (LSTM) are favoured over the traditional multilayer perceptron (MLP) due to the temporal nature of Bitcoin data.

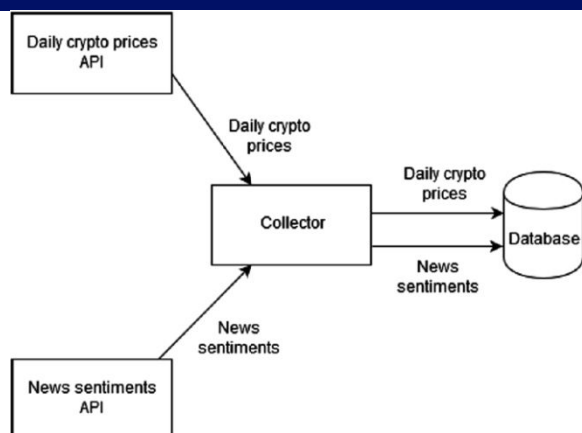


Fig.1: Example figure

The aim of this paper is to investigate with what accuracy the price of Bitcoin can be predicted using machine learning and compare parallelisation methods executed on multi-core and GPU environments. This paper contributes in the following manner: of approximately 653 papers published on Bitcoin, only 7 (at the time of writing) are related to machine learning for prediction. To facilitate a comparison to more traditional approaches in financial forecasting, an ARIMA time series model is also developed for performance comparison purposes with the neural network models. The independent variable for this study is the closing price of Bitcoin in USD taken from the Coindesk Bitcoin Price Index. Rather than focusing on one specific exchange, we take the average price from five major Bitcoin exchanges: Bitstamp, Bitfinex, Coinbase, OkCoin and itBit. If we were to implement trades based on the signals it would be beneficial to focus on just one exchange. To assess the performance of models, we use the root mean squared error (RMSE) of the closing price and further encode the predicted price into categorical variable reflecting: price up, down or no change. This latter step allows for additional

performance metrics that would be useful to a trader in the formation of a trading strategy: classification accuracy, specificity, sensitivity and precision. The dependent variables for this paper come from the Coindesk website, and Blockchain.info. In addition to the closing price, the opening price, daily high and daily low are also included as well as Blockchain data, i.e. the mining difficulty and hash rate. The features which have been engineered (considered as technical analysis indicators) include two simple moving averages (SMA) and a de-noised closing price.

## 2. LITERATURE REVIEW

Research on predicting the price of Bitcoin using machine learning algorithms specifically is lacking. implemented a latent source model as developed by to predict the price of Bitcoin noting 89% return in 50 days with a Sharpe ratio of 4.1. There has also been work using text data from social media platforms and other sources to predict Bitcoin prices. investigated sentiment analysis using support vector machines coupled with the frequency of Wikipedia views, and the network hash rate. investigated the relationship between Bitcoin price, tweets and views for Bitcoin on Google Trends. implemented a similar methodology except instead of predicting Bitcoin price they predicted trading volume using Google Trends views. However, one limitation of such studies is the often small sample size, and propensity for misinformation to spread through various (social) media channels such as Twitter or on message boards such as Reddit, which artificially inflate/deflate prices. In the Bitcoin exchanges liquidity is considerably limited. As a result, the market suffers

from a greater risk of manipulation. For this reason, sentiment from social media is not considered further.

analyzed the Bitcoin Blockchain to predict the price of Bitcoin using support vector machines (SVM) and artificial neural networks (ANN) reporting price direction accuracy of 55% with a regular ANN. They concluded that there was limited predictability in Blockchain data alone. Also used Blockchain data, implementing SVM, Random Forests and Binomial GLM (generalised linear model) noting prediction accuracy of over 97% however without cross-validating their models limiting the generalisability of their results. Wavelets have also been utilised to predict Bitcoin prices, with noting positive correlations between search engine views, network hash rate and mining difficulty with Bitcoin price. Building on these findings, data from the Blockchain, namely hash rate and difficulty are included in the analysis along with data from the major exchanges provided by CoinDesk.

Predicting the price of Bitcoin can be considered analogous to other financial time series prediction tasks such as forex and stock prediction. Several bodies of research have implemented the MultiLayer Perceptron (MLP) for stock price prediction. However, the MLP only analyses one observation at a time. In contrast, the output from each layer in a recurrent neural network (RNN) is stored in a context layer to be looped back in with the output from the next layer. In this sense, the network gains a memory of sorts as opposed to the MLP. The length of the network is known as the temporal window length. Notes that the temporal relationship of the series is explicitly modelled by the internal states contributing significantly to model effectiveness.

successfully took this approach in predicting stock returns combining a RNN with a genetic algorithm for network optimization.

Another form of RNN is the Long Short Term Memory (LSTM) network. They differ from Elman RNN in that in addition to having a memory, they can choose which data to remember and which data to forget based on the weight and importance of that feature. Implemented a LSTM for a time series prediction task finding that the LSTM performed as well as the RNN for this task. This type of model is implemented here also. One limitation in training both the RNN and LSTM is the significant computation required. For example, a network of 50 days is comparable to training 50 individual MLP models. Since the development of the CUDA framework by NVIDIA in 2006, the development of applications that take advantage of the extremely parallel capabilities of the GPU has grown greatly including the area of machine learning. Reported over three times faster training and testing of its ANN model when implemented on a GPU rather than a CPU. Similarly, reported an increased speed in classification time to the magnitude of eighty times when implementing a SVM on a GPU over an alternative SVM algorithm run on a CPU. In addition, training time was nine times greater for the CPU implementation. Also received speeds that were forty times faster for training a when training deep neural networks for image recognition on a GPU as opposed to a CPU. Due to the apparent benefits of utilising a GPU, our RNN and LSTM models are implemented on both the CPU and GPU.



## 3. METHODOLOGY

Prediction of mature financial markets such as the stock market has been researched at length. Bitcoin presents an interesting parallel to this as it is a time series prediction problem in a market still in its transient stage. Traditional time series prediction methods such as Holt-Winters exponential smoothing models rely on linear assumptions and require data that can be broken down into trend, seasonal and noise to be effective. This type of methodology is more suitable for a task such as forecasting sales where seasonal effects are present. Due to the lack of seasonality in the Bitcoin market and its high volatility, these methods are not very effective for this task.

### Existing System

- Bitcoin presents a curiosity to this as it is a time series prediction problem in a market still in its short stage.
- Traditional time series prediction methods such as Holt-Winters exponential smoothing models rely on linear assumptions and require data that can be broken down into trend, seasonal and noise to be effective.
- The prediction of Bitcoin price can be considered correspondent to other financial time series prediction tasks such as forex and stock prediction. Distinct bodies of research have implemented the Multilayer Perceptron (MLP) for stock price prediction. Reported over threetime's faster training and testing of its ANN model when implemented on a GPU rather than a CPU.

### Disadvantages of Existing System:

- As the number of bitcoins in circulation approaches closer to the limit, bitcoins become increasingly harder to mine.
- One of the problem that analysts and researchers faced was to implement a system capable of accurately predicting the prices.

### Solution:

we attempt to predict the Bitcoin price accurately taking into consideration various parameters that affect the Bitcoin value. For the first phase of our investigation, we aim to understand and identify daily trends in the Bitcoin market while gaining insight into optimal features surrounding Bitcoin price. Our data set consists of various features relating to the Bitcoin price and payment network over the course of five years, recorded daily. For the second phase of our investigation, using the available information, we will predict the sign of the daily price change with highest possible accuracy.

### Proposed System:

- The main goal of my project is to investigate with what accuracy the price of Bitcoin can be predicted using machine learning. To help in illustrating more traditional approaches in financial forecasting, an ARIMA time series model is developed for performance illustration purposes with the neural network models.
- Closing price of Bitcoin in USD from the Coin desk Bitcoin Price Index is considered

as a separate variable for this study. I have taken the average price from five major Bitcoin exchanges: Bit stamp, Bitfinex, Coinbase, OkCoin and itBit.

- To estimate the performance of models, we choose root mean squared error (RMSE) of closing price and further put into code the predicted price into categorical variable reflecting: price up, down or no change.
- It allows increased performance in metrics that would be useful to a merchant in the design of a trading approach: classification, accuracy, specificity, sensitivity and precision. Also dependent variables for this paper is taken from the Coindesk, and Blockchain.info. The closing price, the opening price, daily high & low are also included as well as Blockchain data (hash rate).

## Advantages of Proposed System:

- Models possess great potential to change opportunity into revenue.
- High Performance and accuracy.

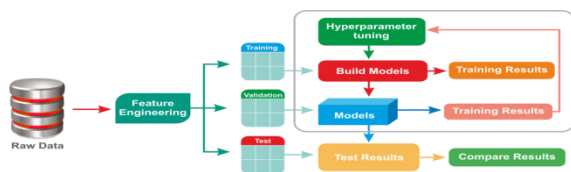


Fig.2: System architecture

## MODULES:

### 1. Strategy: matching the problem with the solution

In the first phase of an ML project realization, we are mostly outline strategic goals. They assume a solution to a problem, define a scope of work, and plan the development.

### 2. Dataset preparation and preprocessing

Data is the foundation for any machine learning project. The second stage of project implementation is complex and involves data collection, selection, preprocessing, and transformation. Each of these phases can be split into several steps.

### 3. Dataset splitting

A dataset used for machine learning should be partitioned into three subsets — training, test, and validation sets.

## 4. IMPLEMENTATION

RNN The first parameter to consider was the temporal length window. As suggested by supporting literature these type of networks may struggle to learn long term dependencies using gradient based optimisation. An autocorrelation function (ACF) was run for the closing price time series to assess the relationship between the current closing price and previous or future closing prices. While this is not a guarantee of predictive power for this length, it was a better choice than random choice. Closing price is correlated with a lag of up to 20 days in many cases, with isolated cases at 34, 45 and 47 days. This led the grid search for the temporal window to test from 2 to 20, 34, 45 and 47 days. To ensure a robust search,

larger time periods of up to 100 days were also tested in increments of five. The most effective window temporal length was 24.

In addition to the temporal window, some hyperparameters also need tuning: Learning rate is the parameter that guides stochastic gradient descent (SGD), i.e. how the network learns. Similarly, momentum updates the learning rate to avoid the model falling into local minima (in terms of error) and attempts to move towards the global minimum of the error function. We used the RMSprop optimiser to improve on SGD, as it keeps a running average of recent gradients and as a result is more robust against information loss. According to Heaton, one hidden layer is enough to approximate the vast majority of non-linear functions. Two hidden layers were also explored and were chosen as they achieved lower validation error. Heaton also recommends for the number of hidden nodes to select between the number of input and output nodes. In this case, less than 20 nodes per layer resulted in poor performance. 50 and 100 nodes were tested with good performance. However, too many nodes can increase the chances of overfitting, and significantly increase the time needed to train the network. As 20 nodes performed sufficiently well this was chosen for the final model. An activation function, a nonlinear stepwise equation that passes signals between layers, is also needed. The options explored were Tanh, ReLu, and Sigmoid. Tanh performed the best but the differences were not significant. The final parameters for selection are batch size and number of training epochs. Batch size was found to have little effect on accuracy but considerable effect on training time when using smaller batches in this case. The number of epochs tested ranged from 10 to 10000, however, too many

training epochs can result in overfitting. To reduce the risk of overfitting, dropout was implemented as discussed above. Optimal dropout between 0.1 and 1 was searched for both layers with .5 dropout the optimal solution for both layers. A Keras callback method was also used to stop the training of the model if its performance on validation data did not improve after 5 epochs to prevent overfitting. Generally, the RNN converged between 20 and 40 epochs with early stopping.

**LSTM** In terms of temporal length, the LSTM is considerably better at learning long term dependencies. As a result, picking a long window was less detrimental for the LSTM. This process followed a similar process to the RNN in which autocorrelation lag was used as a guideline. The LSTM performed poorly on smaller window sizes. Its most effective length found was 100 days, and two hidden LSTM layers were chosen. For a time series task two layers is enough to find nonlinear relationships among the data. 20 hidden nodes were also chosen for both layers as per the RNN model. The Hyperas library<sup>2</sup> was used to implement the Bayesian optimisation of the network parameters. The optimiser searched for the optimal model in terms of how much dropout per layer and which optimizer to use. RMSprop again performed the best for this task. In the LSTM model, activation functions weren't changed as the LSTM has a particular sequence of tanh and sigmoid activation functions for the different gates within the cell. LSTM models converged between 50 and 100 epochs with early stopping. Similar to the RNN, batch size was found to have a greater effect on execution time than accuracy. This may be due to the relatively small size of the dataset.

## Model Comparison

A confusion matrix representing the ratio of true/false and positive/negative classifications is used to derive the ratings metrics. Accuracy can be defined as the total number of correctly classified predictions (price up, down, and no change). To combat inherent class imbalance (bitcoin price predominately increases) the metrics sensitivity, specificity and precision are also analysed. Sensitivity represents how good a model is at detecting positives. Specificity represent how good the model is at avoiding false alarms. Finally, precision represents how many positively classified predictions were relevant. Root Mean Square Error (RMSE) is used to evaluate and compare the regression accuracy. To instrument the evaluation of models, a 80/20 holdout validation strategy is used. In order to facilitate a comparison of the deep learning methods to more traditional methods we built (and optimised) an ARIMA model, as they have been extensively used in price prediction problems. The ARIMA forecast was created by splitting the data into 5 periods and then predicting 30 days into the future. The data was differenced before being fit with several ARIMA models. The best fit was found by auto.arima from the R forecast package.

## 5. EXPERIMENTAL RESULTS

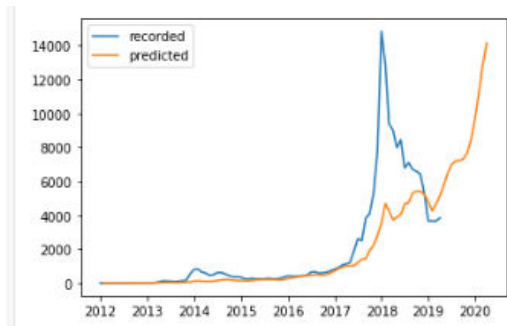


Fig.3: Output screen

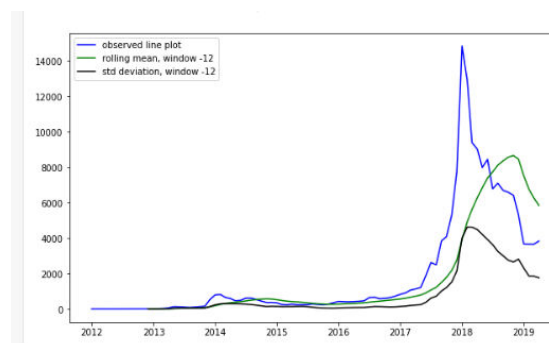


Fig.4: Output screen

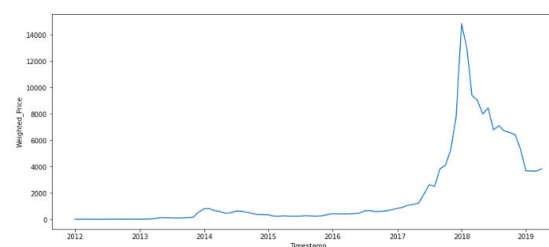


Fig.5: Output screen

LSTM achieved the highest accuracy while the RNN achieved the lowest RMSE. The ARIMA prediction performed poorly in terms of accuracy and RMSE. Upon analysis of the ARIMA forecast, it predicted the price would gradually rise each day. There were no false positives from the model. One reason for this may be due to the class imbalance in predictive portion of the ARIMA forecast (the price tends to always increase). This contributed to the specificity and precision being so high (specificity, precision= 100%). This does not necessarily suggest good overall performance, but rather that it does a decent job at identifying price direction change(s).

## 6. CONCLUSION

Deep learning models such as the RNN and LSTM are evidently effective for Bit coin prediction with the



LSTM more capable for recognizing longer-term dependencies. However, a high variance task of this nature makes it difficult to transpire this into impressive validation results. As a result it remains a difficult task. There is a fine line between over fitting a model and preventing it from learning sufficiently. Dropout is a valuable feature to assist in improving this. However, despite using Bayesian optimization to optimize the selection of dropout it still couldn't guarantee good validation results. Despite the metrics of sensitivity, specificity and precision indicating good performance, the actual performance of the ARIMA forecast based on error was significantly worse than the neural network models. The LSTM outperformed the RNN marginally, but not significantly. However, the LSTM takes considerably longer to train.

The performance benefits gained from the parallelization of machine learning algorithms on a GPU are evident with a 70.7% performance improvement for training the LSTM model. Looking at the task from purely a classification perspective it may be possible to achieve better results. One limitation of the research is that the model has not been implemented in a practical or real time setting for predicting into the future as opposed to learning what has already happened. In addition, the ability to predict using streaming data should improve the model. Sliding window validation is an approach not implemented here but this may be explored as future work. One problem that will arise is that the data is inherently shrouded in noise.

In terms of the dataset, based on an analysis of the weights of the model the difficulty and hash rate variables could be considered for pruning. Deep

learning models require a significant amount of data to learn effectively. If the granularity of data was changed to per minute this would provide 512,640 data points in a year. Data of this nature is not available for the past but is currently being gathered from Coin Desk on a daily basis for future use. Finally, parallelization of algorithms is not limited to GPU devices. Field Programmable Gate Arrays (FPGA) are an interesting alternative to GPU devices in terms of parallelization and machine learning models have been shown to perform better on FPGA than on a GPU.

## REFERENCES

- [1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.
- [2] M. Briere, K. Oosterlinck, and A. Szafarz, "Virtual currency, tangible return: Portfolio diversification with bitcoins," *Tangible Return: Portfolio Diversification with Bitcoins* (September 12, 2013), 2013.
- [3] I. Kaastra and M. Boyd, "Designing a neural network for forecasting financial and economic time series," *Neurocomputing*, vol. 10, no. 3, pp. 215–236, 1996.
- [4] H. White, "Economic prediction using neural networks: The case of IBM daily stock returns," in *Neural Networks, 1988., IEEE International Conference on*, IEEE, 1988, pp. 451–458.
- [5] C. Chatfield and M. Yar, "Holt-winters forecasting: some practical issues," *The Statistician*, pp. 129–140, 1988.

[6] B. Scott, “Bitcoin academic paper database,” suitpossum blog, 2016.

[7] M. D. Rechenstein, “Machine-learning classification techniques for the analysis and prediction of high-frequency stock direction,” 2014.

[8] D. Shah and K. Zhang, “Bayesian regression and bitcoin,” in Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on. IEEE, 2014, pp. 409–414.

[9] G. H. Chen, S. Nikolov, and D. Shah, “A latent source model for nonparametric time series classification,” in Advances in Neural Information Processing Systems, 2013, pp. 1088–1096.

[10] I. Georgioulas, D. Pournarakis, C. Bilanakis, D. N. Sotiropoulos, and G. M. Giaglis, “Using time-series and sentiment analysis to detect the determinants of bitcoin prices,” Available at SSRN 2607167, 2015.