Using RL to Predict Crypto and the effect of COVID -19

IPHS300, Spring 2021, Kenyon College, Prof Elkins and Prof Chun

Anav Dutt

Introduction

The stock market has always been subject to speculation. Oftentimes it is volatile but it is still an important mode of investment for many people. It carries great weight and is reflective of the performance of an economy. In macroeconomic courses in college, students are taught that in times of crises, CASH is KING, not assets in the market. Naturally, one would assume the COVID-19 pandemic would cause people to flock to holding monetary assets in cash. However, recently, as a result of the pandemic there has been a rise in the demand for cryptocurrencies. This might be because people expected the government to inject money into the economy in order to stimulate demand, and therefore people expected a rise in inflation and a fall in the value of money. So they took to a form of currency that doesn't fall under governmental control. Be it Bitcoin or Ethereum or even the rise in "meme coins" such as DOGE. There has been a shift in how economics predicts how people behave..

Project Overview

The stock market has always been a battle of potential risk and return in order to design a profitable strategy. A lot of times these risks are taken based on human speculation. With human speculation, there is always room for error, which is why large financial institutions often pay a hefty sum of money in order to gain a slight advantage over their competitors.
Reinforcement learning has the potential to change how speculation in stocks is undertaken. For this project, RL has been used in order to predict the crypto prices for 3 currencies, 1) Bitcoin, 2) Dogecoin and 3) Ethereum. This compares a pioneer in crypto with Bitcoin, a "meme coin" (DOGE) that has recently gained traction and Ethereum a coin that has been predicted to have high upside as well as important implications for the future. The goal of the project is to see whether RL models are able to understand the fluctuation in the prices of these cryptos and how Covid-19 has affected investment in cryptocurrency. We also look at, if with the use of Reinforcement Learning, humans are able to predict the changes that may arise in the crypto market and see how this functionality can be broadened.

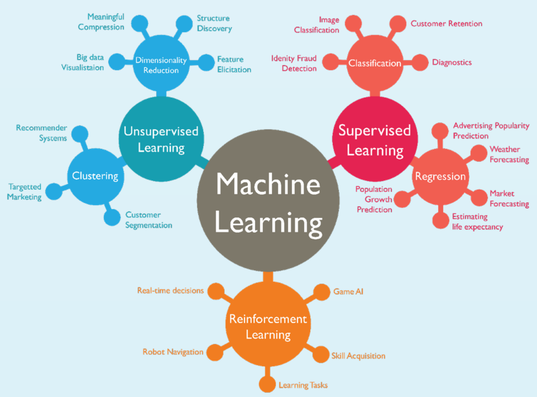

Technology and Background research

In the past, **fundamental analysis** (Graham and Dodd 1934), **technical analysis** (Murphy 1999), **and algorithmic trading** (Chan 2009) **as well as** attempts to use Supervised learning such as Regression and Classification models have been used in order to predict stock prices, however Reinforcement Learning (a subset of Machine Learning) looks to be the way of the future.
Reinforcement learning is a lot like the way humans learn The machine learns from the rewards it receives from the result of its actions. The feedback is inspired by humans. RL usually takes part in human domains such as decision making, incomplete information, etc.
Reinforcement Learning is based on a reward system that it receives from its actions. When predicting stock prices, the benefit of Reinforcement Learning is that it can take the data that it has been trained on to forecast the next shifts. The task is to maximize the reward by taking actions with respect to the information available. RL doesn't need large labeled training datasets, which helps cut down time to clean and sort data. Uses a reward function in order to optimize the reward in the future, unlike a ML classification model that predicts the probability of outcomes. RL also solves the optimization problem by maximizing the future rewards.
The notebook used for this project was only trained once. Multiple iterations and trainings would be ideal for the machine to train itself and perform optimally. Also in order to improve the performance and predictions of the trading bot, custom indicators should be added to the notebook. The metrics used to analyze the effectiveness of the RL bot, it is the monetary payoff it generates.
We can see that the use of RL is and will not be limited to a few niche firms. Institutions such as JP Morgan, "**announced it started utilizing Deep Neural Network for Algo Execution (DNA) for boosting its FX trading algorithms…DNA uses RL to assess the performance of individual order placement choices". It is clear that RL will be used in mainstream trading sooner rather than later as a result of its effectiveness to learn on it's own and be responsive to human feedback to generate better models. It seems as though, RL is being used in predictive models as of now, but there does seems to be a market for trading bots trained on RL as well.**

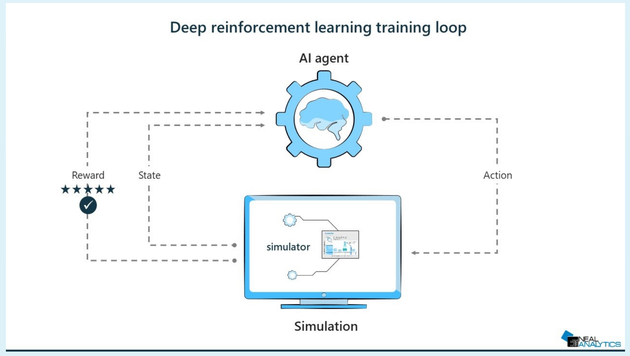

Results and Analysis

Bitcoin prediction

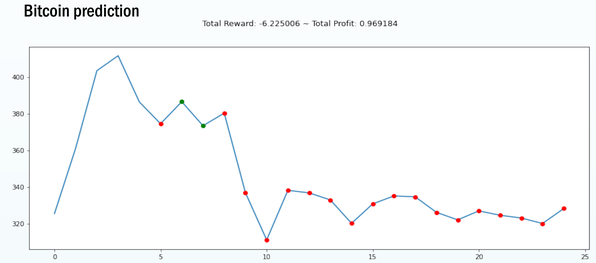

Dogecoin prediction

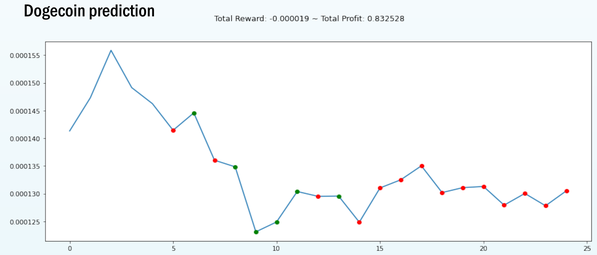

Ethereum prediction

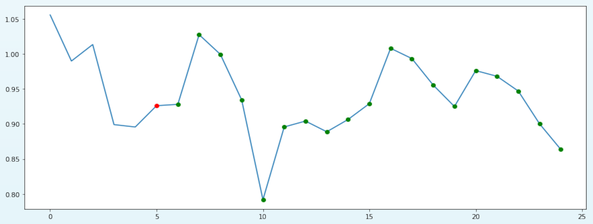

Taking a look at the data above, Bitcoin and Doge appear to have a stronger correlation compared to the prediction for Ethereum. This close correlation between Doge and Bitcoin could be explained by a rise in their popularity over the pandemic. They were often discussed within the same context, as Doge gained traction as a "meme", investors flocked to it, with experts terming it as a 'gamblers rush', assuming it to be a cheaper form of crypto compared to Bitcoin. The rise in investors and so-called financial advisors on "tik-tok" who fueled the fire by portraying false information and pumping up Doge to be similar to Bitcoin, which could explain their close correlation. It also appeared to have the backing of Elon Musk. However the correlation doesn't cover the prices at which these cryptos trade at. Bitcoin is much more expensive than Doge which is another interesting implication that is visible from this model, where it appears that the Doge graph appears to have the same min/max bounces and their peaks and valleys are correlated – potentially as a result of investors who don't have the knowledge looking at Doge as a substitute for Bitcoin.
All 3 graphs appear to fall extremely low at about the same point which can be explained by the post crypto surge of 2018, an initial boom leading to a fall in the prices after investors ducked out.
The trading bot we created is interesting as often, when it should short the crypto, it buys it. The buys are represented with green dots and shorts with red. However, since this bot was only run once, it doesn't have all the data required to train, which is why we see a profit that is less than what we started with. We can read the profits, for example, the Ethereum profit is = 0.917, so if you start out with a dollar, the bot returns 0.917$. This could also be a result of the speculative nature of crypto, where there doesn't seem to be a clear trend, and could be hard to predict as it is based on human expectation.

Ethical Implications

Throughout history, there are a few who benefit off the back of others. Let's look at who has profited the most from the crypto boom from the start of the pandemic. It is mainly the bigger firms and corporations along with those with high income, whereas the people investing with lower wages suffer as they are not able to pay the high prices, or are not privy to market information.
Since a lot of big investors are buying in, this makes people believe there is real value to crypto. Pension funds, retirement funds, college students and other vulnerable investors might get lured into what may be a trap, a scheme to bump up the price and get out, leaving the crash on the hands of the vulnerable. It is quite easy to see that reinforcement learning can be used to exploit those who may not have the resources or the in depth knowledge of the stock market or crypto. This exploitation will primarily target those with more to lose. There may be a need for regulation when it comes to RL and predicting stock prices in the near future in order to avoid the exploitation of those vulnerable to its power.
Covid 19 has resulted in a massive transfer in wealth. The wealthy profited on the hard situations faced by common people during the pandemic. This is clear in the crypto market as well. Institutional investors made investments into reputable cryptos such as Bitcoin, while common folk probably don't have the capital to fund the purchase of Bitcoin and invest in coins such as Doge. As mentioned in the technology section, what happens to normal investors when institutions such as Goldman use RL to its full potential. What is often left out of consideration is also the jobs that may be lost with improvement in RL. Jobs of analysts may be rendered irrelevant.

Conclusion and Considerations

Looking into the future cryptocurrency is beginning to have strong institutional acceptance from firms such as Tesla and other big figures such as Venmo, Uber and Goldman Sachs to name a few, however, it has the potential to be extremely volatile as a result of speculation and it's true value since cryptocurrency trades like a commodity, but behaves like a currency.
At this moment in time there is not much regulation when it comes to cryptocurrencies or using RL in pricing models, so it will be interesting to see how this rise in it's demand will raise its regulation by governmental bodies.
The power for RL extends far beyond crypto. It can be applied to other financial and economic indicators, such as stock trading, helping the government maintain a steady rate of inflation, avoid liquidity traps, help provide insight to monetary and fiscal policy analysis and more.
The market is always subject to turbulent fluctuations as can be seen in the crisis in 2008 and the pandemic as well. It will be interesting to see how RL responds to such changes, it might not be able to predict it, but the rate at which it changes its prediction will be interesting.
While RL does a good job, there are a lot of other factors that need to be considered when predicting the value of crypto and human feedback is essential. However, the plain fact is that RL models are effective at predicting pricing in crypto and it is clear from the models here that covid has caused certain correlations and fluctuations in crypto.

Citations

- Jon Chun, Kenyon College, chunj@kenyon.edu - https://www.kenyon.edu/directory/jon-chun/
- Articles – https://jfds.pm-research.com/content/2/2/25
- https://www.sciencedirect.com/science/article/pii/S0020025520304692
- https://towardsdatascience.com/deep-reinforcement-learning-for-automated-stock-trading-f1dad0126a02
- https://news.ycombinator.com/item?id=24344295
- https://news.ycombinator.com/item?id=24344295
- https://www.sia-partners.com/en/news-and-publications/from-our-experts/future-cryptocurrency-2021-and-beyond-industry-milestones
- Github links - https://github.com/AI4Finance-LLC/Deep-Reinforcement-Learning-for-Automated-Stock-Trading-Ensemble-Strategy-ICAIF-2020
- Kaggle links- https://www.kaggle.com/sudalairajkumar/cryptocurrencypricehistory