

Team 1: Initial Project Description

Mason Wilde, Jacob Hegna, Gage Benne, Andy Monroe, Grant Jurgensen

Project Name

Lawrence Trading Co.

Project Synopsis

An algorithmic cryptocurrency trading platform utilizing Tensorflow machine learning and deep market data to produce successful short-term trading strategies.

Project Description

The field of algorithmic trading is already strongly established within the realm of traditional stock exchanges. Typical strategies rely on extremely quick algorithms and low latency connections to an exchange in order to make short-term trades. Often, algorithms are implemented on a field-programmable gate array (FPGA) due to their speed, and the need to trade as quickly as your competitors to maintain a viable trading strategy. These factors combine to create a high barrier of entry into algorithmic trading on traditional exchanges.

In contrast, algorithmic trading of cryptocurrency is relatively undeveloped, and therefore presents a strong opportunity to enter while competition is thin. Cryptocurrency exchanges are hosted on the internet, which inherently creates a more even playing field than traditional exchanges. Where trades can be made in a matter of nanoseconds on traditional exchanges with direct connections, internet latency will be much higher for all parties. This also means we will need to make strong changes to established techniques in algorithmic trading, as we cannot assume low latency access to the exchange, and we cannot assume quicker access than our competitors.

The end result will be a program which trades cryptocurrencies through an internet exchange based on a machine learning algorithm, trained on market history data we have collected.

Project Milestones

First Semester

1. Collection of raw cryptocurrency market trading data
2. Manual modelling and analyzing / determining market indicators
3. Tensorflow, and general machine learning integration beginnings

Second Semester

1. Production and validation of generated algorithms from collected data
2. Cryptocurrency market integration
3. Production and validation of generated algorithms from real-time cryptocurrency market
4. Putting money on the line and verifying results

Project Budget

This project will primarily utilize Google's open source machine learning framework: Tensorflow. Open source programming languages coupled alongside the machine learning framework such as Python and C++ mean very little implementation costs. That being said, there may be costs associated with hosting certain components of the project on Amazon Web Services, as well possible costs involving automated trading at high frequencies on the open cryptocurrency markets. (There is also a stage in which we put our own money in the algorithmic trading platform, we will not include this in the budget.)

Work Plan

As of this point, clear roles have yet to be determined due to the nature of the project's exploratory intent. This project was undertaken in attempts to become more familiar with interesting growing fields in computer science not taught in our undergraduate careers. That said, strengths of each team member will be a guide to who will take on what component of the project. Mason has had experience with machine learning algorithms, and Jacob has had experience with exchange markets. The entire team has had extensive experience in many facets of computer science, and the team will build on these strengths to determine the most optimal work allocation.

Gantt Chart

Task	Start	Finish	Assigned	Duration
Project proposal	early-October	mid-October	all	< 1 week
Video project proposal preparation	mid-October	mid-October	all	< 1 week
Collection of raw cryptocurrency market trading data	mid-September	mid-February	Jacob / all	ongoing
Cryptocurrency exchange research	mid-October	early-November	all	< 2 weeks
Manual modelling	mid-October	mid-November	Jacob, Andy / all	< 4 weeks
Analyzing manual modeling	mid-October	mid-November	all	< 2 weeks
Researching market indicators	early-November	mid-November	Grant, Mason / all	< 2 weeks
Tensorflow research	mid-November	late-December	all	~ 1 month
Tensorflow and general machine learning integration	mid-November	late-February	all	~ 2 months
Data formatting to support machine learning attempts	mid-December	mid-January	Grant, Jacob / all	< 2 weeks
Assisted algorithm development from collected data	mid-January	mid-February	all	~ 1 month
Validation and correction of generated algorithms	mid-January	mid-February	all	~ 1 month
Cryptocurrency market integration	early-February	late-March	Gage, Andy / all	< 2 months
Production and validation of generated algorithms from real-time cryptocurrency market	early-March	late-March	Andy / all	< 2 weeks
Putting money on the line and verifying results	late-March	early-April	all	< 1 week
Project documentation	mid-October	mid-April	Gage, Mason / all	ongoing
Public presentation preparation	mid-Mar	mid-April	all	< 1 week

	Early Sep.	Mid Sep.	Late Sep.	Early Oct.	Mid Oct.	Late Oct.	Early Nov.	Mid Nov.	Late Nov.	Early Dec.	Mid Dec.	Late Dec.	Early Jan.	Mid Jan.	Late Jan.	Early Feb.	Mid Feb.	Late Feb.	Early Mar.	Mid Mar.	Late Mar.	Early Apr.	Mid Apr.
Project proposal				█	█																		
Video project proposal preparation				█	█																		
Collection of raw cryptocurrency market trading data		█	█	█	█	█	█	█	█	█	█	█	█	█	█	█	█	█	█				
Cryptocurrency exchange research				█	█	█	█	█															
Manual modelling				█	█	█	█	█															
Analyzing manual modeling				█	█	█	█	█															
Researching market indicators							█	█	█														
Tensorflow research								█	█	█	█	█	█										
Tensorflow and general machine learning integration									█	█	█	█	█	█	█	█	█	█	█				
Data formatting to support machine learning attempts										█	█	█	█	█	█								
Assisted algorithm development from collected data														█	█	█	█	█					
Validation and correction of generated algorithms														█	█	█	█	█					
Cryptocurrency market integration																█	█	█	█	█	█		
Production and validation of generated algorithms from real-time cryptocurrency market																			█	█	█	█	
Putting money on the line and verifying results																						█	█
Project documentation				█	█	█	█	█	█	█	█	█	█	█	█	█	█	█	█	█	█	█	█
Public presentation preparation																					█	█	█

Preliminary Project Design

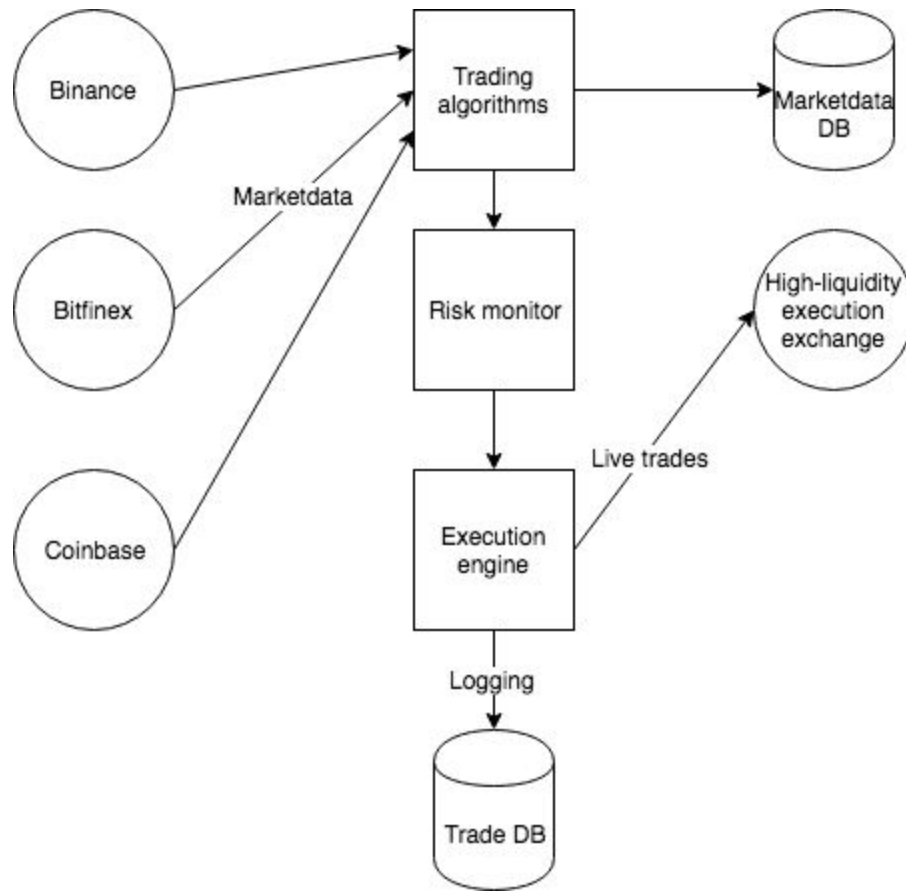
We have set some technical limitations on ourselves in order to keep the project focused and on topic. We have decided to focus strongly on C++ and Python for our languages. C++ was chosen due to its power, efficiency, and popularity, both within the group and industry-wide. This is the language we will use for our computationally significant and time-sensitive components. Python will be used in less computational sections of the project on account of its ease of use, as well as its incredible external libraries such as TensorFlow. Note that TensorFlow and many other Python libraries are written in lower level languages and merely expose a Python calling interface. Likewise, we will use PyBind to establish our interoperability between C++ and Python.

In particular, we will aim to have much of the trading infrastructure be implemented in C++, and do market/strategy research in Python. For example, we will implement:

- Market data feeds/API access
- Trade execution
- Risk (keeping track of the amount of money exposed to various cryptos, with the ability to override trades from the execution engine if they over-expose us. For example, if two different trading strategies each decide to be long on Ripple, the risk platform might stop the second trade because too much money would be dependent on that asset)

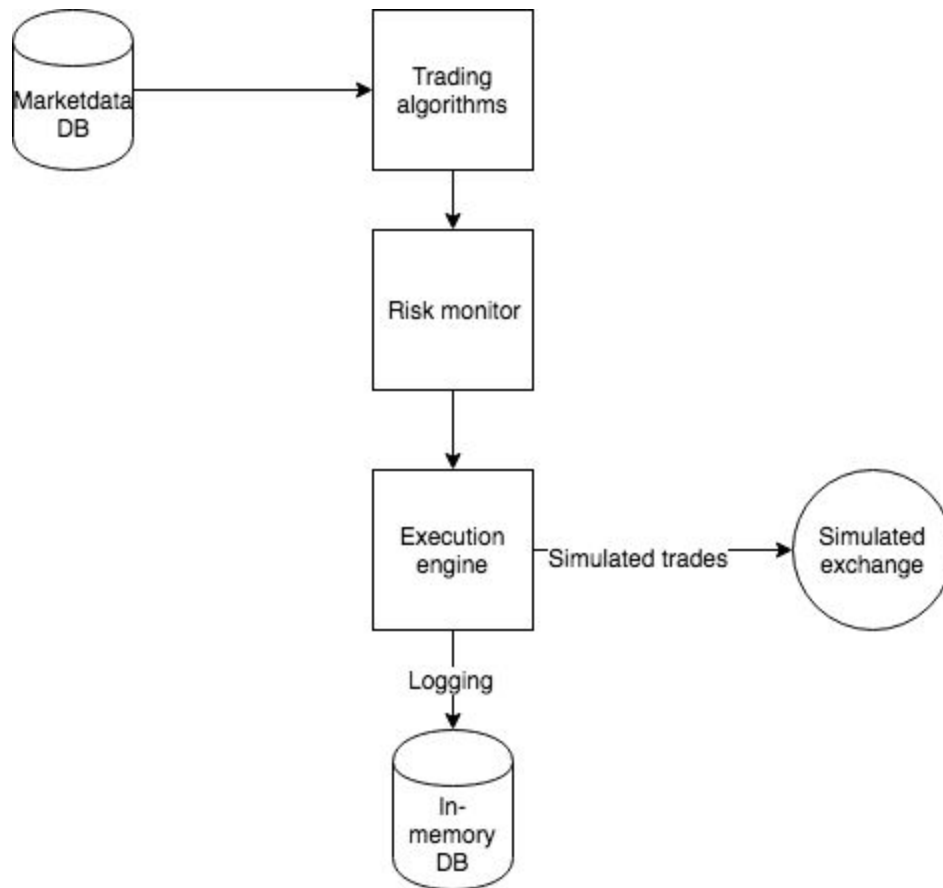
In C++. We will aim to implement things like a simulated trading environment, historical backtesting, and database interaction in Python, as it is easier to get up-and-running in Python and those features depend significantly less on low-latency code.

In live-trading, our infrastructure will resemble:




In reality, we will likely have different/more exchanges we listen to for marketdata. We will use (at least) two databases, one which records the trades that our system made, and one that records every piece of market data using a uniform format, so we can simulate trading days later. Through live testing, we will determine which exchange has the deepest liquidity for our trades (this is important, as in a low-liquidity exchange, incoming price data might not match outgoing execution prices).

Our simulation environment will look nearly identical, with minor changes:



The trading algorithms, risk monitor, and execution engine will be identical to our live environment (the only difference is that will instantiate the execution engine and trading strategies with the testing objects, instead of their live counterparts. Dependency injection allows us to abstract away information about the actual exchanges, and instead only worry about the incoming/outgoing data). We will just be replaying historical market data from the database. We will use a simulated exchange that records things like profit/loss, slippage, Sharpe ratio, and volatility of our profits. The historical testing will be not only for training our ML models, but as to do basic regression analysis. A core trading strategy will also be exploiting correlations in different crypto prices. For example, an early iteration of our code produced this graph (using matplotlib) which shows a price trend within Ethereum and Ripple (adjusted for the scale of the prices):



Which shows the prices are clearly correlated. If we can find a lagging indicator for one of the prices in terms of the other, we can predict price changes and trade accordingly.

In addition, we will attempt to support the Linux and macOS operating systems. This restriction is not self-imposed, but rather a reflection of our development environments and final production environment. The team is working on a combination of Ubuntu and macOS, and we hope to run this software on Amazon Web Services, which supports a variety of Linux distributions. Fortunately, these operating systems are relatively simple to co-develop for on account of their shared heritage and near POSIX compliance.

Finally, we are constrained by the libraries available to us. Thankfully, there is a myriad of powerful open source libraries such as TensorFlow that are available to be used freely. Furthermore, there are libraries such as matplotlib and pybind which allow us to streamline data analytics and interaction between C++ and Python codebases.

On the business side, we are constrained by our schedule. The project has predefined deadlines, and we are also limited by the availability of the team members and external time commitments. Together, these set the scope of our project, and disallow us from expanding it any further. Many trading teams take years to set up their core infrastructure and testing/learning environments - clearly, we will not be able to have this much flexibility, so we will have to skip some features. For example, many trading teams have a system which automatically takes the recorded market data from that day of trading, updates their trading models overnight, and then executes the new trading models the next day with no programmer interaction. This frees the developers to work on new strategies. In contrast, we will likely have

to manually re-train our models on new data, and re-deploy the updated strategies to the trading server.

Our budget will determine where we are able to run our code, the third party software which we can license, and our ability to trade cryptocurrencies. We hope to acquire the funding to deploy and maintain our project on Amazon Web Services. We will likely be able to rely totally on open source and similarly licensed free software. However, we will need to interface with cryptocurrency exchanges, which typically have fees associated with making a trade. Finally, we will require initial funds in order to purchase cryptocurrency for our program to participate in real trades. In order to save money, we will attempt to tune and train our algorithms throughout development by simulating trades. Using real data, we would record trades the algorithm would make in real-time, but without actually trading. We would then go back over our trades and evaluate how successful the algorithm would have been. This is a significant constraint, because we cannot perfectly simulate how a real trade would play out. We will need to make assumptions on latency, and how our trades would actively affect the exchange, if at all.

Ethical and Intellectual Property Issues

Ethically, the primary issue of the project relates to its contributions to society as a whole. At its core, the goal of the project is to maximize profitability for the people using the algorithm to trade cryptocurrencies. In this regard, the project will only directly benefit its creators and will have little to no positive effect on the rest of society. Concerns regarding negative effects to society are also present. Simply put, the cryptocurrency market is a competitive environment, and one party cannot move up without another moving down. Some refer to algorithmic trading as a tax on retail investors, as the bots can get the best deals faster than any human ever could. As the project is contained to cryptocurrencies rather than traditional stocks, this issue is less concerning, but it still warrants mention. Ultimately, the potential negative impacts of the project are minimal, while the benefits are potentially very high, if even for only a small group of people, so we believe it is ethical.

Intellectual property will be slightly complicated for the project. There are two primary areas that will require care to avoid any infringement. First, the trading algorithm itself will need to be fairly unique. There are numerous sources of trading algorithms for the stock market available online in various developer blogs, research articles, and other media. Most of this material is not obviously licensed in any way, but it is also generally only a high level description of the algorithm or technique. As long as no actual code is reused, use of the ideas available online will be safe. Second, the machine learning algorithms used to train and predict information about trades will be a mix of available tools and original code. Luckily, ML tools such as Scikit and Tensorflow are very generous with their licensing, however, it must be verified that any tools we use allow financial gain from their use.

Updated October 22, 2018: Finalized project proposal by adding project budget, work plan, gantt chart, preliminary project design, as well as ethical and intellectual property issues.