Worcester Polytechnic Institute
Digital WPI

Interactive Qualifying Projects (All Years)

Interactive Qualifying Projects

April 2019

Forex Trading Systems: An Algorithmic Approach to Technical Trading

Andrew David Vanosten
Worcester Polytechnic Institute

Benjamin John Anderson
Worcester Polytechnic Institute

Nicholas Sylvester Nugent
Worcester Polytechnic Institute

Thar Min Htet
Worcester Polytechnic Institute

Follow this and additional works at: https://digitalcommons.wpi.edu/iqp-all

Forex Trading Systems:

An Algorithmic Approach to Technical Trading

An Interactive Qualifying Project

Submitted to the Faculty of

WORCESTER POLYTECHNIC INSTITUTE

In partial fulfillment of the requirements for the

Degree of Bachelor of Science

By

Benjamin Anderson (REL 1804)

Thar Min Htet (REL 1802)

Nicholas Nugent (REL 1810)

Andrew VanOsten (REL 1808)

Date: April 24, 2019

Submitted to:

Professors Reinhold Ludwig and Gbetonmasse Somasse of Worcester Polytechnic Institute

Abstract

In financial trading, emotion can often obstruct clear decision making. The goal of this project is to build a system which can overcome this by trading foreign currencies autonomously. Three systems were created: two relying on neural networks, and a third on pattern recognition of candlestick charts. A fourth system has been designed to allocate funds to the others using utility theory. Though the algorithms were not profitable, a powerful interface was built, connecting Python scripts to MetaTrader 4 for trading.

Acknowledgements

We would like to thank Professors Reinhold Ludwig, Gbetonmasse Somasse, and Hossein Hakim for their insight and guidance throughout this project.

We would also like to thank Worcester Polytechnic Institute for allowing us the opportunity of working on the project.

Executive Summary

Introduction

Many people spend decades of their lives working in order to make and save money. It can be quite easy to equate money with comfort or power. Consequently, it is not hard to be caught up in the idea of maximizing one's assets. Compensation for employment can be quite rigid, and often lower than one may desire, leaving employees searching for other sources of income. One way that this could be accomplished is through long-term investment in equities and federal bonds. When carefully thought out, this approach can provide a relatively conservative extra income over a long period of time. For those who desire a greater reward, albeit at an increased level of risk, there exists an alternative: short-term trading of equities, currencies, commodities, or other assets.

Decades ago, short-term trading was not accessible to the average person. However, the advent of the internet, and the associated "Information Age" has brought brokerage platforms into the hands of the general public. With a potentially small initial investment, this type of trading can be an attractive means of generating funds. Without careful planning and analysis, it is very easy to lose money in any active market. To consistently make a profit over time, we hypothesize that one must have a very disciplined and scientific system of trading methods.

While no system can work perfectly in any situation, a common source of error is the person doing the trading. Humans are easily swayed by emotion and, as a result, tend to not be very good at following the recommendations of a system. People frequently believe they are smarter than the system, and feel that they can make better decisions, in the pursuit of a greater profit. Some algorithmic trading systems aim to solve this by executing a set trading strategy, while others aim to make a profit by running at extreme speed with the goal of exploiting tiny inconsistencies in a market [1]. This project focuses on the development of carefully planned algorithmic trading systems, which aim to incur a profit by carrying out predefined actions based on technical analysis. When driven by careful analysis, we see this algorithmic approach as an optimal solution for making money in a stock or currency market.

Goals

The goal of this Interactive Qualifying Project is to eliminate the psychological and emotional aspects of manual trading by creating an autonomous trading system which is capable of making short-term trades of foreign currencies. The system will be driven by a combination of scientific and thoughtful technical analysis, and by neural networks. We do not believe that it

is possible that a single, realistic algorithm, which would trade successfully in any situation, can exist. The goal of this project is therefore to create a set of algorithms that can each function under certain conditions. Together, these systems create a more holistic solution to the problem than any individual algorithm could.

Results and Conclusions

In order to accomplish the goals set forth above, three systems were implemented. Two of these systems planned to experiment with using neural networks. Because of this, we decided to write our trading systems in Python as it has a large set of neural network and machine learning libraries including TensorFlow, Scikit, and Numpy. We also decided that we wanted to use MetaTrader for testing these systems because of its robust trading and backtesting framework. In order to use Python and MetaTrader together, we created a pair of programs, one in MetaTrader's MQL4 language and one in Python, which pass messages to facilitate trading. MQL4 passes tick data to Python which, if it chooses, can return a command to place a buy or sell order.

The following sections summarize each of the three trading systems, as well as the asset allocation system.

An Approach with Recurrent Neural Networks using Long Short-Term Memory

Developed by Thar Min Htet

This trading system is implemented using Recurrent Neural Networks (RNNs) with Long Short-Term Memory (LSTM). This design was chosen over other types of neural networks because of its ability of RNNs to handle time series data, such as financial data, very well. LSTM is an add-on to improve the performance of RNNs. This model is not designed to trade currencies live; it is designed to be incorporated into other trading strategies and give support to different systems. With further development, such as implementation of the connection to MetaTrader, this system could be used for live trading of currencies.

The workflow for the system is summarized as follows.
1. Remove noise from raw currency data using a wavelet transform
3. Train the network on a subset of the processed dataset
4. Tune Hyperparameters
5. Test the accuracy with both training and testing data
6. Make a prediction of the currency price at the next timestep

Automated Momentum Indication via Candlestick Patterns

Developed by Andrew VanOsten

Looking for patterns in candlestick charts is a common practice among traders. Just like any other indicator, these patterns cannot always give a perfectly clear signal to the trader. However, they can still provide a valuable indication of trend reversals. These patterns are very simple to search for with an algorithm as they are defined by the open, high, low, and close prices for each bar—data that are easily available. By using somewhat agreed-upon definitions of common patterns, an algorithm can quickly search an incoming stream of data for potential signals.

These signals each represent the potential for a change in the direction of a trend. However, most traders would not base a trade on this signal alone; similarly, this algorithm takes a slightly deeper approach. A market which is in a downtrend is likely to see several bullish signals as it nears its low point, before reversing and moving upwards. This algorithm generates a small positive signal for each bullish signal, and a slightly negative signal for each bearish signal. Different patterns can be given individual weights to allow for differing reliability of each pattern. These small signal values are simply added to an overall momentum "score" which is used as a signal to buy or sell. As time moves forward this score will decrease exponentially, as the signals seen previously will be less accurate as they grow further into the past. The final value of the score at the nth time-step is defined as:

$$M(t_n) = D \cdot M(t_{n-1}) + \varepsilon$$

where $M(t)$ is the momentum score being generated, t_n is the current time-step, t_{n-1} is the previous time-step, D is the decay constant where $0 < D < 1$, and ε is the signal being generated—a constant proportional to the reliability of the pattern that has been observed. This function was created in order to model the desired behavior of the system. When $M(t) > 1$, one lot of the selected currency is bought, and when $M(t) < 1$, a lot is sold. When $|M(t)| < 0.5$, any currently held orders are closed, as this represents a loss of confidence in the indicator. Only one position of one lot may be held at any time.

Unfortunately, this system was not able to produce a profit over time. However, there are countless combinations of currency pairs, timeframes, and system tunings that could potentially produce better results. At the very least, this system could act as a basis or source of ideas for future research.

Q-Learning using Neural Networks

Development by Ben Anderson

This system attempted to use Q-learning along with neural networks to assign values to potential states which the market could have. The first attempt used pure Q-learning. A

feed-forward neural network was fed information about the past 20 bars in the market, as well as the current trade, if any, which was in place. Before being fed into the system, data was preprocessed—reformatted in order to represent each bar relative to the previous bar, or to the open price of the current bar. Representing the data in this way allowed the system to consider each state independently without having to account for differences in the range of input data. Once all data was fed in, the network was trained using the following formula:

$$loss = |Q(s_t) - (r_t + \gamma \cdot Q(s_{t+1}))|$$

where Q(s) is used to represent the predicted reward of a certain state, r is the received reward, and ɣ is the decay factor which was set to 0.9. The neural network was trained to minimize this loss function for 5 million iterations.

A slightly different approach of using Q-learning was used to make a second neural network. This design came from the realization that the neural network had to do a multiplication of the current amount traded and the predicted change in the market. This was necessary in order to go from making a prediction about the future behavior of the market to, predicting a reward. This is possible, but unnecessarily complicated when it is known that the multiplication is necessary. In light of this fact, a second neural network was designed under the assumption that one lot would be bought. The value of a state could be calculated by multiplying the actual trade size by the value of the state where one lot was owned. The same formula was used to calculate loss for this network, but r was replaced with the change of the next bar. This network was trained for 10 million iterations.

Unfortunately, this system was not able to make profitable trading decisions. The change in the market during an upcoming bar and the predicted value of owning 1 lot are not related ($R^2 = 0.003$). Figure 6.12 helps explain why this problem occurs. If the market is going to go up, buying should be desirable and selling should not be. If the market is going to go down the opposite should be true. However, the system finds that the more desirable owning 1 lot is, the more desirable owning -1 lots is ($R^2 = 0.6357$). This is unfortunately a contradiction which prevents this system from ever being able to make accurate predictions.

Using Expected Utility Theory for Asset Allocation

Developed by Nicholas Nugent

Part of the goal for this project was to devise a way to combine the algorithms into an overarching system. To accomplish this we first examined using the Von Neumann–Morgenstern utility theorem(VNM) to calculate the expected utility of trades based on their take profit(A), stop loss(B), and probability of success(p). The equation describing the lottery under VNM is:

$$L = pA + (1 - p)B$$

We can now rank them by expected utility and allocate our money according with the highest expected utility getting the largest allocation. We could also reject trades that have a negative expected value. This system performed very well in a proof of concept test when compared to a naive allocation, however proved to be impractical to implement for the system. The requirements of simultaneous trades and defined outcomes, in the forms of take profits and stop losses, meant it could not function with the other algorithms developed in this report. From there we decided to look at as a portfolio allocation problem were the algorithms are the assets we can invest in and found the below utility function that could be used to evaluate the algorithm based on there historical performance.

$$U = E(r) - 0.5 * A * \sigma^2$$

Here *E(r)* is the expected return of the asset. *A* is a risk aversion coefficient from 1 to 5 with 5 being most risk averse, and σ is the volatility of the asset.

Table of Contents

List of Figures

List of Tables

1 Introduction

Many people spend decades of their lives working in order to make and save money. It can be quite easy to equate money with comfort or power. Consequently, it is not difficult to be caught up in the idea of maximizing one's assets. Compensation for employment can be quite rigid, and often lower than one may desire, leaving employees searching for other sources of income. One way that this could be accomplished is through long-term investment in equities and federal bonds. When carefully thought out, this approach can provide a relatively conservative income over a long period of time. For those who desire a greater reward, albeit at an increased level of risk, there exists an alternative: short-term trading of equities, currencies, commodities, or other assets.

Decades ago, short-term trading was not accessible to the average person. However, the advent of the internet, and the associated "Information Age" has brought brokerage platforms into the hands of the general public. With a potentially small initial investment, this type of trading can be an attractive means of generating funds. Without careful planning and analysis, it is very easy to lose money in any active market. To consistently make a profit over time, we hypothesize that one must have a very disciplined and scientific system of trading methods.

While no system can work perfectly in any situation, a common source of error is the person doing the trading. Humans are easily swayed by emotion and, as a result, tend to not be very good at following the recommendations of a system. People frequently believe they are smarter than the system, and feel that they can make better decisions, in the pursuit of a greater profit. Some algorithmic trading systems aim to solve this by executing a set trading strategy, while others aim to make a profit by running at extreme speed with the goal of exploiting tiny inconsistencies in a market [1]. This report focuses on the development of carefully planned algorithmic trading systems, which aim to produce a profit by carrying out predefined actions based on technical analysis. When driven by careful analysis, we see this algorithmic approach as an optimal solution for making money in a stock or currency market. To that end, this document provides relevant background information in Chapter 2, a more fully defined statement of the project's goals in Chapter 4, details the development of several algorithms in Chapter 5, and describes the results that were produced in Chapter 6.

2 General Background

2.1 Trading versus Investing

There are two different genres of wealth creation in equity markets: trading and investing. A major difference between trading and investing is the duration for which the asset is held. In trading, one puts money in a certain asset for a relatively short amount of time, whereas, in investing, one operates by buying and holding an asset for a long period. This period could be months, years or even decades. Short-term traders aim to profit through fluctuations of a market on the scale of minutes, hours or days. Such quick fluctuations are insignificant to investors. Another large difference between trading and investing lies in the amount of risk taken as compared to the amount of return that could potentially be generated. In general, trading involves higher risk and the potential for higher return relative to investing. The decision of which strategy is "better" is an entirely personal one, and therefore will not be discussed further in this report.

2.2 Economic Effects on Trading

Macroeconomics, the branch of economics that studies the behavior of the aggregate economy, can play a large role in forex trading. By studying inflation rates, national income, gross domestic product (GDP), unemployment rates, or other macroscopic indicators, one may be able to gain insight into the long-term trend of a currency. On these broader time scales, a forex market can be driven by macroeconomic factors associated with the country or countries which use and support a certain currency [2]. The economic health of a nation's economy is an important factor in the value of its currency. It is shaped by numerous events and information on a daily basis, contributing to 24 hour, constantly-changing nature of the international foreign exchange market.

These economic statistics are generally obtained through data releases. Releases happen at predetermined and advertised times so that all interested parties gain the information at the same time. This is done because the statistics can have large effects on currency markets. If the datapoint released is surprising (as compared to expert predictions), markets can make large movements in very short periods of time [3]. Because of this, staying up to date with these data can be very important in trading.

2.3 The Four Asset Classes

2.3.1 Equities

Also known as stocks, equities represent shares of ownership in publicly held companies. Stocks are traded on stock exchanges such as the New York Stock Exchange (NYSE) or the Nasdaq Stock Market. At the most simplistic level, stocks can be profited from by buying shares at a certain price, then later selling at a higher price. Though it is less common, one can also choose to short-sell a stock, a process where shares are borrowed and resold. In this case, the seller would hypothetically profit from a drop in price [4]. Profit can also be made through dividends—cash or stock payouts given to shareholders by companies as an appreciation for their investment. Rather than tracking national economic data, fundamental analysis of equities is generally focused on factors associated with the company in question. Though equities are a popular asset for trading and investing, they are not the focus of this report, and therefore will not be discussed further.

2.3.2 Bonds

A bond is a loan made to a government or corporation, which provides fixed income to the investor in the form of an interest rate charged on the borrowed funds. Among asset classes, bonds represent the extreme of low risk and low return. Unless the institution to which money is lent collapses, the investor should continue to receive a fixed income from the asset. Though bonds can be traded, they were not the focus of this project [5]. This section is provided for context and contrast.

2.3.3 Commodities

As an asset, commodities are physical goods such as gold, copper, crude oil, natural gas, wheat, corn, and a host of others, which are invested in or traded [6]. These are traded on commodity exchanges, through which orders are executed. Commodities are tightly intertwined with futures contracts, which are legal agreements to trade a certain commodity at a future time [7]. In this case, the investor must be rid of the contract before it matures, to avoid delivery of the physical good. Once again, this section is provided only for context.

2.3.4 Currencies

A currency is a medium of exchange, usually represented in bank notes or coins. The usefulness of a currency to a trader or investor comes from the exchange rates by which one is converted to another. When currencies are considered as pairs, their exchange rates can be analyzed—from a technical perspective—just like an equity or other fast-moving market. These

markets are collectively known as FOREX markets, short for "Foreign Exchange". The FOREX market has much in common with stock markets: currencies can be bought and sold through brokers, and they can be analyzed technically, fundamentally, or sentimentally (described further in the next section). Whereas a major market like the NYSE has thousands of equities, the FOREX market has only a seven major currency pairs, Euro (EUR), Japanese Yen (JPY), British Pound (GBP), Swiss Franc (CHF), Canadian Dollar (CAD), Australian Dollar (AUD), and sometimes the New Zealand Dollar (NZD), each paired with the United States Dollar (USD) [8]. While there are plenty of other currencies, and even more pairs to be created within the "majors", these pairs between USD and the other major currencies are the primary pairs traded. Unlike stock markets, the FOREX market is decentralized, meaning it is mainly comprised of large brokers buying and selling directly. This leads to the markets being active 24 hours per day, from Sunday afternoon to Friday afternoon in US time zones [8]. Because of interest from team members, FOREX was chosen as the focus of this project.

2.4 Technical, Fundamental, and Sentimental Analysis

There are three main ways to analyze any market: technical, fundamental, and sentimental. Technical analysis is the examination of charts and price movements of a certain asset; fundamental analysis looks into the broader economics of the nation or nations in question (in the context of FOREX); and sentimental analysis is the study of collective opinions of other market participants [9]. All of these approaches can be fairly subjective. Because technical analysis is data-driven, it is often straightforward to implement technical methods in an algorithm.

Technical trading focuses on analyzing the past market data in order to forecast the future price movements. When done manually this typically involves study of price charts to attempt to predict some future movement. Indicators, values derived from price data, are a quantitative means of gaining a different perspective on the data being reviewed. These are often used to help gain insight in manual, technical trading. Similarly, indicators can be used by an algorithm to gain the same insight. The methods described in Chapter 5 are generally based on technical analysis.

2.5 Manual Trading versus Algorithmic Trading

On the surface, the differentiation between manual and algorithmic trading seems fairly straightforward, but this delineation is worth expanding upon. Though manual trading is fundamentally being performed by a human, there are many aspects of the process which are being controlled by computers. Many traders use indicators in order to make decisions. The calculation and display of these indicators, along with the display of the price data (such as the generation of a candlestick chart) is all being done by a program. When the right combination of

conditions occurs, the person may choose to make the appropriate trade. At this point, the rest of the process goes back into the hands of a computer for the completion of the order. This shows just how narrow the divide is between manual and algorithmic trading—whether the final decision of the trade is made by a human, or by a computer. If the trader follows a specific system that they have devised, then these decision points are quite well defined, making them often fairly straightforward to turn into an algorithm. It is that process of concentrating trading concepts into succinct instructions that constitutes a portion of what has been carried out in this project.

Of course, not all traders follow definitive systems; many have such experience that they can simply look at a chart and intuitively know how to trade. Their decisions are based on recognition of patterns developed over a vast amount of input over a long period, and therefore are nearly impossible to model as a process. This is where the use of neural networks comes into play. As the name implies, they are designed to act somewhat as the brain would; recognizing patterns and correlations over a large dataset. This is the inspiration for attacking these problems with a neural network based approach, in addition to more traditionally-minded algorithms. As one might imagine, attempting to generate decisions as a human would is a very challenging problem. Profitable manual trading is challenging enough, and conversion of that type of thinking to a program is far from trivial.

2.6 Machine Learning Approaches

Machine learning is currently one of the most highly researched area of computer algorithms, and this research includes the area of trading. The idea behind machine learning is that a program is given a large number of sample inputs and a desired behavior, and the program learns how to produce the desired behavior based on the input. Machine learning is a very broad and powerful subject which has been used effectively in a large variety of applications. A common way of classifying machine learning algorithms is between supervised and unsupervised. Supervised learning involves input/output pairs being given to the algorithm, which tries to replicate the output based on the given input. For example, an algorithm could be given the recent behavior of the market and asked to predict the next change the market will make. Unsupervised learning involves only input being given to an algorithm which tries to find patterns in the given data. A subcategory of unsupervised learning is reinforcement learning where an algorithm, often called an agent, is given inputs such as the recent behavior of the market, a variety of actions it can take such as buying and selling, and a reward which changes based on which actions the agent takes such as trade profit. The agent then tries to learn which actions it can take to maximize the reward it receives.

2.6.1 Neural Networks

Neural networks are a type of machine learning algorithm which make use of linear algebra and various transformative functions to extrapolate information about a set of given inputs. A vector of values, usually a decimal from 0 to 1, is input at one end and put through a series of layers to reach the output of the network. On each layer there are many virtual neurons. For each neuron on a layer, the output of each neuron on the previous layer is multiplied by a value called a weight, and the sum of those values is put through a transformative function, usually a sigmoid function. In the end this produces a vector of values from 0 to 1, each of which represents a different derived attribute of the given inputs. In order to train a neural network, a backpropagation algorithm is used to adjust the weights between layers in order to make the actual output more similar to the desired output. The goal of a neural network is to reduce cost, which is based the difference between the actual and the desired output.

The simplest kind of neural network is a feed forward neural network. This type of network simply puts its inputs through a finite series of layers and returns the calculated output. A more complicated type of neural network is a recurrent neural network, or RNN, which is used to predict sequences. The output of each use of the network is fed back into the network along with the next input vector. Since this method is used to predict sequences, this allows the network to feed itself information about the past of the sequence, meaning it does not need to be input again as is necessary with a feed forward network.

2.6.1 Q-Learning

Q-learning is a type of unsupervised learning in which an agent attempts to predict the expected reward to be gained over an indefinite period of time. It does this by attempting to assign a value to states of a system. For example, an agent might choose to assign a positive reward to a state in which the market is trending upwards and a buy position is open. The exact value of the reward which the agent tries to predict is shown in Equation 1:

$$Q(s_t, a_t) = \sum_{n=0}^{\infty} \gamma^n \cdot r_{t+n} \qquad (1)$$

The function Q represents the reward predicted by the agent. On the right, r represents the reward received by the agent after taking action t+n, and ɣ represents the decay of the reward over time and is less than 1. This formula puts more emphasis on attempting to receive rewards sooner since as prediction goes further in the future, smaller and smaller proportions the possible reward received will be received. Once an agent has gained the ability to predict rewards well enough, the optimal action according to the agent will be the one which maximizes the predicted award.

2.7 Utility Theory

Utility is a concept within economics and finance that is used to model value or worth. In modern usage, it normally refers to a utility function that represents "consumers" preferences over a set of choices [10]. A utility function allows for a quantitative ranking of those choices and thus a consumer can make the mathematically best choice. In the context of trading and investing, the trader or investor is the consumer and the set of choices are trades or assets they could invest in. This is usually a function of both money and risk. The function can then be formulated based on the desired risk profile, most traders are risk-averse, meaning that safer money tends to be more valuable than riskier money. Because the outcomes of trades and investment are uncertain, we must use a special type of utility theory, called Expected Utility. With expected utility the outcomes of the possible choices are not known with certainty, only a probability that they might occur.

Mathematicians John von Neumann and Oskar Morgenstern devised the Von Neumann–Morgenstern utility theorem (VNM)[10] to maximize the expected utility of risky choices. They pose the notion of a simple lottery with two possible outcomes A and B with corresponding probabilities p and $1-p$. The equation for a lottery in the case with two outcomes is shown below in Equation 2:

$$L = pA + (1 - p)B \tag{2}$$

This can be generalized for any number of outcomes using Equation 3 so long as the sum of all the probabilities is equal to one.

$$L = \sum_i p_i A_i \tag{3}$$

We can now rank the preference of each lottery based on it's value. For an example of how this applies to trading we can think of a case where we are presented with two potential trades. One trade would result in a 10% profit if successful and a 5% loss in failure with a 57% chance of success. The other trade would result in a 5% profit if successful and a 3% loss in failure with a 75% chance of success. Now evaluating the expected value(EV) of each trade we find the first trade has an EV of 3.25% profit and the second trade has an EV of 3% profit. This means that a trader should choose the first trade over the second. Equation 4 is the VNM and means that a rational agent will preference lottery M to L if and only if the expected utility of M is greater than that of L.

$$L < M \ iff \ E(u(L)) < E(u(M)) \tag{4}$$

With regards to trading and investing, the VNM gives a user the ability to customize their strategy in a few ways. The first of these ways is that a user can establish a "threshold trades" whereby the EV of any proposed trade must be greater than to be worth it. An example of this geared for longer term investments is to compare a proposed investment against the expected value of a risk-free asset. Another way we can customize this approach is by adding a utility function to adjust the value of our "lotteries" based on our desired risk profile.

3 Literature Review

3.1 Prior Work at WPI

This project is, in many ways, a continuation of previous research conducted at Worcester Polytechnic Institute [11], [12], [13], [14]. The following chapter is a brief summary relevant related research.

3.1.1 Trading Strategies

In Forex Trading System Development (2018) [15] Ziyan Ding devised a "Range And Trend System". The theory behind the system was to use, Keltner Channels, Triple Exponential Moving Average, and Relative Strength Index, indicators to classify the market as either trending or ranging and trade accordingly. The entry/exit conditions follow the conventional triggers for the indicators, however their exact combination is an invention of the author. The system features built in parameters that can be optimized over time to conform to current market behavior. Monte Carlo simulations of the system show an impressive 32.86% profit, 2.881% maximum drawdown, with 95% confidence.

In Automated Foreign Exchange Trading System (2011)[16] Saraffconn et al, tested four trading strategies based on well known technical indicators. The four strategies were:
- ➢ Double moving average crossover, which looked for times when moving averages of two different lengths crossed over.
- ➢ Bollinger Band/ Keltner Channel Counter-Trend, which looked to trade when the price hit support or resistance. For example if the price exceeded the upper Bollinger Band it would sell.
- ➢ CCI Counter-Trend, this strategy traded on the traditional triggers for the CCI indicating, buying if oversold and selling if overbought.
- ➢ CCI and Trade Volume Breakout, this strategy tried to detect a price breakout based on the CCI and volume oscillator and ratio indicators.

The authors backtested the strategies over the course of one month and found the CCI and Trade Volume Breakout strategy to perform the best. It achieved a profit factor and 3.12 with 72.73% of trades ending profitably.

4 Objective of Research

The goal of this Interactive Qualifying Project is to eliminate the psychological and emotional aspects of manual trading by creating an autonomous trading system which is capable of making short-term trades of foreign currencies. The system will be driven by a combination of scientific and thoughtful technical analysis, and by neural networks. Building on the background information provided in the preceding chapters, this report details several solutions which use neural networks and other artificial intelligence approaches.

There are limitless numbers of algorithmic approaches to trading, however it does not seem possible that a single, realistic algorithm, which would work in all cases, can exist. The goal of this project is therefore to create a set of algorithms that can each function under certain conditions. If time allows, then an overarching system may be created to coordinate these individual algorithms into a single platform which can function under a broader set of circumstances than its constituent systems otherwise would.

5 Methodology

5.1 Setting up a Trading Framework

As a part of our research we wanted to experiment with using neural networks for some of our trading systems. Because of this, we decided we wanted to be able to write our trading systems in Python which has a large library of neural network and machine learning libraries including TensorFlow, Scikit, and Numpy. We also decided that we wanted to use MetaTrader as our broker because of its robust backtesting system which was better than any other broker we looked at. In order to achieve both of these goals we needed to build an interface between MetaTrader and Python.

As a starting point we used a basic system made by Darwinex which had an interface for sending prices between MetaTrader and multiple other languages [17]. This interface was made to only send price information, not to make trades. It made use of the ZeroMQ library, available in multiple languages, to create TCP ports through which the programs could communicate. A server running in MetaTrader would receive requests for data in one TCP port and send replies to those requests in another. A client running in Python could then send requests and read the data received. We decided to expand on this system to fit our needs.

Figure 5.1 shows the final architecture of our interface between Python and MetaTrader. The server, written in MetaTrader, is on the left, and the client, written in Python, is on the right.

Connecting them are two TCP ports: a push/pull socket and a push socket. There are two main types of methods available to the client. The first is a push/pull method, such as the tick method which will be discussed later. This type of method sends a message to the push/pull socket and waits until the server sends a response, which the method uses appropriately. The second type is a push method, which again sends a message to the push/pull socket and waits for a response. However, instead of the contents of the message received through the push/pull socket being used, the fact that the message was received is used as a signal that information is available in the push socket which the method pulls and uses.
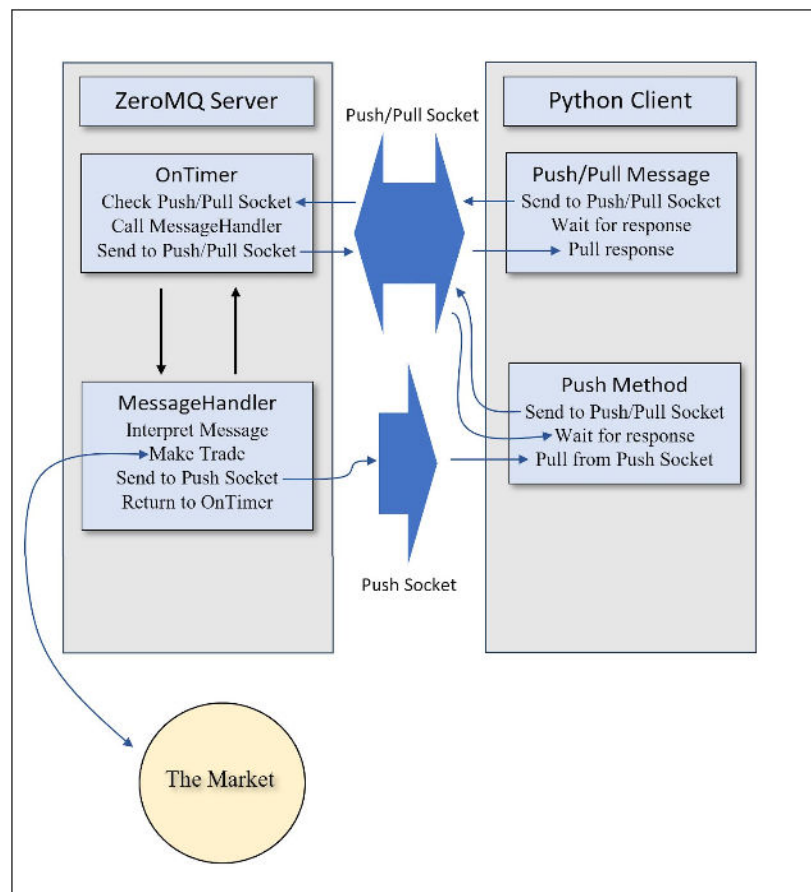

Figure 5.1: Python/MetaTrader Interface

In order to fulfill the requests from the client, the server checks the the push/pull socket for messages on a timer of 1 millisecond. When a message is received, the server sends the message to a message handler which makes trades or sends information to the push socket depending on what the request from the client is. Once the message handler has completed the task requested, it returns to the timer method which sends a response to the client.

The most difficult part of creating this system was ensuring that it worked both when testing on the live market and when backtesting using MetaTrader's built-in backtesting system.

This is difficult because the onTimer method does not execute when backtesting is happening. Instead, a method called onTick is run, and once that method finishes running a new tick of historical data is simulated and the onTick method is run again. To get around this, a method called tick was added to the client whose behavior differed on the server depending on whether it was running live or in backtesting mode. When the server is running in backtesting mode it waits for messages requesting data or trades from the client without stopping the onTick method, preventing new ticks from being simulated. Once the client executes the tick method and the tick message is received by the server, the server allows the onTick method to complete, causing a new tick to be simulated. When the server is running on the live market it operates as before, with ticks happening in real time and the tick message not changing that.

The tick method is also used to send information about ticks as they occur. The tick method maintains a backlog of the last 100,000 ticks and the last 10,000 bars which have happened. The tick and bar information is sent by the server through the push/pull socket when a tick message is received. When the server is running in backtesting mode, information about exactly one tick is sent when the tick message is received, as well as information about the previous bar if it was just created by the newest tick. When the server is running live, a list of all ticks and bars which have happened since the last time tick was called is maintained and sent when tick is called. This is necessary because, since the data is coming in real time, any number of ticks may happen between calls to the tick method depending on how fast the client processes.

This method of handling historical data is important for several reasons. First, it allows for the client to run much faster. Instead of the client having to request every piece of data it needs from the server, most of the important data is stored on the client where access is much faster. Even the simplest trading strategies need data from the recent past, but with the amount of data stored by this system, it is very rare than a client will need to manually request any data from the server. Second, there is no good way to request data when the server is backtesting. There are methods which can be used to request data in MetaTrader, but they do not work properly in backtesting mode, meaning this is the only way that clients made to backtest can run. Third, there is no other way for a client to find out how much time has passed when it is running in live mode. If a client requests the current price of a pair twice in a row and receives the same value both times, it has no way of knowing whether the pair changed in value then changed back during the time it took to ask for the current price again or the pair simply has not changed at all. However, with the server tracking and storing every tick, the client can find out exactly what happened since its last call of the tick method with no uncertainty or loss of information. For these reasons, we believe we have developed a robust, but easy to use system for interfacing Python and MetaTrader which could be used by others in the future to research more complex trading strategies in an environment familiar to most coders and with many libraries which are more versatile and of higher quality than those available in the MQL language.

5.2 Implementation of Trading Systems

In following with the goal of the project, multiple algorithms were developed in an individual, but coordinated fashion so that they could potentially be combined into a single system at a later date. This section details the development of these approaches.

5.2.1 Approach with Recurrent Neural Networks using Long Short-Term Memory

Developed by Thar Min Htet

Overview

This trading system is implemented using Recurrent Neural Networks (RNNs) with Long Short-Term Memory (LSTM). The reason for choosing this over other types of neural networks is the ability of RNNs to handle time series data such as financial data very well. LSTM is an add-on to improve the performance of RNN. The details of RNNs and LSTMs will be described in the following section. This model is not designed to trade currencies live, but the purpose of this model is to be incorporated into other trading strategies and support different systems. With further development, such as extending into MetaTrader, this can be used to live trade currencies.

The workflow for the system is summarized as follows:

1. Acquire the currency price data.
2. Preprocess the data.
3. Extract features using Autoencoders.
4. Train the data on LSTM and RNN.
5. Tune Hyperparameters.
6. Test the accuracy with both training and testing data.
7. Make useful prediction on current currency prices.

This flow of events is depicted in the following flowchart, see Figure 5.2. It begins with finding and downloading currency data that will be used during training. Then, as part of preprocessing, the data is denoised using a wavelet transform formula, and transformed using autoencoders, which automatically extracts features. Then the data is divided into training and testing. Next, the training data is fed into the LSTM model, which is fine tuned until the desired accuracy on the training is achieved. If satisfactory, the model is tested on the testing data. At this step, if the model fails to achieve satisfactory accuracy, all the steps starting from feeding the LSTM model is repeated.
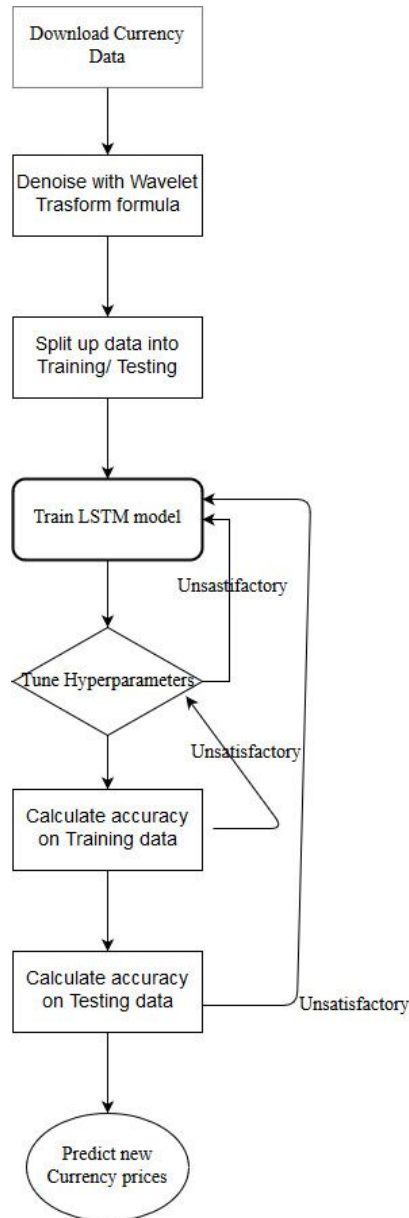
Figure 5.2: The complete flowchart of training currency data using LSTM

Data and Features

As the saying goes in engineering "Garbage in. Garbage out", one of the most essential parts of a machine learning model is the quality and the relevance of the data. A model is only as good as the data it is fed. The system is trained on an old dataset of EUR/USD. The reason was purely because the transformed data which was processed by a professional was found along with the original data, giving the developer a lot more precision and convenience.

In terms of features, common features such as time, open, high, low, close, and volume are used to train the models. However, more advanced features such as values from indicators such as RSI and MACD can be used for potentially more accurate and relevant predictions.

23

Nonetheless, since such data isn't as widely available as simple price data, automatic features are used instead using the autoencoder. The details behind the autoencoders, however, will not be included since they were found online.

Data Preprocessing

Due to the complex nature of currency markets, data is often filled with noise. This might distract the neural network in the learning process of the trends and structure of the price. One solution to this issue is to preprocess the data by removing some of the noise, "cleaning" the data. In this case, a Wavelet Transform is used because it helps preserve the time factor of our data, and an open-sourced implementation of the formula was found:

$$X_\omega(a,\ b)\ =\ \tfrac{1}{\sqrt{a}}\ \int\limits_{-\infty}^{\infty} x(t)\psi(\tfrac{t-a}{a})\ dt \tag{5}$$

where $a \in R^{+}$, $b \in R$, t represents time and $\psi(t)$ is a continuous function in both the time domain and the frequency domain called the mother wavelet and the overline represents the complex conjugate operation. The preprocessing of data starts by transforming the data using a Wavelet Transform. Then, coefficients that are more than a full standard deviation away from the mean are removed. Then the inverse transform is performed on the data to give the cleaned result.

Technology behind the system

RNNs are often proficient in tasks that involve sequences. Since these datasets are time sequences of financial data, an RNN is a logical choice. LSTM is a small, complicated add-on to an RNN in order to prevent technical errors such as exploding and vanishing gradients, which essentially cause the predictor to fail due to the parameters being too small (vanishing) or too big (exploding). Moreover, LSTM has the ability to memorize some value in the past of the sequence because of the memory cells in it, allowing the prediction to relate further into the past of the sequence. Figure 5.4 shows the general structure of an arbitrary LSTM. Figure 5.5 shows a single LSTM cell in details.

Figure 5.3 shows the architecture of a common LSTM model with LSTM as a black box. After the network is trained, the prediction of the sequence works as follows. Firstly, the first value of the input to the sequence ($a^{<0>}$) is defaulted to 0. Then, the first LSTM cell, takes in the input and multiply with the trained weights, and feeding the outcome to an activation function (softmax, in this case), resulting in the outcome for the first value of the sequence $y^{<1>}$. $y^{<1>}$ is used as input to next value in the sequence $x^{<2>}$. And hence, $y^{<t>}$ is always predicted using $y^{<t-1>}$ and the activation value $a^{<t-1>}$ and the pre-trained weights of the network.

Figure 5.4 depicts the inside of a single LSTM cell. In this diagram $h^{<t>}$ is the same as $a^{<t>}$ from Figure 5.3, representing the output value from the cell and also input to the next cell. $c^{<t>}$ is

a notation that is used to represent the memory of a cell. The cell takes in $c^{<t-1>}$ which is the current memory of a certain piece of information and decides whether to keep carrying it or not and outputs $c^{<t>}$ to the next cell. Please also be aware of the multiplier and adder units in the figure which correspond to the operations between the outputs from activation functions and $c^{<t-1>}$ to produce $c^{<t>}$ and $h^{<t>}$.
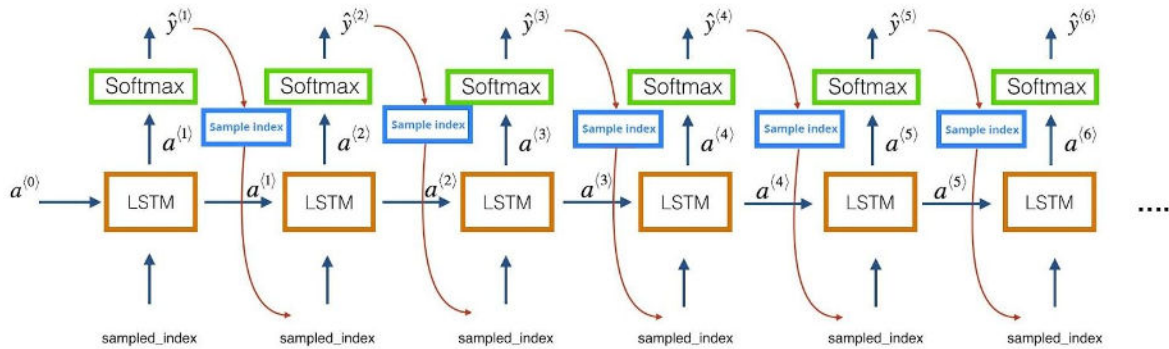
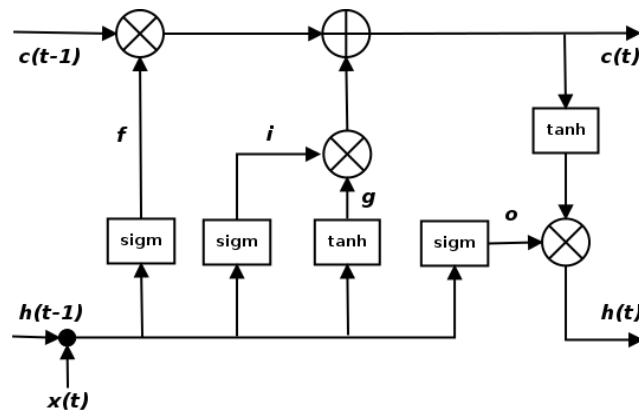

Figure 5.3: Architecture of Long Short-Term Memory

Figure 5.4: A single cell of an LSTM

In this system, the LSTM model has five layers in total, with three hidden layers. The first layer has 50 neurons, and applies sigmoid activation function. The second and third layers both have 20 neurons each and applies tanh activation function. The fourth layer contains 100 neurons. The last layer contains 1 neuron, which is the output value. These are all implemented using a popular machine learning framework called Keras [18].

To give some background about activation functions, Softmax is an activation function that takes as input a vector of K real numbers, and normalizes it into a probability distribution consisting of K probabilities. Sigmoid produces a value between 0 to 1 or -1 to 1.

Tanh a hyperbolic function whose output ranges from -1 to 1.

Training the Model

  Before training the model, the data is split into the training and testing sets. The model follows a popular convention of an 80/20 ratio of training and testing data. This was accomplished by using popular python libraries for data analysis such as Panda [19] and scikit-learn [20]. The following is the code snippet for splitting the data.

```
X_train, X_test, y_train, y_test =
sklearn.model_selection.train_test_split(X, y, test_size=0.2)
```

Where X =[Time, Open, High, Low, Volume] , y = [Close], and test_size = 0.2 refers to a 20% test data set size.

  The other popular way of splitting data would be train/validation/test. However, for simplicity, only train/test is used. Figure 5.6 shows one method of splitting a dataset.
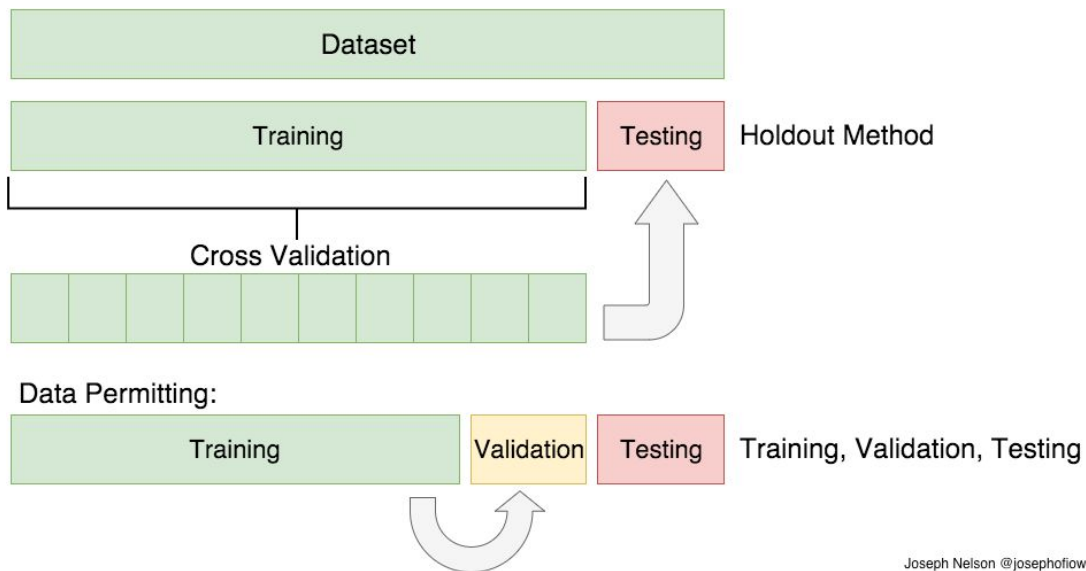

Figure 5.5: Visualization of splitting datasets [21]

  Figure 5.5 also features validation sets. Validation set can be considered as part of training set that is used as test set during training of the model. The advantage of using validation is the prevention of overfitting from hyperparameter tuning. However, this is not implemented in the system for simplicity.

5.2.2 Automated Momentum Indication via Candlestick Patterns

Developed by Andrew VanOsten

Looking for patterns in candlestick charts is a common practice among traders. Just like any other indicator, these patterns are never perfectly accurate, but they can still provide a valuable indication of trend reversals. These patterns allow the trader to quickly see certain movements of the market through easy to remember shapes. Just as these are relatively easy to read by human eyes, they are very simple to search for with an algorithm. They are defined by the open, high, low, and close prices for each bar—data that are easily available and easy for an algorithm to compare. By using somewhat agreed-upon definitions of common patterns, an algorithm can quickly search an incoming stream of data for potential signals.

These signals each represent the potential for a change in the direction of a trend. However, most traders would not base a trade on this signal alone; similarly, this algorithm takes a slightly deeper approach. A market which is in a downtrend is likely to see several bullish signals as it nears its low point, before reversing and moving upwards. This algorithm looks generates a small positive signal for each bullish signal, and a slightly negative signal for each bearish signal. Different patterns can be given individual weights to allow for differing reliability of each pattern. These small signal values are simply added to an overall momentum "score" which is used as a signal to buy or sell. As time moves forward this score will decrease exponentially, as the signals seen previously will be less accurate as they grow further into the past. The final value of the score at the n^{th} time-step is defined as:

$$M(t_n) = D \cdot M(t_{n-1}) + \varepsilon \qquad (6)$$

where $M(t)$ is the momentum score being generated, t_n is the current time-step, t_{n-1} is the previous time-step, D is the decay constant where $0 < D < 1$, and ε is the signal being generated—a constant proportional to the reliability of the pattern that has been observed. This function was created in order to model the desired behavior of the system. When $M(t) > 1$, one lot of the selected currency is bought, and when $M(t) < 1$, a lot is sold. When $|M(t)| < 0.5$, any currently held orders are closed, as this represents a loss of confidence in the indicator. Only one position of one lot may be held at any time.

Disregarding the effects of further input from generated ε values, adjustment of the D value changes the rate at which the $M(t)$ value decays. This, in effect, changes how conservative the algorithm is; a high D value makes the program reluctant to close an order (putting more trust in the score, at the cost of higher risk), while a low value causes held positions to be closed after a shorter period, limiting risk at the cost of lower potential reward. The weights given to each candlestick pattern can also be tuned to some extent. In general, a pattern of a single bar is given a low weight, while a longer-period pattern of three bars is given more clout.

As a compliment to this system, an additional option was added to the program to maximize the profit made per day. As each order was closed, the profit or loss generated by that trade was calculated and logged to a Comma Separated Value (CSV) file, along with the time of day, stored as minutes, beginning at 5:00 PM Eastern Standard Time. This timing convention was used as this is the format used by MetaTrader 4. This logging allowed for instant analysis when a test was completed. When profit/loss was plotted against time of day, one could observe the best times of day during which to trade. This was often quite evident when viewed on a plot. With this information, the system can be instructed to only trade during selects time periods. Because the system relies upon the trends of a market, this layer was quite beneficial to avoid trading during low-volume periods.

5.2.3 Q-Learning using Neural Networks

Development by Ben Anderson

This system attempted to use Q-learning along with neural networks to assign values to potential states which the market could have. The first attempt used pure Q-learning. A feed-forward neural network was fed information about the last 20 bars in the market, as well as the current trade, if any, which was in place. In order to make the information about the market more manageable it was rearranged so that all values were relative to the previous bar or other parts of the current bar. The opening price of each bar was represented as the difference between that opening price and the closing price of the previous bar, the high was represented as the difference between the high and the open, the low was represented as the difference between the low and the close, and the close was represented as the difference between the open and the close. Representing these values in this way allowed the system to consider each state independently without having to account for differences in the range of input data. Once all data was fed in, the network was trained using equation 4.

$$loss = |Q(s_t) - (r_t + \gamma \cdot Q(s_{t+1}))| \tag{7}$$

In this formula Q(s) is used to represent the predicted reward of a certain state, r is the received reward, and ɣ was the decay factor which was set to 0.9. The neural network was trained to minimize this loss function for 5 million iterations.

A slightly different approach of using Q-learning was used to make a second neural network. This design came from the realization that in order to go from making a prediction about the future behavior of the market to predicting a reward, the neural network had to do a multiplication of the current amount traded and the predicted change in the market. This is possible, but unnecessarily complicated when it is known that the multiplication is necessary. In light of this fact, a second neural network was designed where it was always assumed that one lot would be bought. The value of a state could be calculated by multiplying the actual trade size by the value of the state where one lot was owned. The same formula was used to calculate loss for

this network, but r was replaced with the change of the next bar. This network was trained for 10 million iterations.

5.2.4 Using Expected Utility Theory for Asset Allocation

Development by Nick Nugent

As covered in Section 2.7 the Von Neumann–Morgenstern (VNM) utility theorem can be used to rank the value of choices, and so we decided to try to use it to combine our individual trading strategies into one system.

First, let us look at the simplest case of this application. Suppose we have a number of algorithms that propose n number of trades at the same time. Each trade has a known take profit and stop loss, and the historical win rate of the algorithm is known. From there, the expected utility of each trade can be evaluated and the trades ranked using Equation 2. Then we can discard any trades with a negative expected utility and size the trading according to the ranking where the top ranked trade is the largest.

We conducted a proof of conduct on this concept to verify it could work. We made a mock up system of "algorithms." These algorithms merely generated trades with random take profits, stop loses, and probabilities of success. For the test shown in Figure 5.6 the profit and loss was determined randomly based on a Gaussian distribution with a mean of 2% and standard deviation of .3% and the probability of success with a Gaussian distribution with a mean of 50% and a standard deviation of 12%. Those trades were then executed twice, once using a naive allocation and once using a expected utility based allocation. For the naive allocation the starting balance, $1000, was split evenly between the three algorithms and they traded independently. Using that approach we see a small profit over 50 trades. For the expected utility based allocation at each step the three possible trades have their expected utility calculated, they are then ranked and allocated a portion of the account balance based on their rank. If the trade has a negative expected utility the money allocated to it is held rather than traded. The result was that the expected utility based allocation performed much better than the naive allocation.
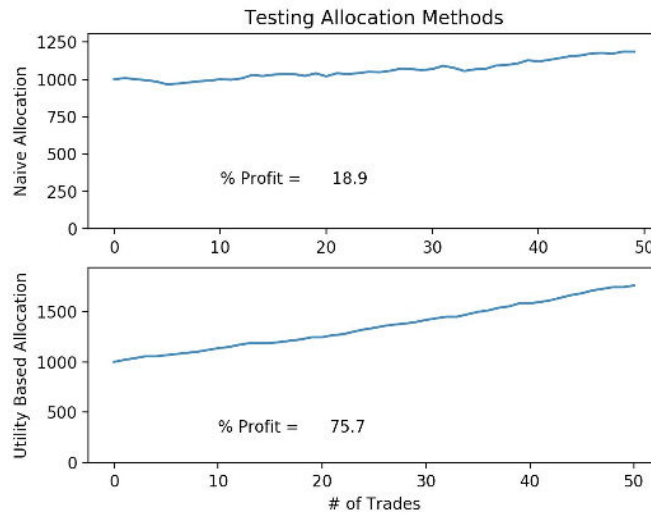
Figure 5.6: Testing Naive vs Expected Utility Allocation.

Even though we have shown that an expected utility based allocation is better than a naive one there are a number of problems with implementing this in our system. The first of which is that unless there was a large enough system of algorithms, simultaneous trade proposals are unlikely to occur, which would nullify the lot sizing aspect of the system. The system would, however, still be of some use. It would act as a filter to block trades with a negative expected utility and thus still increase the performance of the system. The other major problem with implementing this is that it relies on trades with a defined outcome, meaning they must have a take profit and stop loss. For many algorithms, including the ones researched in this report, that is not practical as they have more complex exit conditions.

Because we could not implement expected utility based allocation at a trade-by-trade level we decided to take a step back and look at it as a portfolio investing problem. We essentially have a number of competing assets, the assets being the algorithms, that we want to invest some percentage of our account balance in to maximize our profit. Because we do not have defined outcomes using this method we must modify our utility function so that it can work with historical performance data. There are many utility functions concerning portfolio allocation for a risk averse investor and we will use a simpler one that considers the expected return and volatility.

$$U = E(r) - 0.5 * A * \sigma^2 \tag{8}$$

Here *E(r)* is the expected return of the asset which would be determined from the historical returns. *A* is a risk aversion coefficient from 1 to 5 with 5 being most risk averse. σ is the volatility of the asset, again determined from historical data.

The utility scores of every algorithm would again be ranked and compared to a cut-off value such as a risk-free U.S. Treasury bond to determine if it should be rejected and how much

should be invested. Using this method we would periodically reallocate based on the updated performance metrics of the algorithms.

6 Results

6.1 Approach with Recurrent Neural Networks using Long Short-Term Memory

Since the system was implemented to predict numbers, i.e, future prices, the problem falls under regression, instead of classification. Therefore, the most popular way to evaluate the performance of a regression predictor is to generate charts to compare the values between predictions and training/testing data. Due to the restriction of time and convenience, the system is trained on the old data before 2008. However, the system itself is timeless and should perform in similar way when the most current data is fed to train the system. In Figure 6.1 graphs are generated that highlight the performance of the system. To show the comparison more vividly, both graphs for logarithmic return and market price are generated.
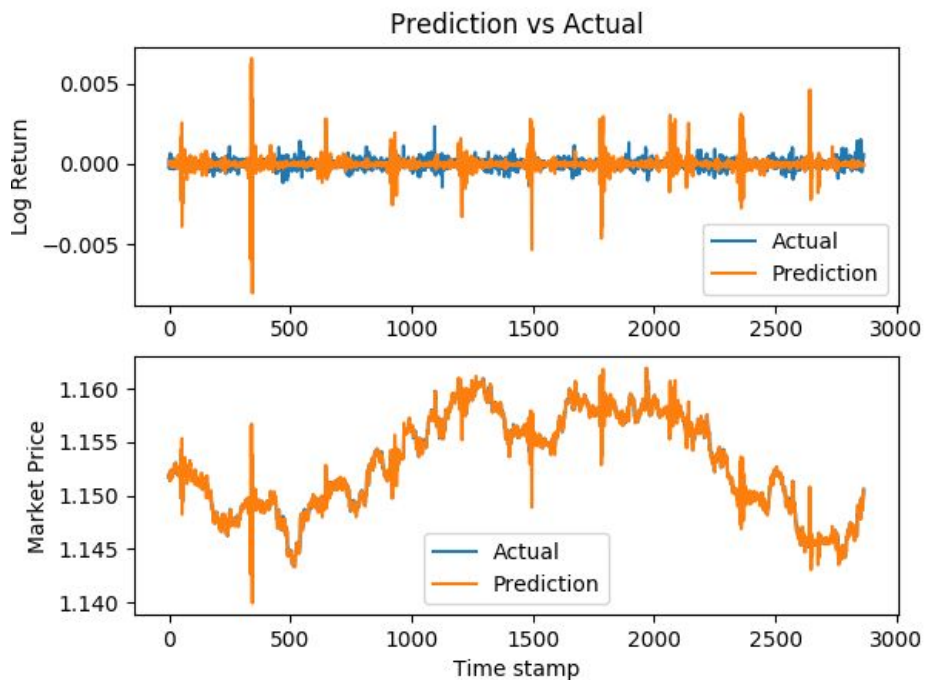
Figure 6.1: Overall comparison between predicted price and actual price

In the comparison from Figure 6.1, although there are some big spikes of noise for the log return, the prediction for market price seems to work pretty well, even overshadowing the actual price. Therefore, to make it clearer, graphs for smaller time duration are generated, as seen in Figure 6.2.

In Figure 6.2 we can observe that the prediction fits fairly well with the actual data in general, except around time stamp 48, which shows a massive gap between the two. This could be caused by the regularization that was done to prevent the model from overfitting the data.
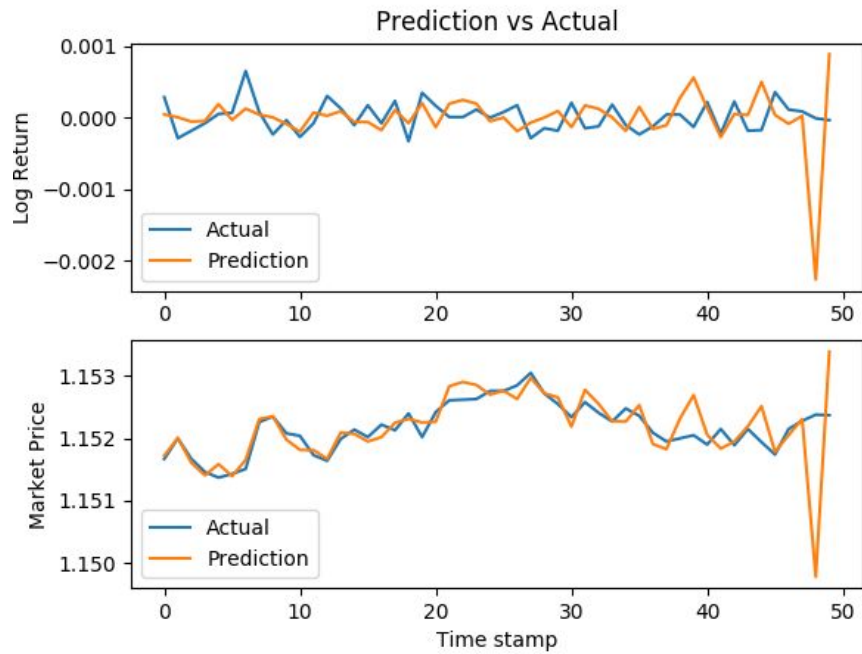
Figure 6.2: Comparison of the first 50 testing examples

Similar to the previous graph from Figure 6.2, in Figure 6.3 the model does a pretty good job fitting the prediction line. The big error at the very start is a continuity from the previous graph, and is likely to be caused by the same reason.
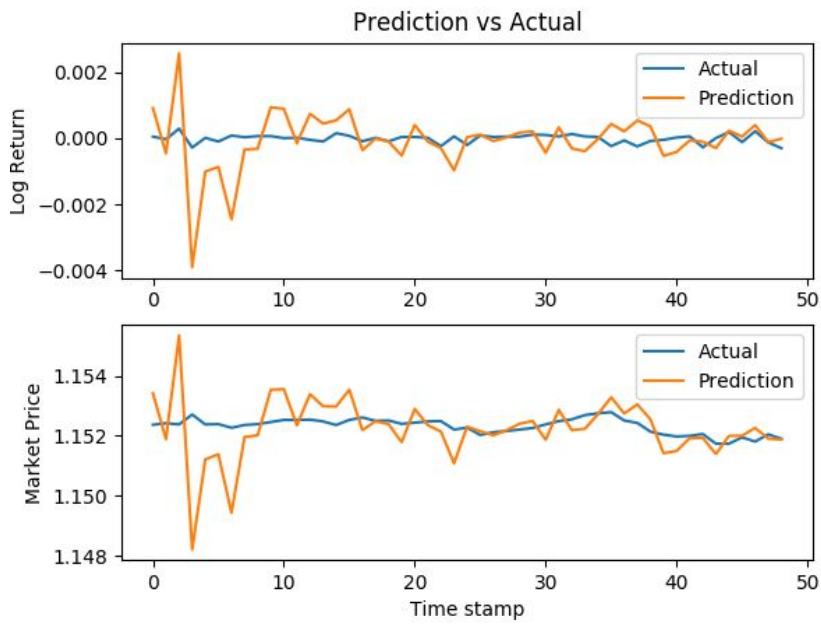

Figure 6.3: Comparison of the second 50 examples

In Figures 6.4 and 6.5, the model performs pretty well at the end of the dataset. However, the prediction for log return (upper subplot in both Figures) seems to be very flat in comparison to the actual data, which has some noise. But because the value from the market price can be used as standalone data in a live trading system, the relatively poor performance of the log system does not matter that much in our case.

As for further concerns, the model seems to be fitting too well, and there is a risk of overfitting. This can be resolved by incrementing the L2 regularization factors and dropout weights, in the machine learning techniques used in this model to prevent overfitting.
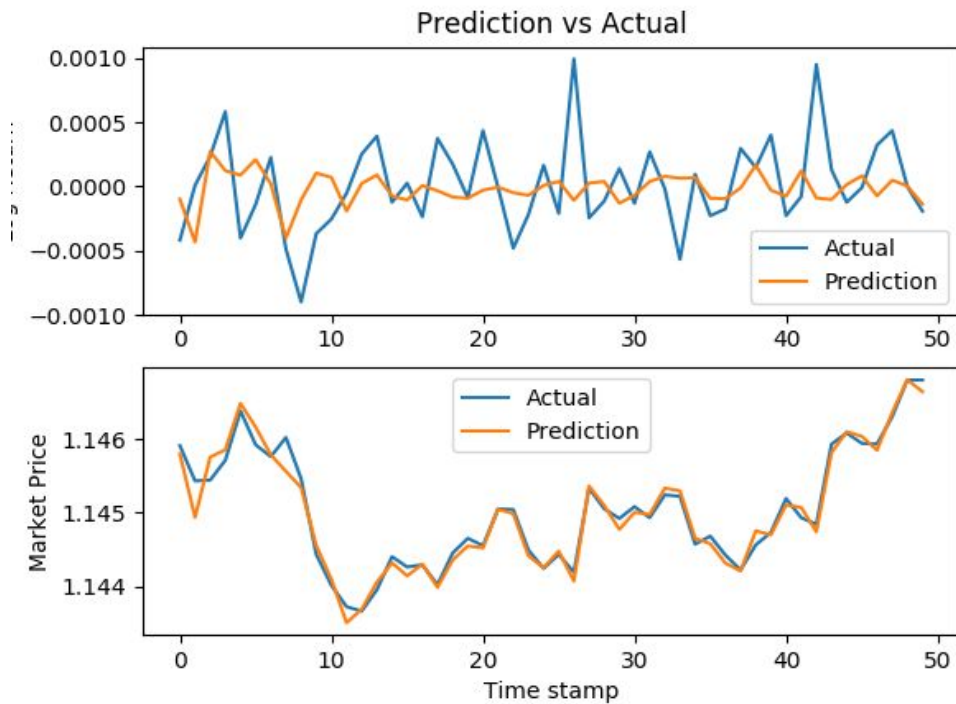

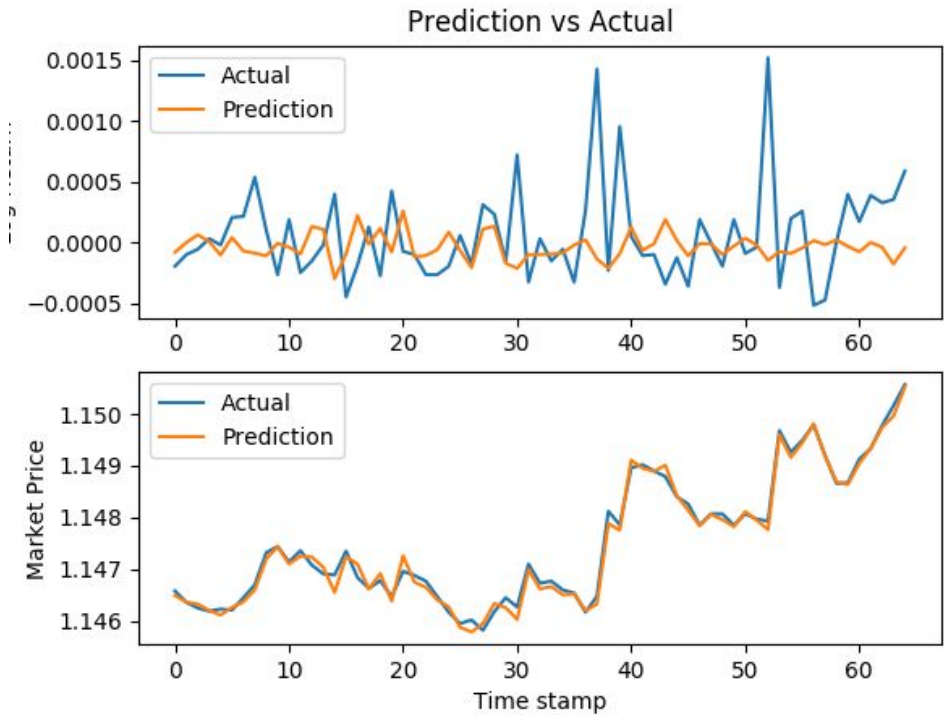Figure 6.4: Comparison of the second to last 50 examples

Figure 6.5: Comparison of the last 62 examples

6.2 Automated Momentum Indication via Candlestick Patterns

As described in Section 5.2.2, a system was implemented using Python and linked to MetaTrader 4 via the interface described in Section 5.1. A variety of single and triple candlestick patterns were experimented with. Table 6.1 below describes the patterns and associated weights which were used via trial and error. Note that the Hammer and Shooting Star patterns are equivalent, but bullish and bearish respectively. The same is true for the Morning Star and Evening Star patterns. This is reflected in the opposite values associated with each. For the sake of comparison, GBPUSD 5 minute data was used for this section. Other currency pairs and timeframes could be used, but would likely require different tuning of ε values.

Table 6.1: Chart Patterns Used and Associated Weights

Candlestick Pattern Name	Associated ε values
Hammer	0.3

Shooting Star	-0.3
Morning Star	0.8
Evening Star	-0.8

To complete the implementation of the equation described in Section 5.2.2, a decay value of $D = 0.9$ was settled on for the final form of the system. In a slight addition to the theory layed out in the previous chapter, one additional condition was added to the flow of the program. If the bar being examined at a given time step is a Doji, which I have defined as a candlestick with open and close prices which are within 0.5 pips of each other, then the $M(t)$ value was multiplied by 0.25. This condition was added because a Doji symbolizes indecision on the part of traders in the given market. Accelerated decay of the momentum value at these points improves the system's model by accounting for this indecision.

With these conditions in place, the following results were generated. Figure 6.6 shows the balance of an account which was controlled by the system from 3/1/2019 to 3/31/2019. The system ran on GBPUSD 5 minute data, and began with an account balance of $10,000 USD. After one month, the balance was $9,574 USD—a return of -4.26%.
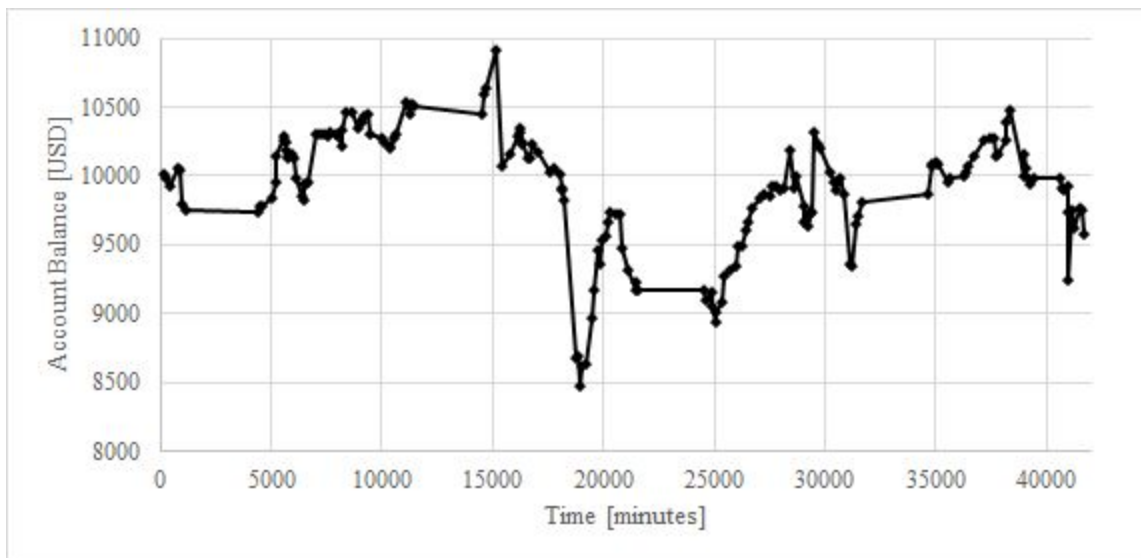

Figure 6.6: GBPUSD-M5 account balance over one month with no time of day restriction

In order to improve performance, the system allows for easy analysis of individual trade profits based on the time of day of the trade. The profit made from each trade in the previously described trial was logged to a CSV file, then plotted vs the time of day represented as minutes since 5:00 PM Eastern Standard Time (EST). The result is shown in Figure 6.7. Each point in the

graph represents a single trade. The horizontal axis represents time in minutes since the simulation started (in this case, 3/1/2019).
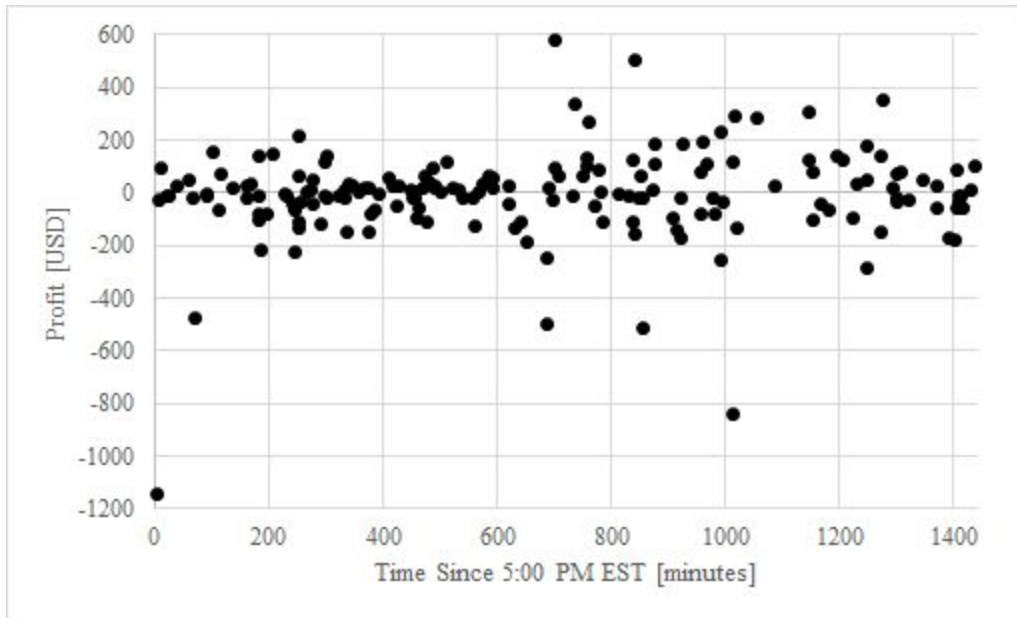

Figure 6.7: GBPUSD-M5 trade profit vs time of day over one month of trading

This figure is not very useful for determining the most profitable times of day to trade, but it does show that the *magnitude* of the profit appears to be greater in the second half of the day (minutes 720-1440). To determine the best way to take advantage of this, the same data can be represented in a different way. Figure 6.8 shows the average profit from trades for various times of day. In this graph, it can clearly be seen that trades in the second half of the time period tend to provide a significantly more positive outcome. These hours, 12 through 24 in the figure, correspond to 5:00 AM to 5:00 PM EST, accounting for most of the London trading session and the entirety of the New York session. This is a logical time for volatility to be relatively high, and for the market to be trending—helpful factors for this system.
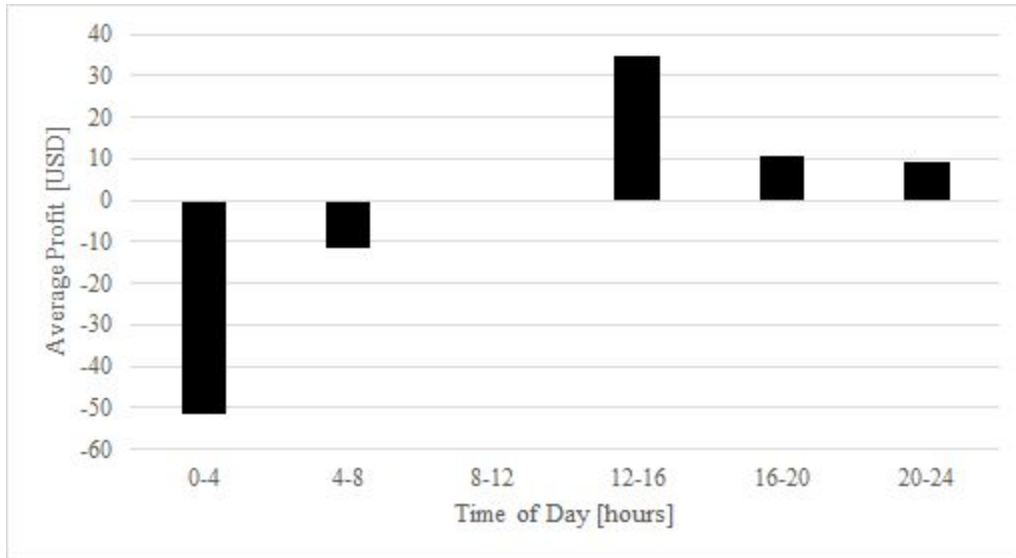
Figure 6.8: GBPUSD-M5 average trade profit vs time of day over one month of trading

The preceding figure provides information that allows the system to be further tuned. One can simply disallow the system from trading during hours 0-12. With this change in place, the account balance for the same time period is shown in Figure 6.9. As shown, the system's account finished the month with $9179 USD, or a return of -8.21%.
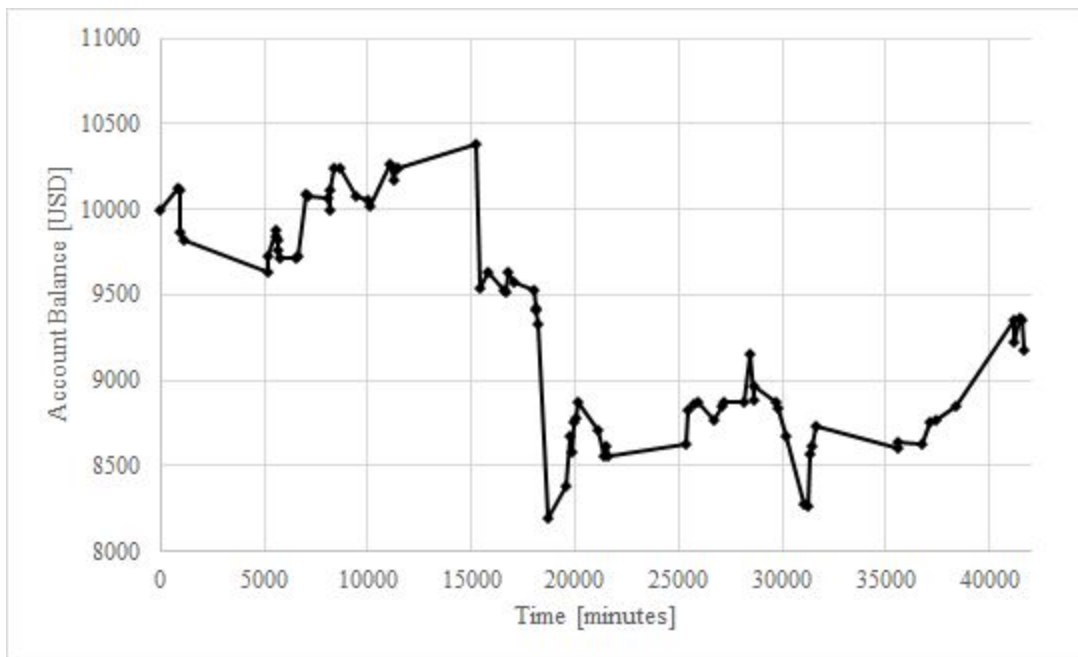

Figure 6.9: GBPUSD-M5 account balance over one month with time of day restriction

In both Figures 6.6 and 6.9, many of the losses seem to happen as large jumps in the graphs, i.e. a handful of highly negative trades. One approach to solving this problem is through the use of a stop-loss. This will cut off any damaging trades before they become major losses. Figure 6.10 displays the results of the system, again on the same dataset, with a stop-loss set at $100 USD (e.g. a trade will automatically close when a loss reaches $100 USD). This trial did not use the time of day restriction.
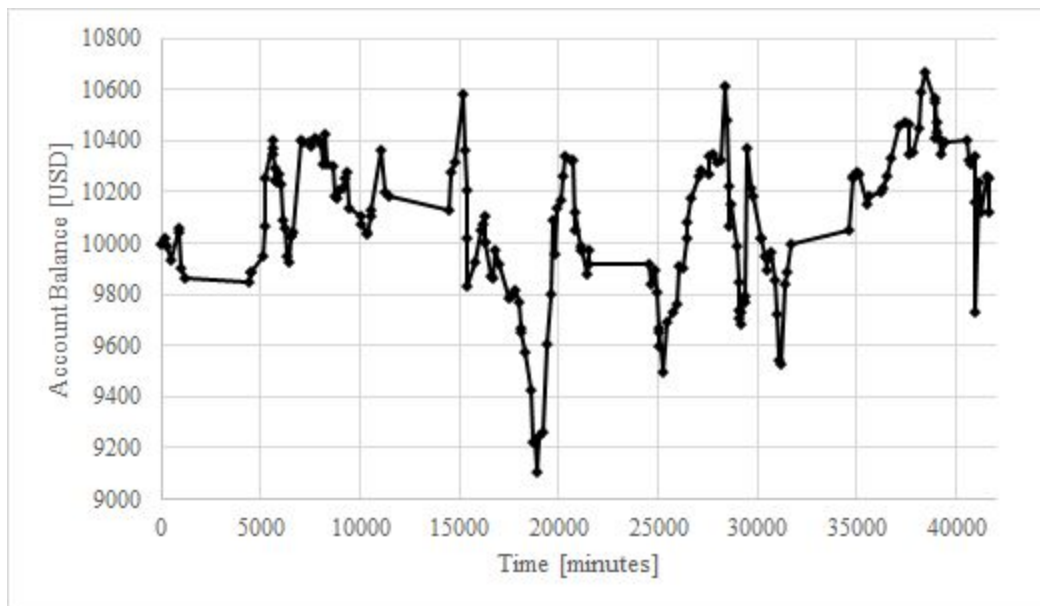

Figure 6.10: GBPUSD-M5 account balance over one month with stop-loss

6.3 Q-Learning using Neural Networks

To test Q-learning using a Feed Forward Network (FFN), the system was allowed to run for 1000 simulated minutes, and allowed to make a trade every bar. No commission or spread was taken into account. For each simulated bar, the system evaluated how many lots it thought was the optimal amount to own. It was then allowed to trade until it had that amount. This was done multiple times on multiple independently trained networks.

The following is the result of a trial. The result of this trial is very similar to the results of all other trials. In total, the agent lost $430. It chose to own one lot 58% of the time, own -1 lots 23% of the time, and chose to own either -0.5, 0, or 0.5 lots the rest of the time. Figure 6.11 shows how little the system was able to predict. The change in the market during the next bar and the predicted value of owning 1 lot are not related ($R^2 = 0.003$). Figure 6.12 helps explain why this problem occurs. If the market is going to go up, buying should be desirable and selling should not be. If the market is going to go down the opposite should be true. However, the system finds that the more desirable owning 1 lot is, the more desirable owning -1 lots is (

$R^2 = 0.6357$). This is a contradiction which prevents this system from being able to make accurate predictions.
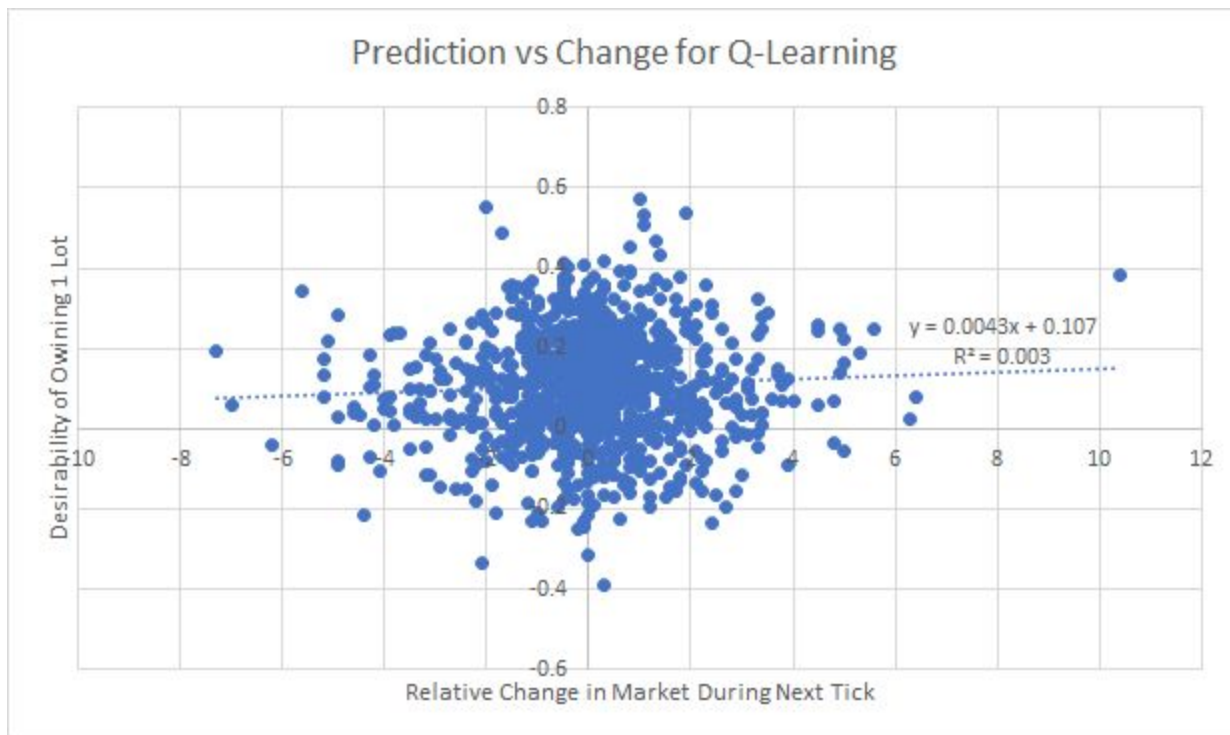

Figure 6.11: Prediction vs change for Q-Learning

Figure 6.12: Desirability of owning 1 vs -1 lots