

AN APPLIED STUDY OF RNN MODELS FOR PREDICTING CRYPTOCURRENCY PRICES

Dane Vanderbilt, Washburn University, dane.vanderbilt@washburn.edu

Kun Xie, Washburn University, kun.xie@washburn.edu

Wenyang Sun, Washburn University, nan.sun@washburn.edu

ABSTRACT

In recent years, cryptocurrency, and its foundational technology, blockchain, have become increasingly popular. As the value of some cryptocurrencies has cataclysmically risen and fallen, many people are curious about what the future of cryptocurrency prices look like, and what means are available to predict that future. We conducted experiments on three RNN models to determine whether one performed better than the others when predicting the price of Bitcoin, Ripple, and Litecoin. We also collected Google trends data for these three cryptocurrencies and ran additional experiments to explore whether this additional data significantly boosted the prediction accuracy of these models. After performing all our experiments, we found that among the three RNN models we tested, none performed significantly better than the others. Additionally, we concluded that supplementing the three models with Google trends data did not significantly boost the prediction accuracy of the models. We discuss implications of this research in the paper.

Keywords: Machine learning, cryptocurrency, cryptocurrency price prediction, Google trends

INTRODUCTION

Cryptocurrency is “a digital currency in which encryption techniques are used to regulate the generation of units of currency and verify the transfer of funds, operating independently of a central bank” (Cryptocurrency, 2019). The foundational technology behind most cryptocurrencies is blockchain, often “defined as a distributed, decentralized, public ledger” (Reiff, 2020). In recent years, these two intertwined technologies have become increasingly popular, as people have become less trusting of large third-party entities in security and finance. During the 2008 financial crisis, when people’s trust was at one of its lowest points, the first cryptocurrency, Bitcoin, was developed and in the following year released to the public. Although the release happened with little fanfare apart from a few early enthusiasts, over the next 10 years it would have a meteoric rise and pave the way for many more cryptocurrencies to emerge. The rise of Bitcoin peaked in December of 2017 when the price of one Bitcoin totaled almost \$20,000. However, the price of the cryptocurrency took a nosedive, dropping over 65% over the next three months, after a trading ban in China, numerous hacks and cryptocurrency exchange thefts (Bitcoin, 2020). Today, the price of Bitcoin and other cryptocurrencies has begun to recover, and many wonder what the future of cryptocurrency prices will look like.

Another field of technology that has experienced rapid expansion over the last few years is machine learning. As computers have become cheaper and more powerful, the ability to run complex models has grown, in addition to the length of time for running those models decreasing. There are many categories of machine learning, such as supervised, unsupervised, semi-supervised, reinforcement, self-learning, to name a few. Besides different categories of learning, there are numerous types of models, such as artificial neural networks, decision trees, support vector machines, regression analysis, and Bayesian networks (Machine Learning, 2020). The application of these machine learning techniques, as well as the types of models, are plentiful. However, certain types fit specific problems better

than others, and rarely is it the case that one type of learning or model works well with every problem. Although there are limitations, machine learning can often provide accurate predictions and keen insights into the problems they are tasked to help solve.

As the popularity of cryptocurrencies and machine learning have grown in tandem, many have attempted to use machine learning to explore whether there exists a viable type of learning or model that could predict the volatile future of certain cryptocurrencies. In this research, we attempted to answer two research questions. First, when comparing the Simple RNN, LSTM, and GRU machine learning models, does one model predict the cryptocurrency price better than others? Second, will adding Google trends data improve the prediction power? To address these questions, we conducted experiments on three RNN models, with, and without Google Trend data, on three cryptocurrencies: Bitcoin, Ripple, and Litecoin.

LITERATURE REVIEW AND RESEARCH HYPOTHESES

Several studies examined and compared machine learning models when predicting cryptocurrency prices. One paper used two regression and two recurrent neural network (RNN) models to predict Bitcoin prices and found that one of the RNN models performed visually better than the others in their experiments (Phaladisailoed & Numnonda, 2018). A research paper selected a LSTM model to predict Bitcoin and Litecoin prices (Yao, et al., 2018), and another used a modified model of a Binary Auto-Regressive Tree (BART), finding their model more accurate than traditional ARIMA-ARFIMA models when predicting the future price of the cryptocurrencies Bitcoin, Ethereum, and Ripple (Derbentsev, Datsenko, Stepanenko, & Bezkorovainyi, 2019).

Other studies have examined whether the addition of cryptocurrency-related data to machine learning models provides a significant boost to the model's predictive abilities. A study explored whether sentiment analysis of news articles and social media (in their study, Twitter) could help their models predict the prices of Bitcoin and Ethereum more accurately. In their study, the final version of their model was able to successfully predict large percentage increases and decreases over their 67-day test period (Lamon, Nielsen, & Redondo, 2017). Another study explored the improvements made to Bitcoin price predictions with the inclusion of sentiment analysis of Reddit data in comparison to Twitter data (Salač, 2019).

Our research aimed to take a new approach and compare different RNN (Recurrent Neural Network) models to each other to see whether any one model performed significantly better than the others. Additionally, during the exploration phase of our research, we were unable to find peer-reviewed and published articles that examined the effect of including Google trends data in a machine learning model. We aimed to fill this void by also exploring whether the addition of Google trends data relating to a cryptocurrency provided a statistically significant boost to a model's predictions. These two aims led to the following hypotheses of our research:

- H₁: Among three RNN models – Simple RNN, LSTM, and GRU - one will predict the price of a cryptocurrency more accurately than others.
- H₂: Adding Google trends data to the input of an RNN model will boost the prediction power of the model when compared to that model without Google trends data.

RESEARCH METHODOLOGY

To start our research, we collected daily trading data of three cryptocurrencies: Bitcoin (BTC), Litecoin (LTC), and Ripple (XRP), from coinmarketcap.com. Google trends data corresponding to each cryptocurrency was collected as well. With the help of the Keras Python library, we implemented three different modified versions of the RNN model: Simple RNN, LSTM, and GRU.

Collecting Google Trends Data

Since Google is the most popular search engine in the world and Google trends data is an indicator showing the popularity of certain keywords on Google, we wanted to explore whether the Google trends data of cryptocurrencies can be an input that improves an RNN model's prediction ability.

Google offers a website, trends.google.com, on which we have to manually request Google trends data for a certain period. Moreover, the extracted Google trends data may have different intervals depending on the requested period. For example, if you request the Google trends data for one or two months, you will receive daily data points for the requested period. However, if you request Google trends data for one or two years, you will only receive monthly data points for the requested period. Additionally, the extracted data has been normalized into a 0-100 range, so data from multiple months cannot be merged directly.

To extract daily Google trend data corresponding in a way that we can use for each cryptocurrency from 2015-01-01 to 2020-03-31, we used an open-source library, pytrends, which offers a set of easy-to-use APIs to extract Google trend data.

First, we extracted the monthly data points from 2015-01-01 to 2020-04-01. Then we extracted the daily data points for each month from 2015 January to 2020 March. Because both monthly data points and daily data points have been normalized into 0-100 range, they can't be merged directly. To remedy this problem, a small script was written to explore the relationship between two groups of data points. We assume there exists a linear relationship. Taking 2020 Jan as an example, denote $df_{daily}['2020-02-01'] = x_1$, $df_{daily}['2020-03-01'] = x_2$, $df_{monthly}['2020-02-01'] = y_1$, $df_{monthly}['2020-03-01'] = y_2$. An equation $ax + b = y$ is assumed to exist, satisfying both $ax_1 + b = y_1$ and $ax_2 + b = y_2$.

Then, we merged daily data points of each month into monthly data points from 2015-01-01 to 2020-04-01 to obtain the daily data points for the long-term period. Figure 1 shows the graph of the monthly data points of Bitcoin's Google trend data, and Figure 2 shows the result of the merged daily data points. This process was repeated during the collection and preparation of Google trend data for Litecoin and Ripple.




Figure 1. Bitcoin Monthly Data Points




Figure 2. Bitcoin Merged Daily Data Points

Setting up the RNN's

Recurrent neural networks (RNN's) are a class of neural networks where the output from the current step will be fed as part of the input to the next step. Thus, the most important feature of RNN is the hidden state, which is the "memory" capturing information about what has been calculated so far.

The formula for the current hidden state can be written as: $h_t = f(h_{t-1}, x_t)$, where h_t is the current hidden state, h_{t-1} is the previous state, and x_t is the current input. Figure 3 shows this process visually. Thus, with the help of the hidden state, the RNN's are very suitable for time series forecasting.




Figure 3. RNN Hidden State Structure (Thomas, 2019)

The Keras library supports multiple different modified versions of RNN models. In this paper, we chose three of them: vanilla RNN, LSTM, and GRU. The main difference between those three versions of RNN's is the structure of the neuron, which controls the way of computing the current hidden state by using the previous hidden state. Figures 7 shows the structures of the neurons in each model.




Figure 4. Neuron Structures of Three RNNs Algorithms (Karim, 2018)

Table 1 shows the list of parameters and corresponding values that were used for all three models. This resulted in a total of 64 different parameter combinations that were used on each of the three RNN models. Additionally, the models were tested with and without the Google Trends data, to determine whether Google Trends data made a statistically significant improvement to the predictions. We used this process for the prediction of Bitcoin, Litecoin, and Ripple to verify that our findings were not skewed simply by the cryptocurrency we had chosen to predict. Using the collected cryptocurrency price data, we train the models on all but 400 days of the data and used the resulting model to predict the price of the cryptocurrency over the next 400 days.

Table 1. Model Parameters and Values

Parameter	Values
epochs	10, 20, 50, 100
batch size	1, 2, 3, 4
neurons	1, 2, 3, 4
activation function	sigmoid
optimizer	adam
loss function	mean square error

RESULTS

For each currency, we tested six groups of observations from three algorithms with and without corresponding Google Trends data. Under different parameter combinations, we have 64 observations for each group.

To measure the prediction accuracy, the root mean square error (RMSE) of each model was calculated by comparing 400-day predicted values and the actual closing prices of each cryptocurrency. Since the RMSE is the square root of the residuals, lower values of RMSE indicate higher prediction accuracy. We got 64 RMSE's for each group. Table 2 displays the descriptive statistics of each group's RMSE. For each group of RMSE's, this table gives its mean, standard error of the mean, minimum value, maximum value, and median.

Table 2. Summary Statistics

Analysis Variable: RMSE						
Currency	Algorithms	Mean	Std Error	Min	Max	Median
BTC	RNN (Recurrent neural networks)	569.0208	84.032299	355.6334	4359.1896	383.6582
	RNN with Google Trends data	569.4191	76.773466	355.6339	3942.1596	376.1405
	LSTM (Long Short-Term Memory)	562.0800	85.839422	356.8653	4637.5279	389.4956
	LSTM with Google Trends data	553.1967	78.108222	355.4359	4254.2921	380.8386
	GRU (Gated Recurrent Unit)	403.2337	9.566334	358.1210	766.1170	377.0558
	GRU with Google Trends data	411.0974	14.988020	356.1217	1047.0533	370.0256
XRP	RNN	0.0242	0.002423	0.0126	0.0997	0.0162
	RNN with Google Trends data	0.0246	0.002264	0.0128	0.0944	0.0171
	LSTM	0.0214	0.001843	0.0126	0.0757	0.0144
	LSTM with Google Trends data	0.0231	0.001714	0.0126	0.0777	0.0173
	GRU	0.0216	0.001496	0.0126	0.0681	0.0164
	GRU with Google Trends data	0.0224	0.001536	0.0127	0.0653	0.0178
LTC	RNN	6.2346	0.689780	3.7382	35.1725	4.2503
	RNN with Google Trends data	5.0593	0.641111	3.5183	26.1380	3.8535
	LSTM	5.7564	0.616805	3.7183	33.2955	4.1191
	LSTM with Google Trends data	4.9071	0.661966	3.4853	21.5038	3.9364
	GRU	4.6481	0.196264	3.7417	11.0714	4.0760
	GRU with Google Trends data	4.4985	0.178672	3.4790	19.8900	3.7987

To explore whether those observations from each group are significantly different from each other, we performed independent samples t-tests comparing each pair of RMSE's using SPSS.

Bitcoin (BTC)

First, we compared each pair of the three RNN models. The independent t-test results for these comparisons are shown in Table 3. None of the resulting p-values were less than 0.05, although two comparisons were close. GRU performs better than RNN (p-value = 0.052). If we use 0.10 as a cut-off for p-value, then GRU performs better than LSTM as well (p-value = 0.068). Therefore, our first hypothesis, H1, is not strongly supported in our experiments with Bitcoin.

Table 3. T-test Results of Pairwise Comparison of RNN Models Predicting Bitcoin Prices.

Comparison	Means	Mean Diff.	t-Value	p-Value
RNN vs. LSTM	569.0208 vs. 562.0800	6.9408	0.058	0.954
RNN vs. GRU	569.0208 vs. 403.2337	165.7871	1.960	0.052
LSTM vs. GRU	562.0800 vs. 403.2337	158.8463	1.839	0.068

Second, we explored whether adding Bitcoin’s Google trends data as extra input could lead to more accurate results. Comparing the RMSE’s for each of the three RNN models with and without Google trends data, the independent t-test values are shown below in Table 4. None of the resulting p-values were smaller than 0.05. Thus, we failed to support our second hypothesis, H2.

Table 4. T-test Results Comparing RNN Models with and without Bitcoin Google Trends Data

Comparison	Means	Mean Diff.	t-Value	p-Value
RNN vs. RNN (with Google trends)	569.0208 vs. 569.4191	-0.3982	-0.003	0.997
LSTM vs. LSTM (with Google trends)	553.1967 vs. 562.0800	-8.8833	-0.077	0.939
GRU vs. GRU (with Google trends)	403.2337 vs. 411.0974	-7.8637	-0.442	0.659

Ripple (XRP)

Similarly, we performed the same procedures on our second cryptocurrency, Ripple. First, we compared each pair of the three RNN models. The t-test results of these comparisons are shown in Table 5. None of the p-values found were less than or equal to 0.05, nor were any of the values very close to 0.05. Therefore, we failed to support our first hypothesis, H1 in our experiments on Ripple.

Table 5. T-test Results of Pairwise Comparison of RNN Models Predicting Ripple Prices.

Comparison	Means	Mean Diff.	t-Value	p-Value
RNN vs. LSTM	0.0242 vs. 0.0214	0.0028	0.923	0.358
RNN vs. GRU	0.0242 vs. 0.0216	0.0026	0.905	0.367
LSTM vs. GRU	0.0214 vs. 0.0216	-0.0002	-0.098	0.922

Again, we explored whether adding Google trends data to the input of an RNN model would boost the prediction power of the model when compared to the models without Google trends data. Comparing the RMSE’s for each of the three RNN models with and without Google trends data, the independent t-test values are shown below in Table 6. Once again, none of the p-Values falls below 0.05, and we failed to support our second hypothesis, H2.

Table 6. T-test Results Comparing RNN Models with and without Ripple Google Trends Data

Comparison	Means	Mean Diff.	t-Value	p-Value
RNN vs. RNN (with Google trends)	0.0242 vs. 0.0246	-0.0004	-.0116	0.908
LSTM vs. LSTM (with Google trends)	0.0214 vs. 0.0231	-0.0016	-.0665	0.508
GRU vs. GRU (with Google trends)	0.0216 vs. 0.0224	-0.0008	-0.347	0.729

Litecoin (LTC)

Likewise, we performed the same procedures on our third cryptocurrency, Litecoin. First, we compared each pair of the three RNN models. The t-test results of these comparisons are shown in Table 7. In contrast to our previous results, one of the resulting p-values from the independent t-tests yielded a value less than 0.05. However, the p-values of the other two comparisons were greater than 0.05. If we use 0.10 as the cut-off, GRU performs better than LSTM (p-value = 0.089). We were unable to strongly support our first hypothesis, H₁.

Table 7. T-test Results of Pairwise Comparison of RNN Models Predicting Litecoin Prices.

Comparison	Means	Mean Diff.	t-Value	p-Value
RNN vs. LSTM	6.2346 vs. 5.7564	0.4782	0.517	0.606
RNN vs. GRU	6.2346 vs. 4.6481	1.5865	2.212	0.029
LSTM vs. GRU	5.7564 vs. 4.6481	1.1083	1.712	0.089

Once more, we explored whether adding Google trends data to the input of an RNN model would boost the prediction power of the model when compared to models without Google trends data. Similar to our previous experiments on Bitcoin and Ripple, our p-Values from our independent t-test comparisons resulted in no values less than 0.05. Therefore, we failed to support our second hypothesis, H2.

Table 8. T-test Results Comparing RNN Models with and without Litecoin Google Trends Data

Comparison	Means	Mean Diff.	t-Value	p-Value
RNN vs. RNN (with Google trends)	6.2346 vs. 5.0593	1.1753	0.381	0.704
LSTM vs. LSTM (with Google trends)	5.7564 vs. 4.9071	0.8493	0.155	0.877
GRU vs. GRU (with Google trends)	4.6481 vs. 4.4985	0.1496	0.466	0.642

DISCUSSION

Our research had the goal of answering two questions. First, does any specific RNN model perform statistically significantly better than other RNN models? The second question was whether Google trends data made any statistically significant improvement to the models' results. The answers we found based on our experiments largely did not support our original hypotheses.

Regarding the first question, although GRU perform statistically better than RNN on Bitcoin and Litecoin (p-value = 0.052, 0.029 respectively), unless we relax the cut-off for p-value to 0.10, we do not see the same trend for GRU when compared to LSTM, therefore, our research cannot conclude that GRU is the best among the three RNN models. When we first found this to be the case during the training and testing of the models on Bitcoin, we looked to test this on other cryptocurrencies, to explore whether this was true for other cryptocurrencies, or just Bitcoin specifically. After we tested two more cryptocurrencies with our models and found similar results, we concluded that we fail to support our first hypothesis. One possible reason our research found this result is because all of our models were RNN's, and although the neuron structure inside the models differed, the models overall are still very similar. Another possibility is that when predicting cryptocurrency, the RNN models we chose performed similarly, but when used on other problems, such as language modeling or classification prediction, some models may perform better than the others.

The answer to the second question was contrary to our original thinking as well. Going into our research, we assumed that introducing the Google trends data to our RNN models would make our models more accurate. However, the results of our experiments were mixed, and in our comparison tests, none of the models with Google trends data performed statistically significantly better when compared to the models without the Google trends data, as well as against each other. Thus, we also fail to support our second hypothesis. We think this was the case because although the interest in a cryptocurrency can initially spike concurrently with its price, the price often continues to increase, even as the interest decreases. As a result, the model's performance is initially boosted, but as the predictions progress, the model's performance is decreased, and the price of the cryptocurrency doesn't correlate as well to the Google trends data. So, the net result is that the Google trends data introduces noisy data into the model that causes a negative or negligible difference in the model's performance.

In addition to the predictions we made using the three RNN models, we also experimented with using a shifted version of the cryptocurrencies price history as the predictions. We used the price from the previous day, as well as the previous

week, as the predicted price of the cryptocurrency. While the predictions using the price from the previous week were poor, the predictions using the price from the previous day had RMSE's very similar or better than some of the predictions from the three RNN models we tested. This brings into question the suitability of RNN's when predicting the cryptocurrency prices and raises the question of what other machine learning models might be better suited to predicting the price of cryptocurrency.

CONCLUSION, LIMITATIONS, AND FUTURE RESEARCH

In this research, we first collected daily closing price data from 2015-01-01-to-2020-03-31 of different cryptocurrencies: Bitcoin (BTC), Ripple (XRP), and Litecoin (LTC) from coinmarketcap.com. Then, with the help of the Python pytrends library, we collected corresponding Google trends data over the same period. Using the Python Keras library, we implemented three modified versions of RNN model: Simple RNN, LSTM, and GRU. With 64 different parameter combinations, we trained those models with daily closing price including or excluding Google trends data for each cryptocurrency. For each parameter combination, we calculated the RMSE by comparing predicted values and actual values. To prove the validity of our hypothesis, we perform pairwise comparisons by using SPSS to do independent sample t-tests. The results showed that our two hypotheses are not supported.

Our research has three implications. First, when predicting cryptocurrency using RNN models, the difference in prediction accuracy between the models is negligible, as shown in our inability to support our first hypothesis. Second, RNN machine learning models may not be well suited for predicting the price of cryptocurrency, as we found in our comparison to simply using the price of the previous day as the prediction for the next day. Finally, the failure to support our second hypothesis showed that Google trends data is not a beneficial data input for RNN models when trying to predict the price of a cryptocurrency.

The main limitation of our study was with the data being used. Since Google does not release an official API to fetch the daily Google trends data, which have been normalized into 0-100 range, the daily data we did collect and combined had a chance of being incorrectly skewed, which would have also skewed our predictions. Another limitation arises from the values we used for the parameter combinations. Due to lack of computational power of our devices, we only chose small values for parameters such as epoch and neurons. We do not know whether large values of such parameters would have a significant influence on prediction accuracy. Furthermore, our work only focuses on three modified versions of RNN model and three different cryptocurrencies. There are more than 1000 cryptocurrencies listed on coinmarketcap.com and there are many other algorithms that can be used to predict prices of cryptocurrencies.

There are several future research opportunities based on our research. Besides Google trends data, news and social media comments from Twitter and Reddit should be investigated to see whether they can possibly boost the prediction accuracy of those algorithms. The scope of this research can be extended to include more cryptocurrencies, as well as using more models such as ARIMA models. It might be interesting to investigate GRU more since it appears to perform better than others in predicting Bitcoin and Litecoin prices. Finally, larger model parameters could be used, and would possibly result in a greater prediction power when predicting cryptocurrency.

REFERENCES

Bitcoin. (2020, May 6). Retrieved from Wikipedia: <https://en.wikipedia.org/wiki/Bitcoin>

- Cryptocurrency. 2019. In Lexico powered by Oxford. Retrieved from <https://www.lexico.com/en/definition/cryptocurrency>
- Derbentsev, V., Datsenko, N., Stepanenko, O., & Bezkorovainyi, V. (2019). Forecasting Cryptocurrency Prices Time Series Using Machine Learning. *Proceedings of the Selected Papers of the 8th International Conference on Monitoring, Modeling & Management of Emergent Economy*, 320-334.
- Karim, R. (2018). Retrieved from <https://towardsdatascience.com/animated-rnn-lstm-and-gru-ef124d06cf45>
- Lamon, C., Nielsen, E., & Redondo, E. (2017). Cryptocurrency Price Prediction Using News and Social Media Sentiment.
- Machine Learning*. (2020, May 6). Retrieved from Wikipedia: https://en.wikipedia.org/wiki/Machine_learning
- Phaladisailoed, T., & Numnonda, T. (2018). Machine Learning Models Comparison for Bitcoin Price Prediction. *ICITEE*, 506-511.
- Reiff, N. (2020, February 1). *Blockchain Explained*. Retrieved from Investopedia: <https://www.investopedia.com/terms/b/blockchain.asp>
- Salač, A. (2019). Forecasting of the cryptocurrency market through social media sentiment analysis.
- Thomas, R. (2019). Retrieved from <https://medium.com/@thomas.rhea9/awd-lstm-909eb75a5152>
- Yao, Y., Yi, J., Zhai, S., Lin, Y., Kim, T., Zhang, G., & Yoonjae Lee, L. (2018). Predictive Analysis of Cryptocurrency Price Using Deep Learning. *International Journal of Engineering & Technology*, 258-264.