NTNU
Innovation and Creativity

Automatic stock market trading based on Technical Analysis

Fredrik Larsen

Master of Science in Computer Science
Submission date: February 2007
Supervisor: Helge Langseth, IDI

Norwegian University of Science and Technology
Department of Computer and Information Science

Problem Description

Technical analysis is the forecasting of market prices based on analysis of historical prices in order to find particular patterns that can tell wether a stock theoretically should be higher or lower priced in the future.

The goal of this thesis is to:
-Automatically detect different patterns in stock prize data
-Implement a virtual agent trading based on a combination of these patterns
-Evaluate the trading agent using real-life data

Assignment given: 20. September 2006
Supervisor: Helge Langseth, IDI

Abstract

The theory of technical analysis suggests that future stock price developement can be foretold by analyzing historical price fluctuations and identifying repetitive patterns. A computerized system, able to produce trade recommendations based on different aspects of this theory, has been implemented. The system utilizes trading agents, trained using machine learning techniques, capable of producing unified buy and sell signals. It has been evaluated using actual trade data from the Oslo Børs stock exchange over the period 1999-2006.

Compared to the simple strategy of buying and holding, some of the agents have proven to yield good results, both during years with extremely good stock market returns, as well as during times of recession. In spite of the positive performance, anomalous results do exist and call for cautionous use of the system's recommendations. Combining them with fundamental analysis appears to be a safe approach to achieve succesful stock market trading.

Preface

This master's thesis documents my achievements of the 10th semester of the Master of Science studies in Computer Science at the Norwegian University of Science an Technology (NTNU). It was carried out at the Department of Computer and Information Science (IDI).

I would like to thank my supervisor Helge Langseth for valuable advice and comments during the process. Furthermore, I would like to thank Morten Erga, Trond Vadseth and Tomas Alf Larsen for insightful information about stock market trading and tips concering technical analysis patterns. My co-students deserves a thank-you for being cheerful company during coffee breaks, as does all friends and family that have supported me throughout the implementation and writing process.

Trondheim, February 14, 2007

Fredrik Larsen

Contents

List of Figures

List of Tables

Chapter 1

Introduction

A stock is a share of a company. As an asset, it can be bought and sold, and the possibilty to make money in this process attracts a range of both professional full-time investors and novices. The two main schools for maximizing profits from such trading are called *technical* and *fundamental* analysis.

1.1 Stocks

By issuing shares of itself and sell them to investors, a company can raise money for further investements. In return, investors, based on how big part of the company they buy, gets paid dividends of the companys income and have a voice when crucial decisions about e.g its strategy and the composition of the board of directors are to be made. In addition, prices for a share in a company perfoming good normally rises, making it profitble to sell stocks in a company after a growth in revenues. The possibility to make huge profits when investing in a well-performing company, and the fact that stocks over time has clearly outperfomed other investement objects (such as bonds, commodities and bank deposits) [1, page 22] has made stock trading an extremely popular occupation for people seeking to increase the amount on their bank statements. The following subsections defines some important aspects of stocks as investement objects.

1.1.1 Valuation of stocks

The value of a stock at a given time is the price it can be sold for. This implies that the existing valuation of a company, called market cap, is the current stock price multiplied with the total number of shares issued. There are different approaches to deciding how a company's value should be decided. The Price-to-earning ratio, short P-E, is the most basic and fundamental yardstock. Since 1870, the average value for this ratio has been 14.5 [1, page 96], meaning that a fair price for a stock is 14.5 times its annual earnings. Other ways to value a stock is:

- Expected future earnings: Widely used during the years 1999-2001, when speculative pricing of information technology stocks made the average P-E ratio more than double its normal value [1, page 96]
- Price-to-book ratio (P-B): Prices are decided based on the summarized book value of the company's assets
- Adjusted shareholders equity: The pricing is based on estimations of the actual value of the company's assets, not the recorded one.
- Dividend yield: How big stake the annual dividend is of the stock price

1.1.2 Stock exchange

A stock exchange is a marketplace for trading stocks, where the pricing of a stock is based upon supply and demand. Stocks listed on the exchange are given a symbol name, called ticker, and can be traded during the opening hours of the exchange. Transactions are carried out when a seller's ask price is matched by a buyer's bid price. Figure 1.1 illustrates how stock trading is performed over a stock exchange, examplified by the FOE stock on the Oslo Børs stock exchange. The first line in the figure lists up the last trade price, the change since last days close in both price and percent, the highest bid price, the lowest ask price, the open trade price, the days highest trade price, lowest trade price and the close price of last trade day. Below this line, the accumulated order book is shown, with bid and ask prices in addition to the size of each trade lot.

Siste	Endring	Endr%	Kjøper	Selger	Åpning	Høy	Lav	Forrige
282.00	5.00	1.81	282.00	282.50	282.00	283.50	280.00	277.0

Akkumulert ordrebok for FOE på Oslo Børs					
Kjøp			Salg		
Ordre	Dybde	Bud	Tilbud	Dybde	Ordre
3	6 150	282.00	282.50	5 300	3
3	10 650	281.50	283.00	4 130	3
1	12 000	281.00	283.50	12 300	8
2	14 400	280.50	284.00	9 150	4
1	200	280.00	284.50	1 400	2
29	76 450	Total kjøpere	Total selgere	102 270	87

Figure 1.1: Intraday trade data for a stock

1.1.3 Initial public offering

An initial public offering (IPO) is a mean used in connection with the process of listing a new company for trade at a stock exchange. It means that the company offers a part of its shares available for public trade at a price within a decided range. This price range is normally set by one or more investement banks after a valuation prosess. A stock exchange demands that a certain amount of stocks is in circulation to allow it being listed for trade, hence the need for an IPO.

1.1.4 Short selling

Short selling is a way to make money in a stock when its price is falling. It is realized by borrowing stocks from someone, selling them in the market and buying them back to the lessor at a later point in time. If prices has declined in the meantime, the buying back of the stocks will cost less than the money earned when initially selling them, thus leaving profits for the short selling investor.

1.1.5 Bear and bull markets

Bear and bull are definitions of a nagative and positive market opinion respectively. A bear market is a market in which prices constantly fall only interrupted by temporary up corrections. Similarly, a bull market is a market where prices are in an overall up moving trend.

1.2 Stock trading strategies

Every stock market investor has her own idea of how the most profitbale stocks can be found and at what time they should be bought and sold. Fundamental and technical analysis constitute the two main schools for picking attractive trade objects and their corresponding trade entry and exit points.

1.2.1 Fundamental analysis

Fundamental analysis involves using economic data to forecast prices or gauge whether the markets are over or undervalued. This data can for instance be stock prices compared to actual earning of each stock, measuring actual value of a companys assets compared to their book value, crop reports or the last developement of consumption spending in a specified country.

1.2.2 Technical analysis

Technical analysis is the studies of price activity with the assumption that all relevant information is known by the market players and baked into a stocks pricing at any given point. Based on this assumption, and the belief that the fundamental information and market opinions reflected by the stock price at any given time, will result in recurring price patterns that, by fitting them to past movement, can foretell the future direction. Technical patterns are believed to be good predictors for trading entry and exit points.

Market psychology

Stock market trading is mainly motivated by two factors: greed and fear. When one of them become predominant, human rationality is often put aside. Periods of abnormal increases in stock pricing, popularly called stock market bubbles, have historically been cut short by steep decreases intesified by massive sell-off from overstrung investors that, only days earlier, had euphoric thougts about their stock portfolios prospects [1, page 320]. There is no doubt that human psychology and its herd instinct is an important contributor to the evolution of stock prices. The goal of technical analysis is to understand which human feelings are steering the price directions, and act accordingly. The premises for its success lies in the fact that market fluctuations seem to repeat themselves and establishes behavioral formations to be recognized at later points.

A single stock market investor cannot stop or slow down the giant mass of investors, and in order to make money in the market, the essential part is to pay attention to what the market seems to be deciding regaring the valuation of stocks and act accordingly. In short, one only needs to know when the market is on its top and on its bottom and trade thereafter. Once a trend is started, it lasts until a trend reversal is established [2, page 15]. Technical analysis is the most well-known method for detemining these trends and their reversals.

1.3 Automatic trading system

Following the evolution in computers, technical analysis has achieved widespread use. This is a result of the possibility to easy present continously updated price data, perform mathematical calculations based on them, and complete the trades within seconds via online stock brokers. This fact has raised belief that computer based systems will replace as much as 90% of todays human stock brokers [3].

With the widespread use of automatic trading systems, the theory of technical analysis has good chances of becoming a self-fulfilling prophecy: when a large range of investors e.g find a stock being overbought and decide to sell it on the basis of this, the subsequent decline is intesified because of the massive sell-off.

Every system interpreting technical signals correctly and early enough is a potential money making machine, hence the extreme popularity of and interest for good working technical analysis software.

Chapter 2

Program

To test the validity of computerized technical analysis, a program able to draw stock charts for visualization and evaluate technical signals from their corresponding data has been developed. The chapter describes important aspects of the software design.

2.1 The program

A program for automatic analysis of technical indicators and patterns has been written in the Java programming language with the purpose of testing the reliability of technical analysis upon trading the stock market. An overall view of the programs classes and some important methods is shown in Figure 2.1. The most important classes of the system are discussed in the following subsections.

2.1.1 The StockTerminal class

The StockTerminal class is the starting point of the program and includes the main method which invokes the different functionality. Based on input parameters, it either starts the graphical user interface, a text-based evaluator of trading agents performances over sets of stocks, or training of agents using genetic algorithms (process described in Chapter 4). It is also responsible for holding a list of stock tickers and trading agents used by the program.

2.1.2 The StockTicker class

Objects of the type StockTicker denotes a trading history for a particular stock, and has methods for importing it from local text files or the internet. The class offeres methods for getting highest, lowest, open, close and volume values

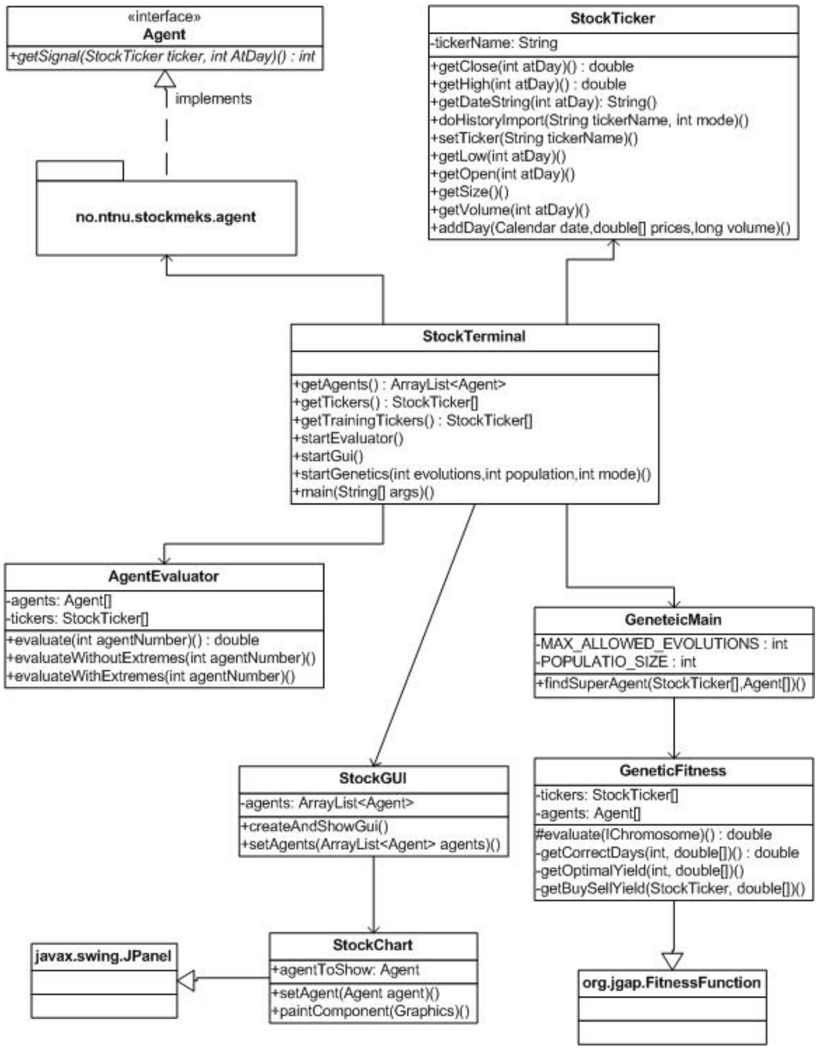

Figure 2.1: Program overview

for each trade day in the stocks imported history in addition to more complex services, like getting the extreme values over specified time periods, whether a specified day was an up close day, and the value of the next future extreme value above a threshold level (if enough future history available), for benchmarking of sets of agent signal over the trade days.

2.1.3 The agent package and Agent interface

The Agent interface, which defines the crucial tasks for each technical indicator or pattern recognizer implemented. The interface defines the following functions:

- **int getSignal(StockTicker ticker)**: Returns the latest signal for a given stock ticker

- **int getSignal(StockTicker ticker, int atDay)**: Get the signal for a given stock ticker at a specified day number in the trading history
- **String getAgentName()**: Returns the logical name of the agent
- **int[] getSignalDays(StockTicker ticker)**: Returns an integer array existing of all historical agent signals for all days of a paramterized stock ticker.

All technical trading rules used by the program implement this interface. It constitutes the minimal necessities of a trading agent and ensures easy adaptability for prospective agent extensions of the program.

All agents used are bundled in the package no.ntnu.no.stockmeks.agent. The package is shown in the programs class diagram, Figure 2.1, and more detailed in Figure 2.2, which also includes some sample agents.

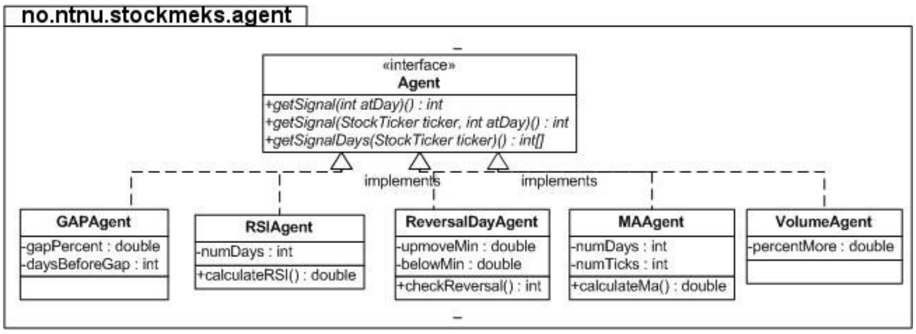

Figure 2.2: The agent packacke and some sample agents

2.1.4 The graphical user interface

The graphical user interface (GUI) class of the program is responsible for the visualization of stock trading histories and the different trading agents appurtentant signals. Its appreance is shown in figure 2.3. The top line of the program has input fields where the stock ticker name and how many trade days ending in the current is to be shown, a drop-down box deciding which import method is to be used (fetching stock history from local text file or internet), and a drop-down box from which a trading agent can be chosen for visualization. A visualized trading agent's sell signal is indicates on the chart by an added little red square at the signal day position, and a buy signal by an unfilled blue circle at the correspoding position. This visualization can be used to benchmark the different agents performance in addition to test their correctness as further described in section 2.1.5.

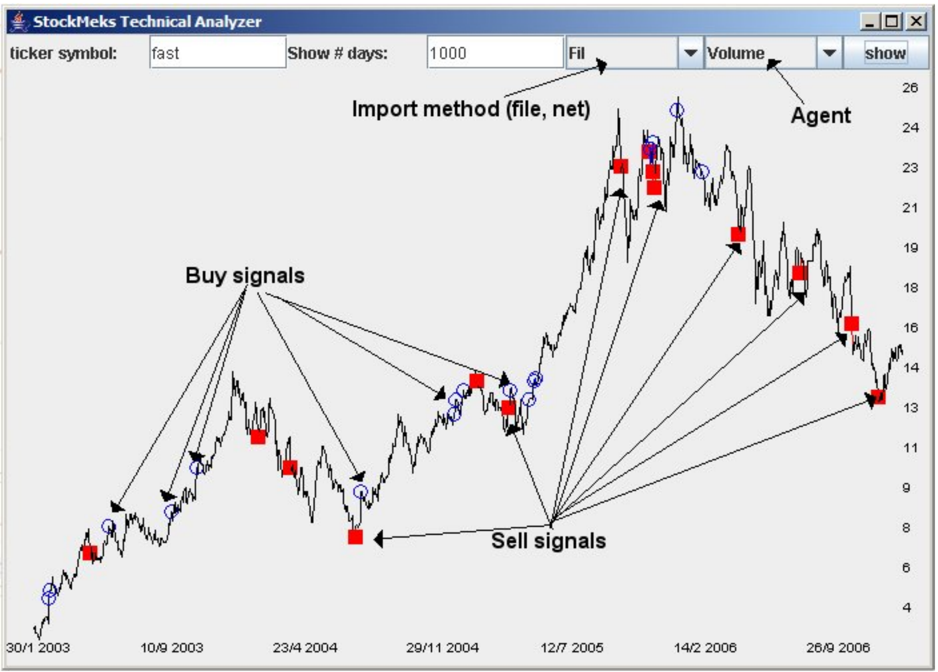

Figure 2.3: The graphical user interface

2.1.5 Testing correctness

Testing of the program and the trading agents has been done using JUnit test cases, the testing framework for Java programs, to assure correct output of computational methods used. For testing of more complex agents, able to recognize chart formations, special stock trade history test cases has been written to ensure that they detect obvious, constructed patterns. In addition, hands-on tests has been performed by visualizing real stock trade histories and testing whether the agents identify the same patterns that a human does. Each signal has also been verified to check that it is not a false positive; if an agent generates a signal, the reason for the signal has been identified in order to justify it. Such a hands-on example including justification is shown for a Double top/bottom pattern (described in Section 3.8) in Figure 2.4. It grounds all patterns recognized but the double top pattern identified inside the drawn ellipse is, however, not recognized with a corresponding sell signal. This can be explained by the formation being slighly crooked, the bottom being somehow lower than the end of the last top before reversing, or the fact that the first top of the formation is formed by a beginning long gapping price move, making it appear bigger than the second top for the agent. A human technical analysist would probably indicate the formation as a Double Top pattern, and this indicates that the implemented agent has some limitations. These observations show the necessity for and the advantages one can derive from testing agents with hands-on examples, including securing correctness and identifying weaknesses, limitations and improvement potential.

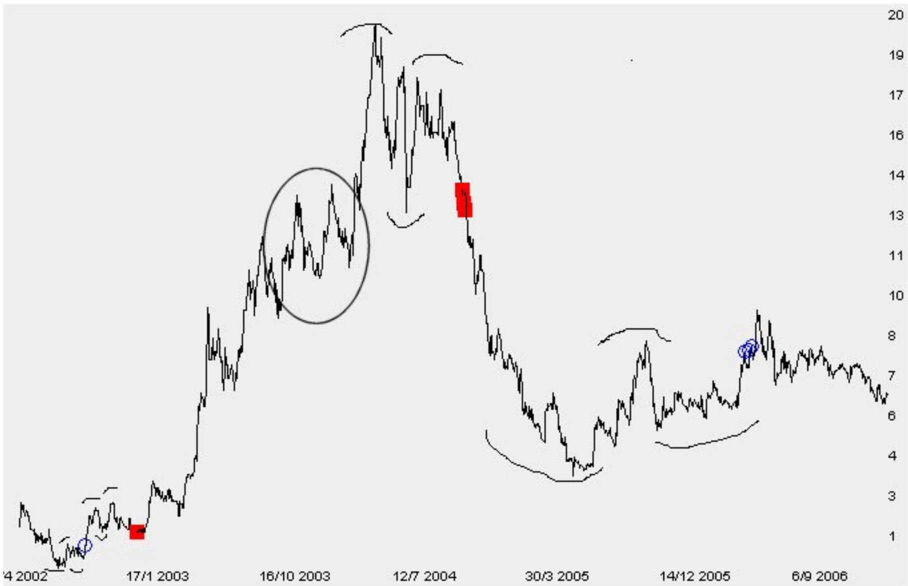

Figure 2.4: Signals and explanation: Double Top for APP stock 2002-2006

Chapter 3

Technical patterns and indicators

To achieve computerized technical analysis of stocks, in total 20 trading agents has been implemented and added to the program, each of them giving trade recommendations when presented with a series of historical price data. Their signals are derived using different calculations or pattern recognizing techniques, described in this chapter. Their parameters for necessary threshold values and indicator calculation basises have been derived from [4] or set using personal assumptions when not defined by theory.

3.1 Trend lines

One of the most well-known sayings of market timers is "Make the trend your friend" [1, page 288], and the drawing of trend line is probably the most well known technical analysis tool used to achieve such a goal. Trend lines is drawn by connecting sets of relative high and low values in a stock's trading history, and the theory states that future prices will move within the range the line connecting sets of local highs and the line connecting sets of local lows constitute [4, page 33]. An uptrend — meaning prices over time constantly moving upwards — is identified as series of higher highs and lower lows. The opposite applies for a downtrend, which is identified by series of lower lows and lower highs. Trend lines are known to be useful, but their importance is often overstated because it is so easy to overestimate the correctness of them when they are drawn with the benefit of hindsight [4, page 35]. Trend lines needs for instance be redrawn as bull or bear markets extends, and penetration of them is often only a signal that the rate of decline or increase in the share price is changing. Still, during a trending market, the increase or decrease in relative extremes gives a good indication of the primary market trend. Up and down price channels for a sample price history made up of trend lines is illustrated by Figure 3.1.

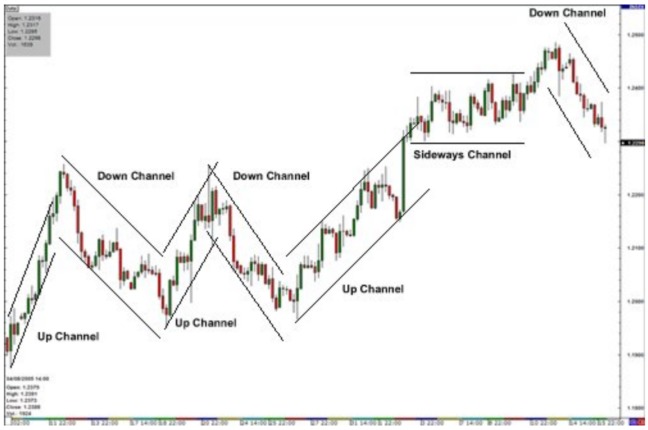

Figure 3.1: Trend lines and -channels [5]

3.1.1 The Trend agent

The Trend agent tries to identify in which direction the market is currently trending. This is done by calculating relative highs and lows and checking whether they are showing signs of moving in one direction. The significant relative highs and lows are found using the following procedure:

Upon traversing the complete trade history: If price at day n is higher than the highest price between the last found relative high (the first trade day in the stock history if no previous relative high is found) and the price at day *n-1*:

- Set the interim highest price to the price at day n

- If the price at day n is at least 30% higher than the previous interim lowest price: Add the previous lowest price to the list of relative lows and set the interim lowest price to the price at day n

An eqvivalent procedure is performed when the price at day n is lower than the previous interim lowest price, adding the existing last relative high to the list of significant highs if the price at day n is at least 30% lower than the previous interim highest price.

When the list of historical significant highs and lows has been gathered, the trend agent checks the direction the previous two relative highs and lows is taking: if they both show an increase, the trend is said to be rising and a buy signal is given. If both the clustered lows and highs are showing a decline, the trend is said to be falling and a sell signal is given. If the high and low trends are not coinciding, a hold signal is returned.

3.2 RSI

Relative strength index (RSI) is a technical indicator based on the momentum, refering to the rate (or speed) at which prices change [4, page 110]. A strong momentum is a signal of a healthy price trend while weakening trends often have stagnant or decreasing momentum. In addition, momentum also highlight shorter-term market extremes or exhaustion points reffered to as overbought or oversold levels [4, page 111]. The logic is that extremely strong and rapid price moves are not indefinitely sustainable and will suffer a at least temporary price reversal. The matemathical definition of RSI is given as:

$$RSI = 100 - \frac{100}{1 + RS}$$

where RS is the average of the total number of N-day up closes divided over the average of the N-day down closes. The number of days used to calculate the RSI decides its sensibility: the shorter the time period, the more sensitive the oscillator becomes and the wider its amplitude [6, page 240]. Figure 3.2 shows a 10 day RSI (bottom chart) for the BHP Billiton stock between 2000 and 2001 at the New York Stock Exchange. Blue dots and red squares has been added to indicate oversold and overbought levels respectively.

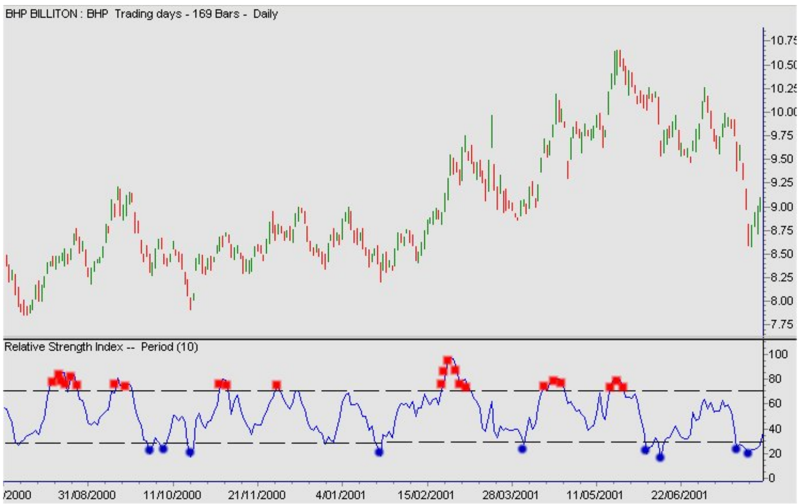

Figure 3.2: 10 days RSI (bottom chart) [7]

3.2.1 The RSI agent

The RSI agent will use the mathematical formula stated above to calculate the N-days RSI value. A RSI value above 70 indicates an overbought stock, resulting in a sell signal. A value below 30 indicates an oversold stock, resulting in a buy signal.

3.2.2 The short-term RSI agent

The short-term RSI agent will calculate the RSI value over the last 9 days, catching the overbought and -sold levels over a little amount of days and therefore risking many false signals. This agent is appropriate for catching minor trends.

3.2.3 The normal RSI agent

The normal RSI agent will calulate the RSI value over the last 14 days. This agent is appropriate for evaluating the relative strengt of the secondary trend.

3.2.4 The longer-term RSI agent

The longer-term RSI agent will calculate the RSI value over the last 21 days, failing to catch volatile short-term movements but is normally more robust in a sideways moving market.

3.3 Gap

A gap day is one in which the lowest trade price is above the previous days highest trade price or the highest trade price is below the previous days lowest and can take different forms [4, page 73]:

- Common Gap: A gap occuring within a trading range and not particularly significant
- Breakaway Gap: Occurs when prices surge beyond the extreme of a trading range, leaving an area in which no trading has occured.
- Runaway Gap: Occurs when a trend accelerates and is defined by the prices gapping in the same direction both before and after the first Gap.
- Exhaustion Gap: Occurs after an extended price move and is soon followed by a trend reversal. The only difference from a runaway gap is hindsight, and when not followed by a trend reversal it is named a continuation gap.

It is necessary to assure that a gap is not filled within the first number of trade days to confirm its validity [4, page 158].

3.3.1 The Breakaway gap agent

The Breakaway gap agent checks whether the following conditions satisfying a breakaway up gap, with 3 days to confirmation of it not being filled, at day n is met:

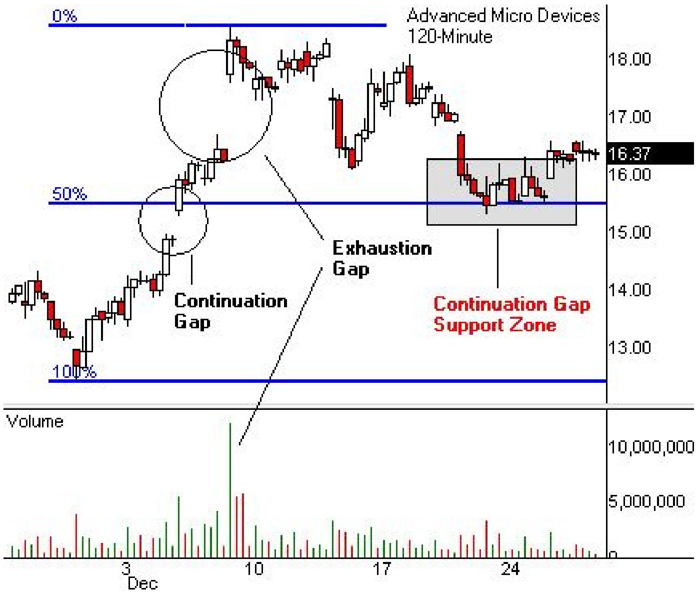

Figure 3.3: Gaps in various forms [8]

- The lowest trade price at day *n-3* are higher than all trades between day *n-4* and day *n-14*
- No trades between day *n-3* and *n* are completed at prices lower than that of day *n-3*

When these conditions are met, the agent returns a buy signal. Analogous conditions are checked for a common down gap: if the highest trade price 3 days ago is lower than all of those 10 days before that and no trades are done the last days at prices higher than the highest 3 days ago a sell signal is returned. Otherwise, a hold signal is returned.

3.3.2 The Runaway gap agent

The Runaway Gap Agent uses the same teqnique for checking for up or down gaps as defined for the breakaway gap agent but adds an extra condition to the equation: the lowest trade price at day *n-3* needs to be at least 1% higher than the highest of the 10 preceding days, thus leaving the gapped price area open.

3.4 Island reversal

Island reversal occur when prices gap, trade one or more days leaving the gap open and then reverses with a gap in the opposite direction. An island top is formed when prices gap higher following a gap low in the preceding days. An island bottom is formed when prices gap lower followed by a gap higher in the preceding days. This sequence is a potent combination which often can signal major trend transitions [4, page 107]. An example of an island reversal is given in figure 3.4.

Figure 3.4: Island reversal [9]

3.4.1 The Island reversal agent

The Island reversal agent uses two breakaway gap agents as defined in Section 3.3 to calculate its signals. The first of them checks for the initial gap that forms the first part of the island reversal and is given the following paramters:

- Minimum days preceding the Gap: 5
- Gap must at least be above: 0%
- Days to check wether the gap has not been filled: 0

The second checks whether a new gap in the opposite direction occurs within a given number of days after the initial one:

- Minimum days preceding the Gap: 0
- Gap must at least be above: 0%
- Days to check wether the gap has not been filled: 1

An island bottom is found when the first gap agent signals sell and the second buy within a range of 10 days counting from the day of the first gap, resulting in a buy signal from the Island reversal agent. This is analogous for island tops which occurs when a buy signal from the first gap agent is followed by a sell signal from the second, resulting in an Island reversal agent sell signal.

3.5 Reversal day

A day that witnessed a new high in an upmove or low in a downmove and then reversed to close below (for upmove) or above (for downmove) the preceding days close is named a reversal day. It indicates a buying or selling climax and is a trend reversal indicator. The reversal day is said to successfully call 100 out of every 10 extremes - it is without a doubt subject to many errors and therefore often considered useless [4, page 80].

A more stric definition is to define a reversal day as a day in which the close is above/below the preceding days high/low instead of its close. In addition, one can also demand that up- or downmove is above a given percentage of the preceding days close. This approach removes many of the false signals created by the standard definition of a reversal day.

3.5.1 The Reversal day agent

The Reversal day agent uses the strict definition of a reversal day as stated above. It checks wether the stock price one particular day has risen above or fallen below 3% of the preceding days close. If this is the case, the agent checks whether the close price is below the preceding days low or high. Given corresponding values for both upmove and close price, the agent outputs a sell signal. Analogous, if the share price was subject to a significant downmove but closed above the preceding days high, the agent outputs a buy signal.

3.6 Moving averages

Moving averages provides a very simple means of smoothing a price series and making trends more discerible [4, page 45]. It is defined as the average price of the last N days, usually the close price. A typical 40 days moving average would be the average of the last 40 days close, ending in the current day. If the moving average is increasing, it indicates an up trend, and when decreasing a downwards moving trend. The degree of how much moving averages smooths a price series is commensurate to its length at the expense of lag introduced. In order to avoid *whipsawing*[1], it is necessary to use minimal threshold values for determening turns in the moving average. Moving averages provides a good tool to spot trends in a trending market. Figure 3.5 shows how trading based on a simple 200 days moving average would lead to an almost perfect timing when trading the Dow Jones Industrials during the years before and after the famous stock market chrash in 1929 (shaded areas indicate a falling moving average).

[1] Trend signal flipping back and forth when moving average is near zero

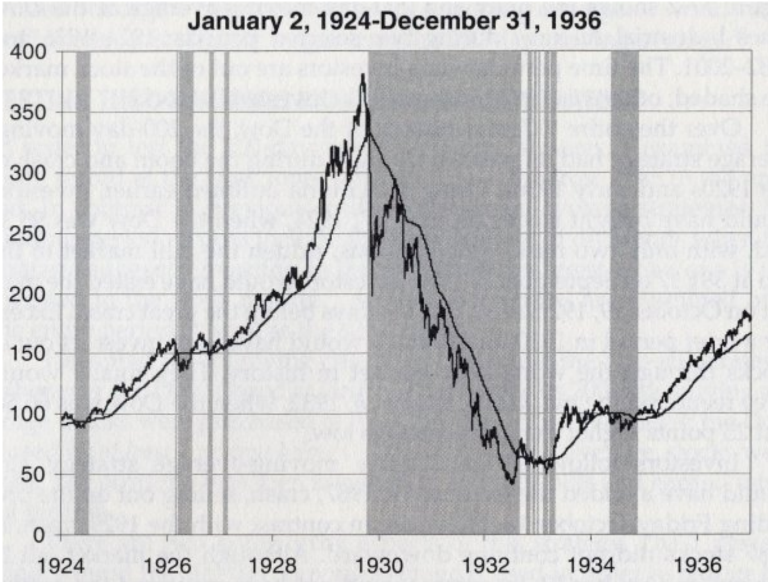

Figure 3.5: 200 days moving average: Dow Jones Industrials [1, page 292]

3.6.1 The Moving average agent

The moving average agent calulate the N-days moving average and check wether an up or down move is the case for a given amount of succeding days, rise in the N-days moving average resulting in a buy signal and decline in the N-days moving average resulting in a sell signal.

3.6.2 The 14 days moving average agent

The 14 days moving average agent is a short-term moving average agent, calculating the moving average over the last 14 days and using threshold value of 5 ticks for calculating signals. This short-term approach is taken to catch the primary trend.

3.6.3 The 40 days moving average agent

The 40-days moving average agent is a normal moving average agent, calculating the moving average over the last 40 days and using threshold values of 10 ticks for calculating signals.

3.6.4 The 200 days moving average agent

The 200-days moving average agent is a long term moving average agent, calculating the moving average over the last 200 days and using threshold values of 10 ticks for calculating signals. This long-term agent introduces much lag in its performance but should catch overall trends with accurancy.

3.6.5 The Crossover moving average agent

The crossover moving average agent uses two different moving averages to calculate signals: When a shorter term (10 days) moving average agent crosses the line of a longer term (100 days) moving average agent and stays above/below for a specified consecutive number of days the agent generates a signal corresponding to the direction of the cross: up crossing results in a buy signal, down crossing in a sell signal.

3.7 Price envelope bands

Price envelope bands is a method for defining support and resistance levels and is derived from moving averages, described in 3.6. The upper band of the price envelope band is defined as the moving average plus a percentage of the moving average, and the lower band is defined as the moving average minus a percentage value of the moving average [4, page 69]. For example, with a percentage level of 3 and a moving average value of 100 the price envelope band would have an upper level of 103 and a lower level of 97. These levels represent resistance and support respectively, and it is exptected that prices stay within them. The price envelope provides a reasonable good indication of when the market may be near a turning point but extended trends can make the prices hug one end of the envelope over a longer period. Figure 3.6 shows a sample price envelope.

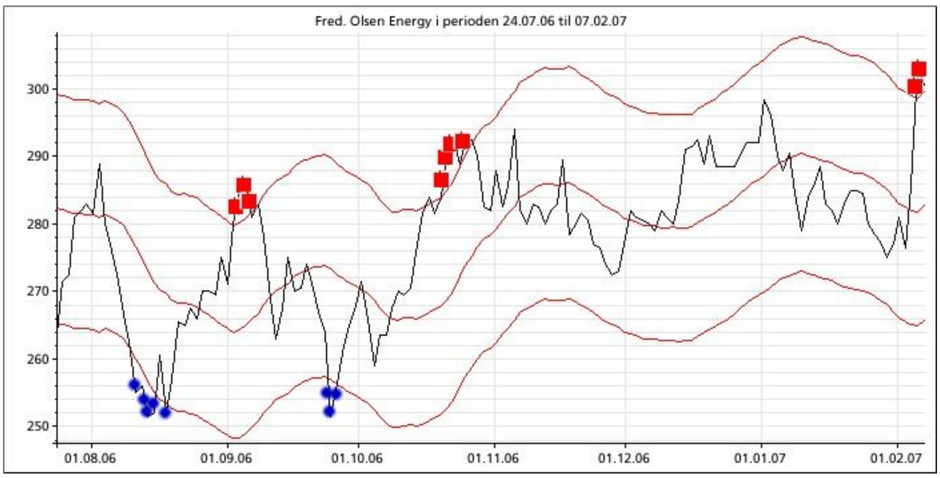

Figure 3.6: 6% price envelope based on 40 days Moving average: FOE stock

3.7.1 The Price envelope band agent

The Price envelope band agent uses the normal 40 days moving average agent described in 3.6.3 to calculate moving averages and adds 3% to create the envelope bands. If prices move above or below 5% of the upper and lower band respectively, a buy or sell signal is given according to which direction the prices are moving.

3.8 Double tops and bottoms

Double top and bottom formations are chart patterns made up of two consecutive price moves followed by reactions back to the initial starting point of the advances. The two tops or bottoms that make up the pattern need not be exactly the same, only in the same general price vicinity [4, page 96]. A double top is considered completed when prices move below the reaction low between the two tops of the formation. A double bottom is completed when price move above the reaction up between the two bottoms of the formation. Figure 3.7 shows a double bottom formation.

Figure 3.7: Double bottom formation: KVE stock 2000-2002

3.8.1 The double top and bottom agent

The double top and bottom agent checks for a double extremes formation determined by the stock trade price move the last day. If this direction is down, the agent checks wether the last trade is the confirmation of a completed double top formation. This is done using the following procedure:

- Identify if there exist two prior trade days with trade price at most 1% lower than that of day n but with prices being at least 25% higher than that of day n in between them
- Check wheter the difference in length of the time period between each of the two minimum point and the highest value between them does not exceed 30%

If these two conditions are met, a double top is identified and sell is signaled. Similarly, a double bottom formation is identified when two bottoms are formed in the same manner as the tops above and a buy signal is returned. Otherwise, a hold signal is returned.

3.9 Head and shoulders

The head and shoulder formation is one of the most well known technical chart patterns. It is formed by a three part formation, consisting of three tops where the middle high is above the high points on either side. A neckline, made up of a straight line between the relative lows on each side of the middle high, needs to be penetrated by the last price move for the formation to be completed. The reversed head and shoulders is analogous and consists of three consecutive bottoms where the middle one is lower than the two on each side. When successfully completed by penetration of the neckline, it signals that prices should move upwards. A head and shoulder pattern is illustrated by figure 3.8 where the formation was finalized in may 2000, indicating a sell signal at price NOK 138 per share. This signal was followed by a steady price decline of more than 50%, bottoming out at NOK 67,50 per share in december 2001.

3.9.1 The head and shoulders agent

The head and shoulders agent checks wether the last price move in a particular stock is the final penetration of a neckline in a head or shoulders or reversed head and shoulders formation. If the price move at day n is a down one, the agent tries to identify a completed head and shoulders pattern:

- Identify the lowest price in each of the 3 last clusters of prices that are between -3% and 3% of the trade price at day n. The 3 lowest prices in each cluster and the last trade price is used to calculate the starting and end points of the shoulders and the head in the formation
- Check that none of the clusters has prices below -3% of the price at day n. The price range between -3% and 3% constitute the neckline of the formation

Figure 3.8: Head and shoulders formation, KVE stock 1999-2000

- Check whether the prices moved above 20% of the trade price at day n between the first and second point and third point and day n. The top values between these two sets of points constiture the shoulders of the formation
- Check wheter the highest point between the second and third point is at least 40% above the trade price at day n and higher than the highest value of the other two tops. This forms the head part of the formation
- Check wether the prices moved upwards at least 10% prior to the first cluster (the beginning of the first shoulder)

If these conditions hold, the agent has identified a head and shoulders formation and returns a sell signal. Similarly, when the price move at day n is a down one a possible completion of a reversed head and shoulders pattern is tested, using high prices and minimum values between clusters of such. If found, this will make the agent return a buy signal. In all other cases, the agent returns a hold signal.

3.10 Hanging man

Candlestick charts are special stock charts formed by the lowest, highest, close and open price for each trading day. It is drawn using the relation between the open and close price and the highest and lowest values. The open and close make up the "real body" - if hollow it means that the close is higher than open price, if solid it means that the open price is higher thant the close.

The hanging man is a special candlestick pattern, suggesting that the market is at the end of its rope and that a turning point is near. The criterias to indicate the hanging man are [10, page 65]:

- The real body is at the upper end of the recent price range
- The lower part of the shadow [2] should be at least twice the height of the real body
- It should have no or a very small upper shadow

The longer the lower shadow and the smaller the real body, the more predicitive the hanging man candlestick is said to be about a near trend reversal. A Hanging Man candlestick type is illustrated by Figure 3.9.

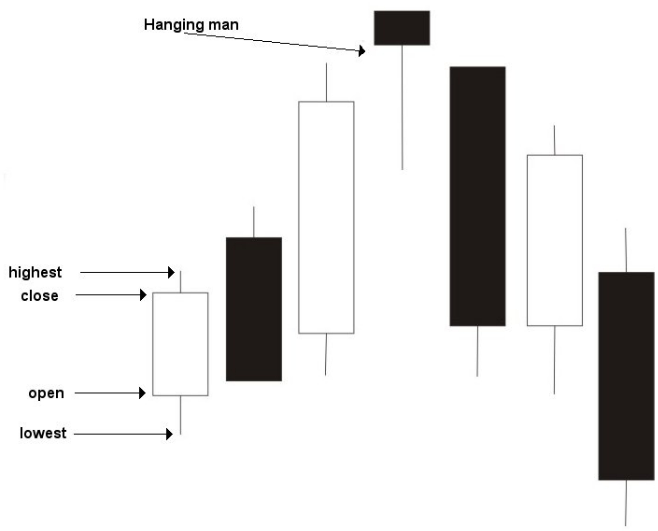

Figure 3.9: Hanging man candlestick

3.10.1 The Hanging man agent

The Hanging Man agent is defined using the following parameters:

- Minimum intraday price range: 5%
- Maximum head (real body) size: 30% of shadow
- Maximum upper shadow: 5%

[2] The thin line made up of the intraday highest and lowest trade price

The Hanging Man agent will check wether the most recent price activity matches the above given parameters. If the shadow is of length above the minimum intraday price range, the real body has a size lower than the maximum head size and the upper shadow is no longer than te maximum parameter value, the agent will output a buy (if the prices has fallen in the recent days) or sell (if the prices has risen the last days) signal, otherwise it will output a hold signal. If the trend the recent days is a sideways one the agent will output a hold signal - such a movement cannot be reversed.

3.11 Volume

Trade volume is the number of stocks that has been traded a particular day. When signifanctly higher than normal, the volume can signify increased attention given to a stock either in positive or negative manner. In addition, many trades performed indicate that a large amount of investors finds the pricing of a stock in the traded range attractive. The trade volume of a stock can therefore be used to confirm price movements.

3.11.1 The Volume agent

The Volume agent calculates the latest 10-days mean volume value. If the volume for a particular trade day is significantly above (more than 200%) the mean of the last 10-days volume, a signal based upon the latest price movement is given. If the price movement that day is an up one, the agent returns a buy signal. If the prices move downwards on high volume, a sell signal is returned.

3.12 Sell in may and go away

“Sell in may and go away“ is one of the best known proverbs in the stock market [11]. It is based on the fact that, historically, the stock market indexes have declined between may and november compared to the increase between november and april. Figure 3.10 is constructed using calculations of the monthly returns for the S&P 500 stock market index in the period between 1950 and 2006, made available by Yahoo Finance [12]. It shows the average monthly return for the index over the time period and clearly supports the phrase with an average return of 0.74% between november and april compared to only 0.16% between may and october. It should be noted, however, that these results are clearly influenced by a couple of isolated occurences: the S&P500 dropped 22% in october 1987, 17% between june and october 2001 and 23% between the same months in 2002. With these occurences removed, the “sell in may and buy in november”-alternative coincides approximately with the fully invested alternative [13]. With the famous and big stock market crashes in 1929, 1987 and 2001 in mind, all which took place in the may-october period, it seems that the “sell in may and go away”-tactic is good to avoid the big down fluctuations but superfluous during the majority of years, in which no such big movements takes place.

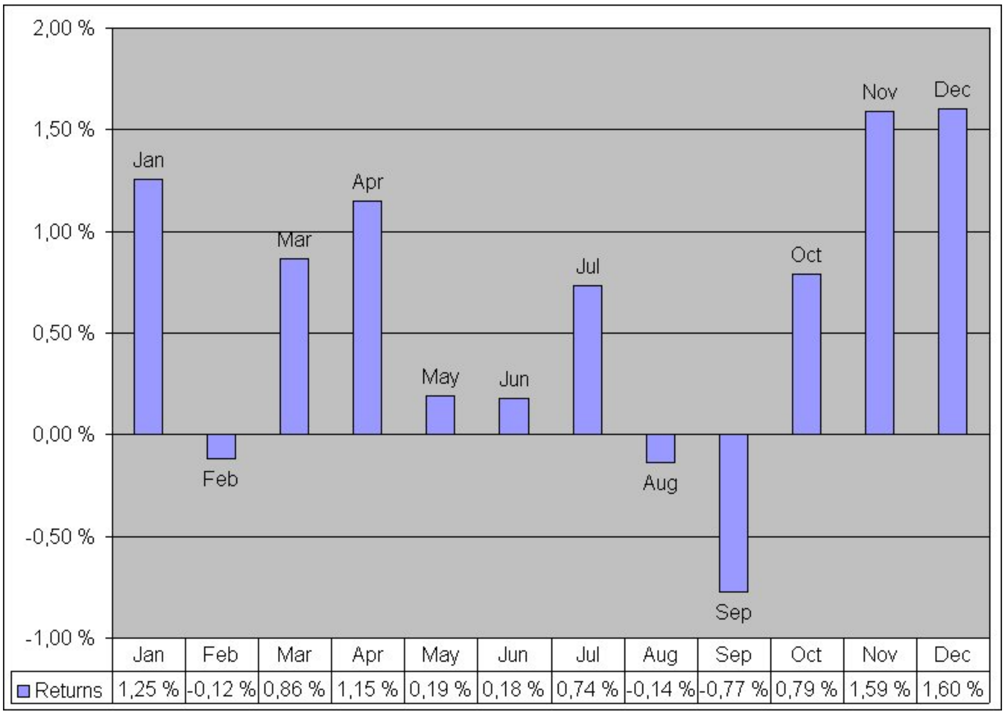

Figure 3.10: The S&P 500 index returns by months, 1950-2006

3.12.1 The Sell in may and go away agent

The sell in may and go away agent signals a negative stock opinion in may and a positive one in november in order to catch the historically profitable november-april period and avoid the historically less profitable may-october period.

3.13 The january effect

From 1925 through 2001, the average return on the S&P 500 index in january was 1.7%, whereas the average returns on small (low market capped) stocks came to 6.5% [1, page 300]. As can be derived from figure 3.10, january is historically a good month to invest in stocks, and as above mentioned, particular for smaller capped stocks. An explanation for this effect can be found in the fact that tax-motivated selling takes place in december, making investors keen to buy stocks that have taken not fundamentally based setbacks in december. This explanation is supported by the fact that prior to the introduction of the American income tax in 1913, there was no January effect and that the

Australian market have abnormally large returns in July with their tax year ending June 30 [1, page 303]. In addition, individuals often receive bonuses on top of the planned income at the end of the calendar year and companies starts their financial year congruent with the calendar year, leaving individual capital and industrial budget surplus available for stock market investing. This accounts for the fact that the january effects seems be most measureable after years in which the stock markets, as a natural effect of increased company earnings, ascended [14].

3.13.1 The January effect agent

The January effect agent will try to catch the historically based probable price advance in january, thus signaling a positive market opinion in december. In order to make the stock agent use the general market trend more accurate by trying to trade only the good years, an extra condition has been added to the agent: it will only signal buy if the year about to end is a year where the stock market main index is trading higher at the end of the year than in the beginning. In all other cases, the agent gives no buy or sell recommendation of stocks.

3.14 Neural networks

A neural network is a collection of interconnected simple processing elements called perceptrons. The output from each perceptron is calculated by multiplying each input with a trained value and adding a bias value to the equation. If the resulting value ends up above 0, the perceptron outputs 1. If not, it outputs -1. Thus, the final output from the neural network will be binary: yes or no. A neural network with n input layers, 3 hidden layers and one output layer is illustrated by Figure 3.11. Several studies have been made concluding with the fact that such neural networks can be successfully applied to time stock market entry and exit [15].

3.14.1 Training a neural net

The R project for statistical computing [16] provides a software framework for training and evaluating neural networks. The incoming weights and biases of each perceptron is in this case fitted using a neural network package implementing the BFGS-algorithm for backpropagation of weigths.

3.14.2 The best and worst patterns

William J. O'Neil suggests in his book "How to make money in stocks: a winning system in good times or bad" [17] that a trader should study the pattern of previously particular successful and unsuccesful stocks at the point before

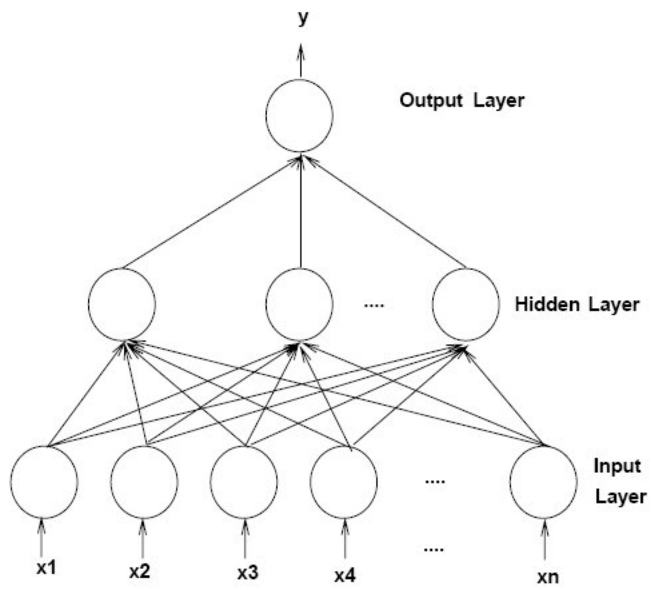

Figure 3.11: A single-hidden-layer neural network [15]

their big price move in order to pick future winners and losers. The book provides weekly trade charts for a selected group of total 93 models of the most successful investments in the United States between 1953 and 1993 and 23 of the more unfortunate ones. The models show the last year of trade for stocks like Northwest Airlines, Kresge, Prime Computer and Ford Motor, ending just before they started advancing up to 1500% in a time period of 8-40 months.

O'Neil suggests that many of the future stocks on the verge of making a big move in one direction with large probability will follow the patterns previous similarly moving stocks has made up to the point of their enormous increase or decease [17, page 159].

By training the perceptrons of a neural network using O'Neils models together with another handful examples for particularly good, bad and normal sideway stock patterns up to the point before they occured, two neural networks can be calculated: one deciding whether a stock should be bought or not and one whether it should be sold or not. Using the model depicted in figure 3.11, the *x1->xn* input layers will be the 52 last weekly trades in a particular stock, and the output y whether that pattern seems to fit that of the most successful/unsuccessful or not.

Since stock price movement vary relative to an initial price, training data needs to be normalized. This is done by adjusting the history by means of relative extremes over the training time period:

$$normalizedprice_n = \frac{price_n - lowest}{highest - lowest}$$

where *highest* is the highest price over the time period to be normalized — *lowest* is the lowest.

3.14.3 The Neural net agent

The Neural net agent uses trained neural network by adding statical perceptrons to a net evaluator. Two networks are being used: one for buy and one for sell signals. Incoming data are normalized using the equation as shown in 3.14.2. The normalized data is evaluated by applying weights from the predefined network. For the 52-3-1 network used by the agent, each of the 3 perceptrons has 52 incoming weights for each signal in addition to one bias value. The output for perceptron x is computed as follows:

$$output_x = bias_x + \sum_{n=1}^{n=weights} normalizedprice_n * weight_n$$

Dependent of wether this value ends up above or below 0, the value 1 or -1 respectively is multiplied with its corresponding weight and summarized with the other two perceptron weighted outputs and the output bias to form the neural net agent signal. If the output perceptron value is above 0, it yields a signal according to which net is being evaluated, buy or sell, or hold otherwise. If both the sell and buy networks evaluates to positive signal, the agent outputs hold on the basis of the agents ambiguity.

Chapter 4

Aggregation agents

Different technical signals can corroborate or invalidate each other. The agents described in Chapter 3 have been combined to form "high-level"-trading agents. These are described in the following subsections.

4.1 Motivation

Trading based on single trend indicators or trend reversal patterns can be hazardous and should instead always be based on different technical models that confirm each other. Only under such conditions can signals be found reliable [2, page 16]. To achieve this, machine learning has been applied in making aggregation agents able to combine the signals from the agents described in Chapter 3.

Figure 4.1 illustrates one of the ways such an aggregation agent can contribute: The Double top/bottom agent returns a buy signal at price NOK 295 per share directly before an enormous advance to a maximum value of NOK 2380 per share. The subsequent sell signal is first given when the share price has fallen all the way below NOK 40 per share, which in this case would result in trading based on the agents signal alone would give a negative yield on the investement although the buying is timed almost perfect. By listening to signals from e.g a normal RSI agent somewhere near the price peaks after the advances, a trader would have been advised to sell its shares based on the fact that the stock was extremely overbought, thus realizing good return on the initial investement insteading of sitting on the stocks all the way through the down period.

4.2 Training data

To train the different aggregation agents, data for 24 stocks listed on the New York Stock Exchange over different time periods between 1970-1998 has been used, totalling up to 68,246 trade days. To secure correctness in the price

Figure 4.1: Double top/bottom agent signals for OPC 1999-2006

data, periods in which a stock has been subject of splits og splices has been omitted. The stocks and in which time periods they have been used is listed in the Appendix A.1.

4.3 The Voting classifier

The Voting classifier is a trading agent that weights the signals given from each of the simple agents and uses this aggregation to decide how it should act. The final signal value is calculated as follows:

$$\frac{\sum_{n=1}^{n=agents} signal_n * weight_n}{\sum_{n=1}^{n=agents} |signal * weight_n|}$$

With a sell signal defined as -1, a hold signal as 0 and a buy signal as 1, agents signaling a neutral stock opinion (hold) will not be taken into account when calulating the summarized voted value. A final voted value above 0.5 results in a buy signal from the Voting classifier and one below -0.5 a sell signal.

This is exemplified by figure 4.2 which shows an example of how such a classifier could look like. To 0 weight of the trend agent indicates that its signal will never be taken into consideration. The weight of 9 of the Double Top Agent compared to the summarized absolute weight of 7 for the other agents will make it able to veto the outcome of the classifier when signaling either buy or sell. When it signals hold, its weight will not be taken into account in the final equation, making any signals from the other agents able to decide the result of the classifier. Since the RSI agent is given a weight lower than zero, a sell signal from it will be added as a positive value in the equation and a buy signal negative. This implies that this examplified Voting classifier defines transgression of the overbought threshold a buy signal.

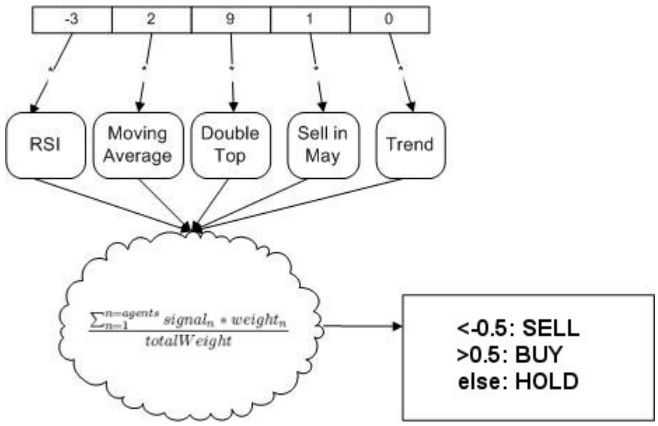

Figure 4.2: A sample voting classifer agent

4.4 JGAP

To make the voting classifier optimal by applying the best weights to each signal, it is being trained using a genetic algorithm. This is realized using JGAP, a java based open source framework for applying evolutionary principles to problem solutions [18]. For supplementary information about genetic algorithms, *An introduction to genetic algorithms* [19] by Melanie Mitchell gives an interesting introduction to the subject.

To configure the JGAP framework for a specified problem domain, one needs to specify a set of genes that make up chromosomes used as parameter in a fitness function, returning values based on how the chromosome performs over the training set. The genes can be compared to the weights in Figure 4.2 and a chromosome to the complete array of the weights in the same figure.

4.4.1 Defining chromosomes and genes

A solution for the problem is represented by a chromosome, set up by an array of genes. For this application, each gene has been defined as the weight for a specified simple agent, being an integer value ranging from -10 till 10. A negative value implies that the agents signals should be reversed, meaning that a sell signal should be regarded as a buy signal and vice versa. Gene values near 0 denotes that the agent's signals should not be given much weight; the importance of an agent's signals increases as it approaches the extreme values -10 or 10.

4.4.2 Choosing the fitness function

The fitness function is used to calculate how good a chromosome (solution) is for the problem domain. It is being used to train the best classifier, where higher fitness values for a chromosome denotes a better agent. For all fitness functions, an "always-in-the-market"-approach is used, meaning that two consecutive buy or sell signals would result in the agent simply holding its first position, and a buy signal followed by a sell signal that the agent would realize its position and take a new short position. The subsequent buy signal would result in the agent buying back its short position and take a new long (buying) position immediately.

The Buy and sell fitness function

The Buy and Sell fitness function evaluates a chromosome based on how much average yield the summarized total of all buy-sell and sell-buy transactions returns for all stocks in the training set. For each normal buy-sell transaction, the yield is meaured by the formula:

$$yield_n = \frac{sellPrice_n - buyPrice_n}{buyPrice_n}$$

For a short selling transaction, it is calculated as follows:

$$yield_n = \frac{sellPrice - buyPrice_n}{sellPrice_n}$$

The overall fitness value is calculated based on the average return for all trades in each stock, both normal and short selling:

$$fitness = \frac{\sum_{n=1}^{n=trades} yield_n}{stocksInTrainingSet}$$

The Correct Signals fitness function

The Correct Signals fitness function measures how many percent of the signals given from the agent are the correct ones. This means that if a buy signal is followed by a price advance, it is given status correct. Similarly, a sell signal followed by price decline and a hold signal followed by a price movement within a range of 1% of the initial price is said to be correct. The resulting fitness value over the training set is calculated as follows:

$$\frac{\sum correctSignals}{allSignals}$$

It is not sufficient, however, to only look at the next day price move to determine the correctness of a signal. Figure 4.3 illustrates this, in which the common gap agents first buy signal is followed by a small price advance and a subsequent large decline. The same applies for the first sell signal that is given before a milder up price correction. It is obvious that these signals should be awarded as two negative ones and not positive as the next-day price direction approach would lead to. The correctness value for a signal at a certain day has therefore been chosen to be based upon the extreme top or bottom value found afterwords, until this value has corrected back at least 15% from its peak or one year has passed after the initial signal. With this prerequisite, the correctness function avoids awarding less fortunate longer term signals based on short-term minor correctness. Figure 4.3 is illustrated with a sample drawn threshold range for both of the discussed signals, indicating the range of 30% from the signal which needs to be penetrated from either the up- or downside to result in a sell or buy signal being classified correct.

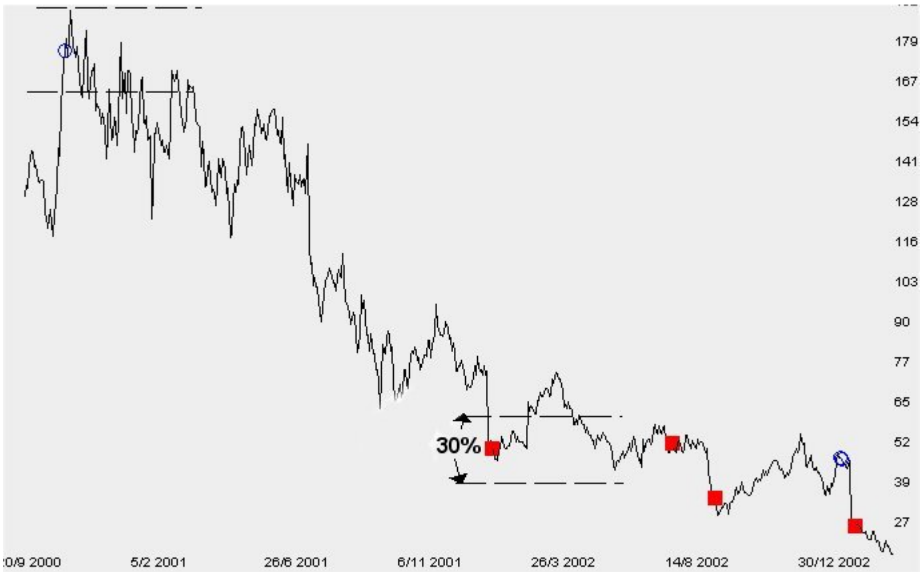

Figure 4.3: Common Gap signals for TOM stock, 2000-2002

The Extreme price fitness function

The extreme price fitness function calculates its fitness using the same approach as for calculating correct signals as outlined in the section above, but adds bias to each signal by using the next relative extreme value. This means that signals that resulted in extreme earnings are more weighted in the total equation than signals that was followed by only minimal price movement. The weight of each buy signal is calculated as follows:

$$weight_n = \frac{nextExtreme_n - price_n}{price_n}$$

A sell signal is given weight as follows:

$$weight_n = \frac{nextExtreme - price_n}{nextExtreme_n}$$

The resulting fitness is the average of all signal weights:

$$\frac{\sum_{n=1}^{n=transactions} weight_n}{transactions}$$

4.4.3 Training the classifier

With chromosomes, their genes with allowed values and a fitness function well defined, JGAP can be started in order to try and find an optimal solution. This is done by specifying a population size[1] and the total number of evolutions. Each gene is initialized by applying random values within their allowed range (between -10 and 10). For each evolution, JGAP will use the defined fitness function to calculate the fitness of each chromosome. The best chromosomes (those with the highest fitness value) will have some of their genes mixed from one another whilst those with lower fitness values will see their genes die out before the next evolution, following the "survival of the fittest"-principle. This, combined with mutations - that a random, small amount of genes accidentaly changes their value in each evolution - should imply that each evolution of the population produces better chromosomes. For each evolution, the best chromosome is saved and this will, after the total number of allowed evolutions is reached, set the weights of the voting classifier agent.

4.4.4 The Voting Classifier agents

Three different Voting classifier agents has been trained based on the fitness functions described above using the standard configuration in JGAP with the option to preserve the fittest chromosome for each evolution set and using a population size of 3000 chromosomes with a total of 10000 evolutions. The goal is to get agents with competence to achieve optimized results for different environments:

- The Buy and Sell voting classifier: Be able to make most money out of constant trading, optimizing the average yield for all transactions
- The Correct signals voting classifier: Be able to give the best signals at all times, recommending sell near at, at and during recessions and buy signals near and just after prices are bottoming out and during an uptrend
- The Extreme price voting classifier: Be able to spot extreme price moves in advance and ensure that one does take the winnings such can produce possibly on the account of the correctness for signals not returning as much of an investement

A confirmation rule for the agents output have been taken to avoid whipsawing — constant shifting positions when no direction is confirmed — as for the Moving average agents: a signal needs to be repeated for at least two consecutive

[1]The number of different chromosomes allowed in each population evolution

days to be put into effect. If the resulting computed signal of the agents shift on a day-to-day basis, a hold signal is return on basis of the instability shown.

4.5 Decision Trees

Decision trees are built up from a set of connected nodes which holds options, where traversing them eventually leads to a leaf node describing which action is to be taken. Its application in the construction of an aggregation is outlined by figure 4.4. In this example, a buy signal from the RSI agent will result in the decision tree indicating that buying should take place. Otherwise, other agents signals are evaluated before a final decision is given at the leaf of each node.

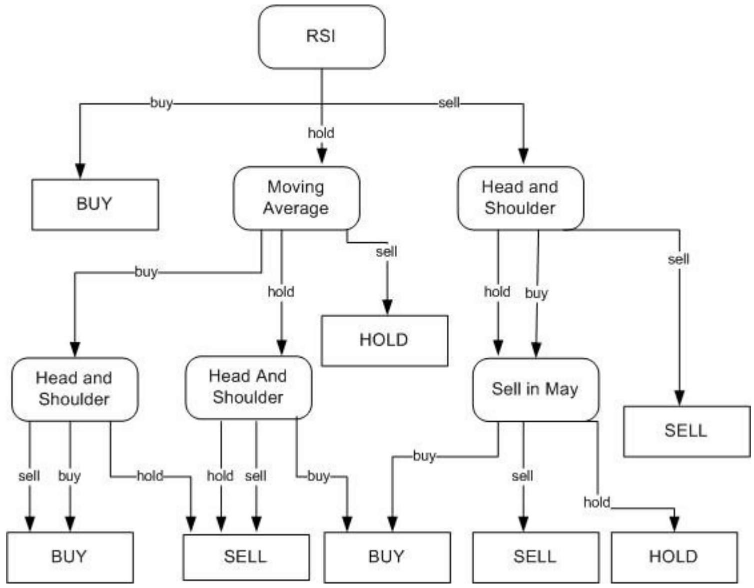

Figure 4.4: A sample decision tree, some nodes interlaced for readability

Compared to the voting classifier, a decision tree has the advantage that it can express connections between only parts of the agents if desired, e.g signaling sell if and only if a RSI and a Head and Shoulders Agent does so. This benefit is gained on expense of the awarness of the dynamical relation between all agents.

4.5.1 Weka

The Weka Machine Learning Project [20] provides an open source machine learning software package for data mining tasks. It uses the Attribute-Relation File Format (ARFF) which defines the problem domain and the data related. Such a file for this specified problem can for instance contain:

```
@RELATION agents
@ATTRIBUTE RSI {sell,hold,buy}
@ATTRIBUTE SellInMay {sell,hold,buy}
@ATTRIBUTE MovingAverage {sell,hold,buy}
@ATTRIBUTE HeadAndShoulders {sell,hold,buy}
@ATTRIBUTE class {SELL,HOLD,BUY}

@DATA
hold,hold,hold,buy,BUY
hold,hold,sell,hold,SELL
hold,hold,hold,hold,HOLD
hold,hold,hold,buy,BUY
hold,hold,hold,buy,HOLD
sell,hold,sell,sell,SELL
```

In this case, the first 4 attributes constitutes the different agents and their possible signals. The last attribute, *class*, is the classification definition, in other words which possible values the final result can take. The data section contains the agents signals for each trade day together with a classification value based on which signal is desired at the particular day. By applying this training set to the machine learning algorithms available in Weka, a decision tree for the classification value can be constructed, where the one illustrated by figure 4.4 could be a solution for this example. In order to build an aggegation agent based on a decision tree, a ARFF file for all 20 simple agents and the classification of sell, buy or hold as attributes has been constructed, together with the agents signals over all days in the training set. The desired classification value has been set based on the particular trade days next extreme value as described in section 4.4.2.

4.5.2 The ID3 decision tree agent

The ID3 decision tree agent uses a decision tree computed using the ID3 algorithm package in Weka to classify wether a set of agent signals should result in a sell, hold or buys signal. The tree has a total of 8766 nodes and correctly classified 61.8% of the instances in the training set.

4.5.3 The J48 decision tree agent

The J48 decision tree agent uses a decision tree computed using the J48 algorithm package in Weka configured with pruning to classify its signals. The tree has a total size of 1075 nodes, 717 of them being leaf nodes. It correctly classified 61.3% of the instances in the training set after completed training.

4.6 Neural network

A neural network, as described in section 3.14, can be used to classify combinations of single agent signals into one unified signal. Using the same training data as described in Section 4.5, a 20-3-3 neural net for signals can be constructed, where the 20 incoming first weights are the different agents signals and the 3 output perceptrons indicates buy, hold and sell respectively.

4.6.1 The Neural net aggregation agent

The neural net aggregation evaluates the trained neural network by applying its weights to the signals from each of the 20 seperate single agent signals and returns the value indicated by the output of the network.

Chapter 5

Results

The implemented agents' profitability has been benchmarked using a broad range of stock trade histories from the time period 1999-2006. This time period covers both major up trending cycles as well as a longer period of decline. The resulting performance measures gives hint of the different agents strengths and weaknesses, and are a good foundation for evlauation their usefulness and reliabilty both in general terms and for special market situations.

5.1 The voting classifiers

Table 5.1 displays the trained agent weighting for each of the three voting classifiers discussed in 4.4.4. The compositions gives a useful indiation of the different simple agents usability.

5.1.1 Overall observations

All agents have assigned little weight to the Island Reversal Agent, the three Moving Average agents, the Head and Shoulders agent and the Neural Network Agent. This indicates that signals from these agents has not been considered useful and implies that the patterns to be recognized or the lack of accuracy in doing so, and the indicators computed, have failed to live up to their theoretical goals. This fact is further discussed when evaluating the test data later in the chapter.

5.1.2 The Buy and sell Agent

The weighting of the Buy and sell agent indicates the following when exploring its composition from a maximum average yield per transaction-perspective:

Agent Name	Buy&Sell	Optimal	Extreme
Double Top/Bottom	0	1	2
Island Reversal	3	0	2
Moving Average 10 days	2	0	-1
Moving Average 40 days	2	1	0
Moving Average 200 days	-1	0	0
Crossover Moving Average 10+100 days	-5	10	-1
Price Envelope Band	-2	-1	-10
Common Gap	-5	0	8
Breakaway Gap	-2	0	9
RSI 9 days	9	-1	-10
RSI 14 days	10	-1	-10
RSI 21 days	0	0	-10
Head and Shoulders	-1	0	0
Sell in May	-10	0	0
January Effect	1	10	0
Trend	-4	0	-1
Volume	7	1	1
Hanging Man	-10	-1	-10
Reversal Day	-9	-10	0
Neural Network	-2	0	0

Table 5.1: Weighting in the different voting classifier agents

- The “sell in may and go away” phrase seem to lack meaning when used on real-world data as the agent want to interpret it with maximum weighting as “buy in may and sell in november”
- The hanging man and reversal day patterns are weighted to notify of the opposite of what they theoretically indicate
- Price movements occuring on extreme volumes compared to normal seem positive for catching profits out of a yield perspective
- Short and normal term RSI indicators gives hints of good trade entry and exit points
- A short term moving average agent crossing the path of a longer term one is a mildly negative for this investement perspective

5.1.3 The Optimal signals Agent

The optimal signals Agent relies heavily on three agents: The Crossover Moving Average 10+100 days agent, the January effect agent and the opposite opinion than that of the Reversal day agent, all of which are given maximum allowed weighting. The summarized weighting of the other agents taken into account, all with minimum allowed weighting apart from 0, can serve to veto a signal from one of the three maximum weighted agents, or decide the outcome when a signal from them are absent. In addition to the fact that the month January after an upturn year seem to give good buying possibilites, this indicates that producing good signals in the area around price peaks instead of demanding them being exactly at them can principally be done quite straightforward, listening mainly to crosses between short- and long term moving averages and interpret reversal day formations as continuation patterns. When such signals are not available, doing the opposite of short- and normal term RSI and price envelopes when all occuring at the same time seem to give good results.

5.1.4 The Extreme signals Agent

The composition of the extreme signals Agent shows an almost perfect coinciding with the theory:

- Both Gap agents are weighted heavily: An extreme price move will normally take the form of a gap several times during its lifetime
- All RSI agents have been given maximum negative weighting: An extreme up move will move into and stay in the overbought range as it is defined as a move where buying holds a concurrent high rate. This analogous for a down move
- The Price envelope band is weighted maximum negative: The Price envelopes will stay constantly penetrated when prices surge as it defines overbought- and oversold ranges similar to the RSI indicators

The maximum negative weighting of the Hanging man agent is, however, quite contrary to what the theory implies. It is supposed to signal trend reversals but, counting opposite from what it is designed for, signals trend continuation instead.

Two agents shows unexpected low weighting compared to what one could expect:

- The Island reversal agent is only weighted slightly negative: Island reversals are normally formed in the beginning and end of a trend major trend and should theroetically be an obvious tool in detecting extreme price moves
- The Neural network agent is trained to detect extreme price moves. The fact that its signals are being totally ignored by the Extreme signals voting classifier indicates it does not work according to its goals. This can be a result of too little training data, poor training rules or the simple fact that the assumptions it has been trained upon are incorrect or insufficent

5.2 Testing

In order to test the agents, stock market data from a time frame containg periods from sideways trading and major up- and down-movements has been chosen and special occurences within the time frame has in addtition been extracted to benchmark agent performance during special market trends.

5.2.1 The well-chosen example

It is possible to find favourable illustrations for all trade systems but basing probable future performance on isolated well-chosen examples from the past does not give a correct evaluation of overall reliability. To determine wheter a system has value it should be tested over an extended time period for a broad range of sectors [4, page 255].

5.2.2 The test data

To avoid falling into the well-chosen example trap, all available stock history for the Oslo Børs stock exchange over the time period of 7 years between 01.01.1999 and 31.12.2006 has been chosen as test data. This ensures that the test data is not influenced by the training data whatsoever (as it is taken from a different stock exchange and a different time frame) and has good sector spread by including all the diversified stocks available at the exchange. The test data has been derived from each available ticker listed at the complete stock exchange list from the online stock broker Netfonds [21] using the option to download historical quotes. The historical quotes has been adjusted for splits and splices by the stock broker, securing compability with the latest list prices. To examine performance during special market movements, the training data has in addition been broken into the trades in hindsight resulting in particular high increases, decreases and waving prices.

Some stocks needed to excluded from the test set for one of the following reasons:

- Illiquidity: A total of 65 stock histories had several sets of longer periods of consecutive trade days in which no trading occured. This results in the stock's trade data taking forms not suitable for technical analysis and mathematical calculations based upon them. In addition, such illiquidity provides for artificial price moves as smaller stock blocks are sold to or bought for the best available price of few avaible for investors with haste getting out of or into the market. Both of these issues are illustrated by figure 5.1. The ticker symbols of the stocks excluded on basis of illiquidity is listed in the Appendix A.3.1.

- Too late listing: A total of 31 stock got listet for trade at the stock exchange so few days before the end of 2006 that it is not sufficent data available to perform technical analysis on them. These stocks has been excluded from the test set and are listed in the Appendix A.3.2.

A listing of the 132 stocks that has been included in the test set is given in the Appendix A.2.

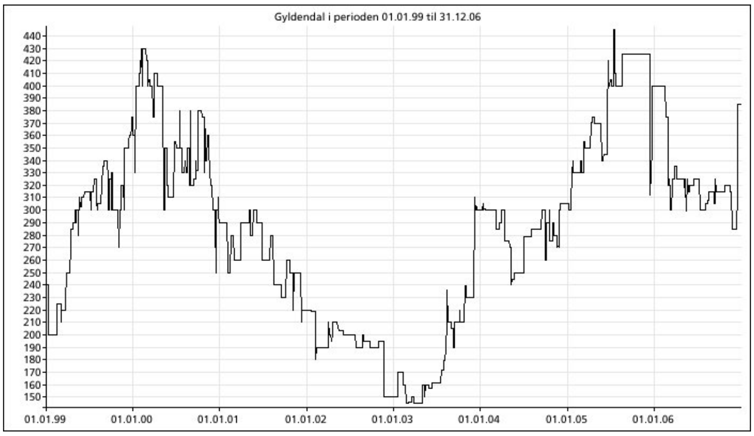

Figure 5.1: The illiquid GYL stock, 1999-2006

The overall data

The overall data contains several different market movements:

- The years 1999-2000 makes the last part of a major uptrend cycle and saw abnormal stock market value increase for a range of speculative stocks.
- The end of the powerfull bull marked in 2000 was followed by a market recession which lasted until march 2003
- The years 2003-2006 has showed an extremely steady and long up move only disrupted by minor corrections on the way

There are no unified index indicating the overall performance of the stocks on the Oslo Børs stock exchange over this entire time period because the exchange switched the definition of its main index in mid 2001 and stopped quoting the old one, named Totalindexen, at the end of 2001 [22]. Illustration of the price movements on the exchange over the complete time period has therefore been divided into two different charts: Figure 5.2, designed using the historical data available at the stock exchange's own home page [23], shows the existing main index for Oslo Børs during the time period 01.01.1999-31.12.2001, and Figure 5.3, collected from the oline stock broker Netfonds [24], shows the current main index, The Oslo Børs Benchmark Index, over the time period 01.01.2002-31.12.2006.

The winning stocks

A total of 49 stocks that has showed particular price increases has been extracted from the test set to examine the agents performance during steep up moves. Periods of longer sideways or minor downwards movement prior to or after the major increases has been erased from this extracted test set, but all movements just before and during the moves has been included. A list of the winning stocks included is given in the Appendix A.2.1.

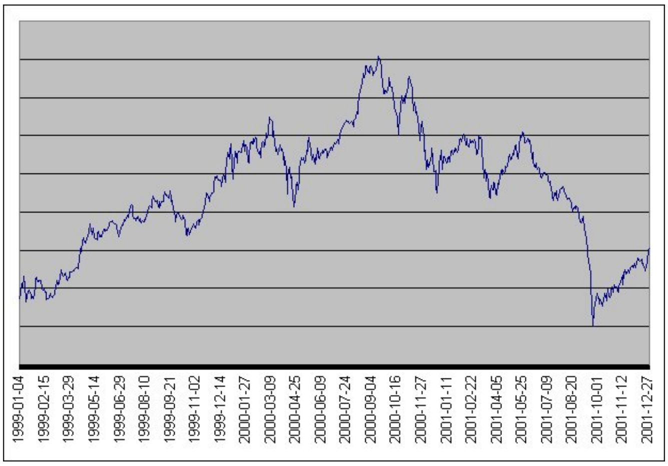

Figure 5.2: Oslo Børs main index, Totalindeksen, 1999-2001

Figure 5.3: Oslo Børs Benchmark index, 2002-2006

The losing stocks

As for the winning stocks, 23 stocks has been extracted from the test set, all of which showed particular negative developement. Similarly to the winning stocks, periods of longer sideways or minor updards movement prior to or after the major decreases has been erased from the extracted test set, but the complete price history just before and during the down moves has been included. A list of the extracted losing stocks is given in the Appendix A.2.2.

The volatile stocks

Volatile stocks has the biggest potential for money making compared to stocks moving in one direction as they are recognized by waves of both up- and down moves during their lifetime. Investors able to buy shares near local bottoms in such formations and sell near or at local tops kan repeatedly catch good profits. Trading volatile stocks can, however, be extremely disadvantageous when price wave catching are not timed correctly. A list of the volatile stocks that has been extracted from the test set is given in the Appendix A.2.3.

5.3 Measuring performace

Three different approaches has been chosen in order to measure the different agents performance. These has been named yield, profitable trades and all-in yield. For all methods, an always-in-the-market strategy has been taken, meaning that when sell is signaled by an agent, it would sell the stocks it owns and additionally short sell stocks for the investing amount, unless the agent already has short sold stocks. Similarly, a buy signal will result in buying back of short sold shares and buy a number of stocks based upon the amount of money available to invest. The reason for using this approach is the fact that it makes the performance easily measurable when comparing to the buy and hold strategy, and the interpretation that a positive market view stays positive until the oppsite is indicated and vice versa.

5.3.1 Yield

An agent's performance is measured using its yield by assuming that it constantly would invest the same amount of money in the market for each trade, taking long positions when it signals buy and short ones when signalling sell. This procedure was outlined in section 4.4.2 and also used to train one of the voting classifiers: Profits and lossess are not reimbursed into the market, instead put aside and added to the final yield. For each normal buy-sell trade, the profit is meaured by the formula:

$$yield_n = \frac{sellPrice_n - buyPrice_n}{buyPrice_n}$$

For a short selling transaction, the profit is calculated as follows:

$$yield_n = \frac{sellPrice_n - buyPrice_n}{sellPrice_n}$$

The overall yield is given as the return on all trades, both normal and short selling and winning and losing, in the stock:

$$yield = \sum_{n=1}^{n=trades} yield_n$$

This approach means that all signals in the history is measured equally as the agent always invests the same amount of money.

5.3.2 Profitable trades

Measuring performance based on profitable trades means measuring how many times an agent's trade is correct — in other words, how many times a pair of buy-sell or sell-buy signals resulted in a positive return of the investment. This approach has been taken instead of counting how many times an agents signal turned out to be correct one (how many times buy signals was followed by an increase and sell signals buy a decline) as the latter is no good indication of the overall performance.

5.3.3 All-in yield

The all-in yield uses the same approach as the yield method with one exception: the agent always invests all available money in each trade. This means that the agent would have more money to invest after a succesfull trade than after an unsuccesful one. The all-in yield is based upon how much money an agent would have after finishing all trades.

$$money_0 = 1.0$$

$$\sum_{n=1}^{n=trades} money_n = money_{n-1} * (1 + yield_n)$$

The yield for each transaction is calculated as described for normal yield in section 5.3.1 above. The final all-in yield is the average amount of money earned for all stocks used to trade:

$$allinyield = \frac{(\frac{money_{trades}}{stocks}) - money_0}{money_0}$$

If a short selling trade ends up with the share more than doubling is initial trade prize upon realization, the agent is given status bankrupt, cut off from further trades and resulting yield is negative based upon how much dept the trade has produced.

This approach has the best basis for comparison to the buy and hold strategy as they both have the premise of equal trade starting- and end points and invested amount.

5.4 Overall result 1999-2006

Table 5.2 shows the different single agents performance over the whole time period 1999-2006 for comparison with the "buy and hold" strategy. The table show that trading based on the Voting classifiers, island reversals, moving averages, volume and gaps were amonst the most profitbale approaches in addition to the buy and hold strategy during these in retrospective upward years and that RSI, double top, hanging man and reversal day were less fortunate trading basises.

Agent Name	Yield	Profitable trades	All-in yield
Buy and hold	238.4%	78.8% (of 132)	238.4%
Double Top/Bottom	-126.6%	46.4% (of 192)	73.2%
Island Reversal	189.1%	42.9% (of 2576)	119.8%
Moving Average 14 (5)	151.9%	42.9% (of 2424)	182.5%
Moving Average 40 (10)	201.4%	42.9% (of 933)	276.2%
Moving Average 200 (10)	64.0%	47.0% (of 330)	37.4%
Crossover MA 10+100	203.3%	48.8% (of 10544)	275.7%
Price Envelope Band	-149.7%	65.6% (of 4219)	-208.7%
Breakaway Gap	147.5%	42.9% (of 1141)	299.7%
Runaway Gap	114.4%	39.8% (of 966)	82.1%
RSI 9 days	-120.4%	57.8% (of 4296)	-125.1%
RSI 14 days	-118.1%	56.7% (of 2692)	-111.6%
RSI 21 days	-163.1%	54.8% (of 1395)	-148.7%
Head Shoulders	10.9%	48.1% (of 285)	40.1%
Sell in may	82.2%	58.3% (of 1189)	48.9%
Trend	42.4%	40.6% (of 244)	6.9%
Volume	122.2%	41.7% (of 3610)	146.9%
Hanging man	-75.9%	53.7% (of 3050)	-50.9%
Reversal day	-53.9%	42.1% (of 1112)	-48.4%
Neural Net	-0.1%	46.4% (of 815)	17.1%
ID3 Tree	32.7%	47.8% (of 42386)	283.8%
J48 Tree	52.2%	46.3% (of 28118)	155.3%
Voting Classifier Buy & sell	-68.8%	42.2% (of 12877)	-85.1%
Voting Classifier Optimal	222.4%	46.8% (of 2692)	606.9%
Voting Classifier Extreme	102.9%	24.7% (of 13280)	164.8%
Neural Superagent	-6.2%	39.6% (of 1162)	-36.5%

Table 5.2: Overall agent performance 1999-2006

5.4.1 The RSI and Price envelope band agents

All three RSI agents and the Price envelope band agent distinguish themselves by producing the largest portions of profitable trades of the agents but at the same time yielding the worst results for continous trading. This is probably the result of a few steep advances continously adding extremely bad yield to the overall result. This fact is further discussed when examining the winning and volatile subsets of the test set in Section 5.5.2 and 5.7.3 respectively.

5.4.2 The Volume agent

In spite of a poor amount of profitable trades, the volume agent results in quite good yield and all-in yield values. This fact suggests that high volume combined with movement in a direction is a good indication of coming steep increases and decreases.

5.4.3 Trend agent

Only yielding slightly positive in a stock market that in retrospective has showed a clear trending direction for several subsequent years is an extremely strong hint that the trend agent has not been designed properly to meet its purpose. Section 5.5.4 discusses this fact further when comparing it to a major up trend, but it may seem that the trend determining threshold values has been set false to catch an overall trend.

5.4.4 Double top/bottom agent

The Double top/bottom agent only performed 192 transactions over the complete trade history for the 132 stocks in the test set. It is not plausible to believe that the charts contained in average only slightly more than one double top formation and this fact implies that the parameters of the Double top/bottom agent are set to strict or that the agent is not good enough in catching patterns.

5.4.5 The Head and shoulders agent

In contrast to the Double top/bottom agent, the Head and shoulders agent managed to perform a sufficent amount of trades, but it doesn't show impressing results with an average yield of 10.9% per stock. The fact that the trained voting classifiers in practice all neglect its signals combined with the proven poor results indicates one of the following or a combination of both:

1. The Head and shoulders formation is not as potent as theory suggests and previous research has indicated

2. The agent detects to many or to few of the present patterns or gets its results corrupted by a few bad seeds

The first reason can be substantiated by the formations renown: professional traders spotting e.g a reversed head and shoulder pattern in its later stages of forming could very well artifically finish the formation by buying larger amounts of the stock in order to finalize the pattern, and take the contrary opinion once the crowd of formation traders rush to trade in the stock when the pattern is finished. If successfully executed, it would give the traders a good profit when selling off their recently bought stocks, and the massive sell-off will result in the stock not moving in the direction the formation indicated. Such formation speculation should only count for parts of the bad results as there are still a clear overweight of technical analysists relying on the head and shoulders formation, a clear indication that the agent still has much improvement potential.

5.4.6 The decision trees

The ID3 and J48 decision tree agents both yield quite good but only slightly above average all-in yield. The fact that the number of transactions they produce is extremely higher than the average of all other agents will make trading based on these agents unfavourable, especially because of the large amount of transaction costs involved.

5.4.7 The Buy & sell voting classifier

The Buy & sell voting classifier gives quite bad results with negative values for both yield and all-in yield and an average of only 42.2% profitable trades. The performance value based on the agents yield is calculated using the same approach used to train the agent and the poor test results of this value indicates little correlation between the train and test data. One of the main reasons for this is probably the training set stocks being big capped ones and a portion of those in the test set more specualtive, smaller ones. The rapid up and down price movement of e.g many of the information technology stocks during the end of the dotcom period (1999-2000) has no counterpart in the train data, and such movements has therefore not been factored in upon training. The strong reliance on the RSI indicators which is hazardous under such extreme movements can therefore partually account for the poor results. The poor test results proofs that the Buy & sell voting classifier is not suitable for general stock market trading.

5.4.8 The Optimal signal voting classifier

An all-in yield of more than double that of the next best on the list and over 60% better than that of the in hindsight very profitable buy and hold strategy, combined with relatively few transactions performed and the best average yield of all agents suggests that the Optimal signal voting classifier agent is a successful one, being able to take advante of both good and bad times.

The agent compared to its components

The Crossover moving average 10+100 days agent, the January effect agent and the opposite signal of that from the Reversal day agent are the main components in the Optimal signal voting classifier agent, with their maximum allowed weighting. The total of 10544 trades for the Crossover moving average agent compared to that of only 2692 for the voting classifier shows that there are many contrary signals being given, ruling out frequent shifting trend opinions. In addition, in times when only one or none of the major weighted agents gives signals, it seems that the summarized value of the 6 other agents with their minor weighting are enough to veto signals when they are contrary, or decide a good direction themselves. The reduction in total trades is in any case of a good one as it strongly decreases the transactions costs and in addition makes the agent more profitable.

The agent compared to buying and holding

Figure 5.4 shows the all-inn yield for both the buy and hold strategy and the Optimal signal voting classifier for the complete test period. It gives a good indication of the agents strength when comparing to buying and holding: it seems to be able to follow the up movements in the market but avoids losing money during the longer recession between mid 2000 and march 2003. The big fluctuations in the price corrections in the beginning of 2004 and autumn 2005 when comparing to those of buying and holding can be seen in comparison with the more investement capital available. It seems that the agent overall follows the market during trends and corrections but manages to keep a negative opinion during longer down periods.

Yield spread

Figure 5.5 is a histogram which accounts for how much each transaction in the trade data produced by the Optimal signal voting classifier has yielded over the complete time period. The most interesting (and pleasant) observation is that the sell signals are all mainly clustered in the lower part of the scale with only 9.3% of them having more than 25% negative yield whilst the positive yielding transactions are spread, seeing more than 29% of them yielding above 25% and about 6% of them returning more than double of the invested money. The poor total of more than half of the transactions beging positive ones is in this case overshadowed by an agent able to minimize the number bigger losing transactions while maximizing the number of bigger winning ones.

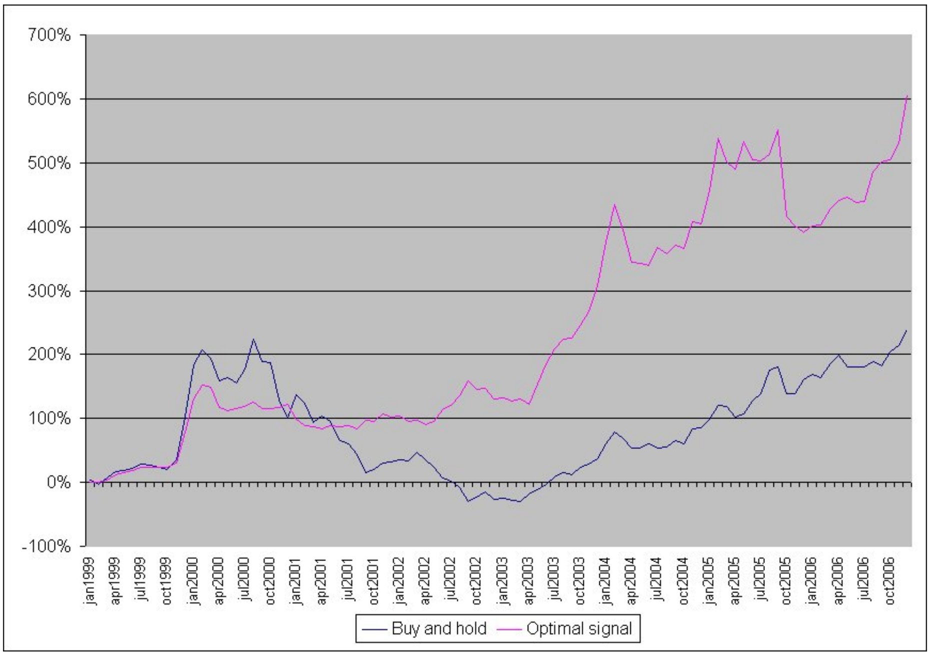

Figure 5.4: All-in yield for buying and holding compared to the voting classifier

5.4.9 The Extreme signals voting classifier

A horrible amount of winning trades, only 24.7%, and still performance in the upper range of the scale; no doubt that the Extreme signals voting classifier has managed to catch what it has been trained for: the big moves. But still, the high number of trades with its appurtenant transaction costs and total yield half that of a simple buy and hold strategy, makes the Extreme signals voting classifier unsuitable for general stock market trading. It can probably be used to indicate that big moves may be on the verge or in its initial phase, but apart from that does not give much value.

5.5 The winning stocks

Table 5.3 shows the different agents performance when trading in stocks that in hindsight has proven to be successful investment objects with the correct timing. The buy and hold-strategy is obviously a good one and in this case it clearly produces the best results with its 476.2% yield. Other agents distinguishing themselves positively are the Moving Average agents, the GAP agents, Island Reversal and the voting classifiers. The Double Top/Bottom, RSI and Price envelope band agents are placed on the losing end of the scale.

Agent Name	Yield	Profitable trades	All-in yield
Buy and hold	493.5%	100.0% (of 49)	493.5%
Double Top/Bottom	-571.9%	51.7% (of 29)	136.7%
Island Reversal	248.5%	42.1% (of 968)	236.5%
Moving Average 14 (5)	146.4%	41.1% (of 856)	200.7%
Moving Average 40 (10)	293.0%	47.5% (of 314)	296.0%
Moving Average 200 (10)	196.1%	50.4% (of 115)	125.3%
Crossover MA 10+100	202.8%	48.5% (of 3104)	219.7%
Price Envelope Band	-148.1%	63.9% (of 1328)	-261.0%
Breakaway Gap	217.2%	45.1% (of 412)	235.5%
Runaway Gap	137.7%	40.0% (of 345)	110.5%
RSI 9 days	-88.4%	58.0% (of 1599)	-74.4%
RSI 14 days	-118.2%	55.7% (of 955)	-101.0%
RSI 21 days	-174.5%	55.1% (of 490)	-141.5%
Head Shoulders	-117.1%	36.4% (of 11)	59.4%
Sell in may	56.3%	59.9% (of 439)	-7.1%
Trend	177.7%	54.4% (of 90)	86.1%
Volume	100.0%	41.7% (of 1248)	229.6%
Hanging man	-78.2%	52.8% (of 1064)	-61.3%
Reversal day	-87.6%	40.2% (of 251)	-95.2%
Neural Net	-39.1%	44.9% (of 214)	2.9%
ID3 Tree	26.8%	48.8% (of 16589)	28.9%
J48 Tree	56.8%	47.2% (of 10191)	59.3%
Voting Classifier Buy & sell	-72.2%	41.0% (of 4683)	-72.5%
Voting Classifier Optimal	193.7%	46.3% (of 903)	258.6%
Voting Classifier Extreme	64.6%	24.6% (of 5146)	115.9%
Neural Superagent	-25.7%	39.8% (of 472)	-61.4%

Table 5.3: Agent performance for winning stocks

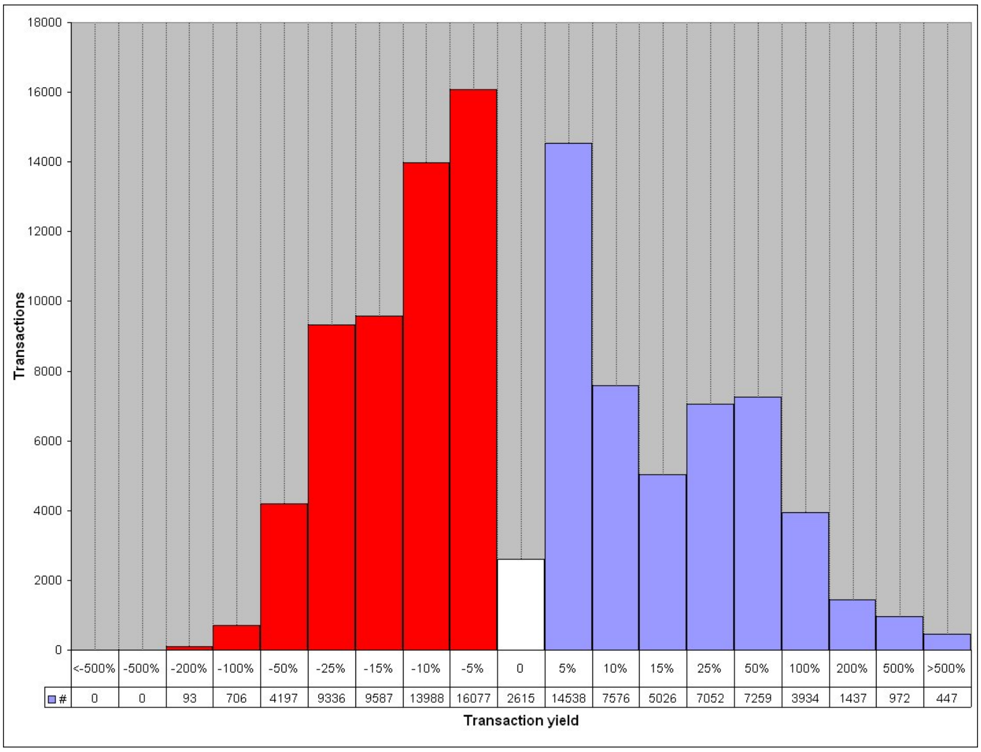

Figure 5.5: Transaction yield histogram for voting classifier

5.5.1 The Double top/bottom and Head and shoulders agents

-571.9% yield places the Double top/bottom agent in extremely bad light compared to the other agents. This result should not be overemphazied as it is based on a total of only 29 trades. False sell signals for a couple of stocks just before extremely steep advances is probably the main reason for these bad overall results. The same applies for the Head and shoulders agent where its 11 trades are not sufficient data to conclude with signs of bad performance.

5.5.2 RSI and Price envelope band agents

The Price envelope band agent and all three RSI agents yield the worst negative results, in spite of the fact that they are amongst the few agents that actually has more than 50% profitable trades. This is a natural consequence of the fact that the agents assumes that the stock prizes moves within a gradual evolving price direction, with only minor deflections from established price intervals over longer time periods. A steep price increase will result in that the agents will conclude with the stock being overbought, thus resulting in a sell signal. The overbought range will consequently hold its line during such an extreme price advance, making trading based on these osciallators very unfortunate in such situations. This is exemplified by figure 5.6 which shows the performance of the RSI 14 days agent for the whole period of 1999-2006. Although the agent performs good during the slightly declining price developement in the end of the period, with buy signals on relative bottoms and sell signals on relative tops, it almost continous signals sell during the extreme price advance between 2002 and late 2004, only interrupted by buy signals at minor relative lows over the period.

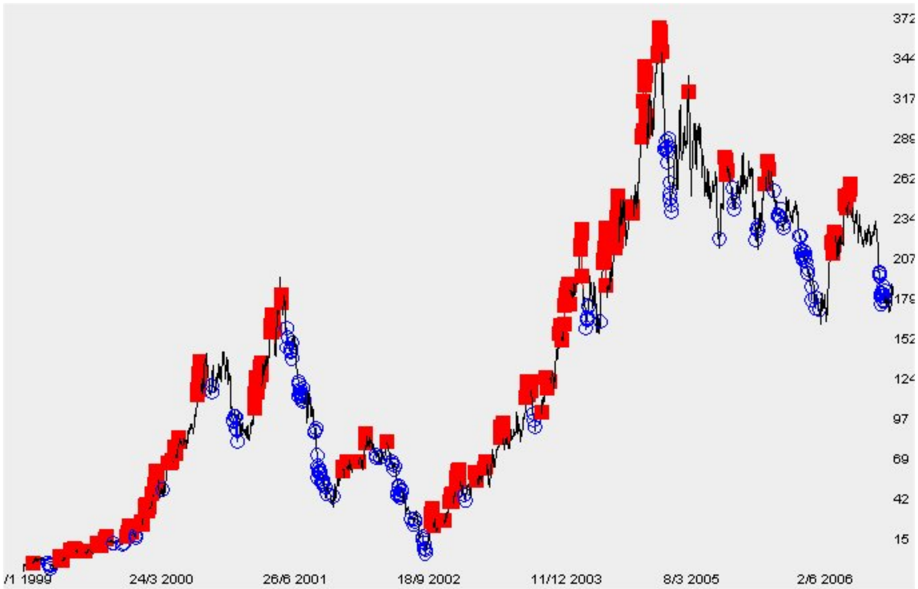

Figure 5.6: RSI 14 days agent signals for FRO-stock 1999-2006

5.5.3 Moving average agents

The moving average and crossover moving average agents are the best performing simple agents in spite of more than half of all their trades being losing ones. This coincides good with the less fortunate behaviour of the RSI and Price Envelope band agents: The moving average agents are designed to spot trends and follow them. As the winning stocks are clearly in an up trend, the agents has managed to follow the major moves and performed badly only under the less hazardous conditions when price moves are only small, quite contrary to the above mentioned agents. Figure 5.7 shows the 40 days moving average agents signals over the same period as the above illustrated RSI 14 days agent. It understates the difference between them, with the Moving average agent as a rule always signaling buy on relative tops and sell on relative bottoms but being able to catch the whole 2002-2004 price advance with buy signals.

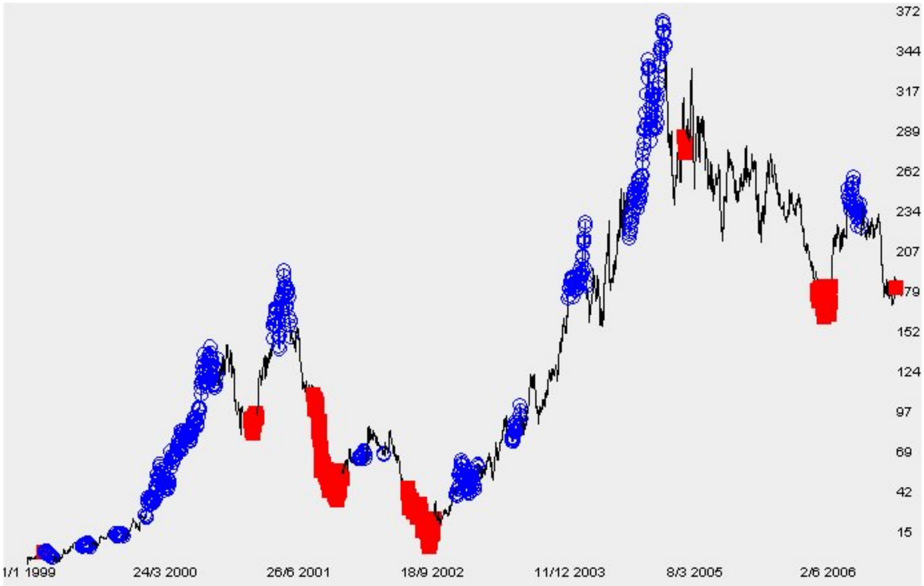

Figure 5.7: Moving Average 40 days agent signals for FRO-stock 1999-2006

5.5.4 Trend agent

The trend agent's all-in yield of 71.8% is not impressive given the fact that it should be able to determine a clear trend such as a major up one. The introduced lag of the agent need to shoulder the responsibity for this fact: the agents needs a set of established higher relative highs and lows in order to determine whether a trend in one direction has been established. Figure 5.8 illustrates this, with the agent first correctly determining that the period december 2003-december 2004 being a downtrend, and needing much of the major uptrend thereafter to be certain that the trend has shifted. The short position between december 2004 and mid-2005 yields a more than 100% negative result, leaving the agent in great dept with an all-in trade strategy, unable to ride the several subsequent good buy signals with a long position. Figure 5.9 tells a slightly different story: the more steep advances of new relative high and lows gives the agent the opportunity to decide trend direction slightly earlier, thus yielding much better from trading in the stock. This difference accounts for the fact that the trend agents overall performance is quite good but shows weak signs, particular for the all-in approach. In addition, the average of less than 2 trades per stock also indicates that the agent does not shift its market positions often, giving few transactions to realize winning trades.

5.5.5 Gap agents and Island reversal

The good performance of the Gap agents and the gap based Island reversal agent can be explained by the fact that major up and down moves seem to be started with gapping prices, making the agents able to time entry and exit

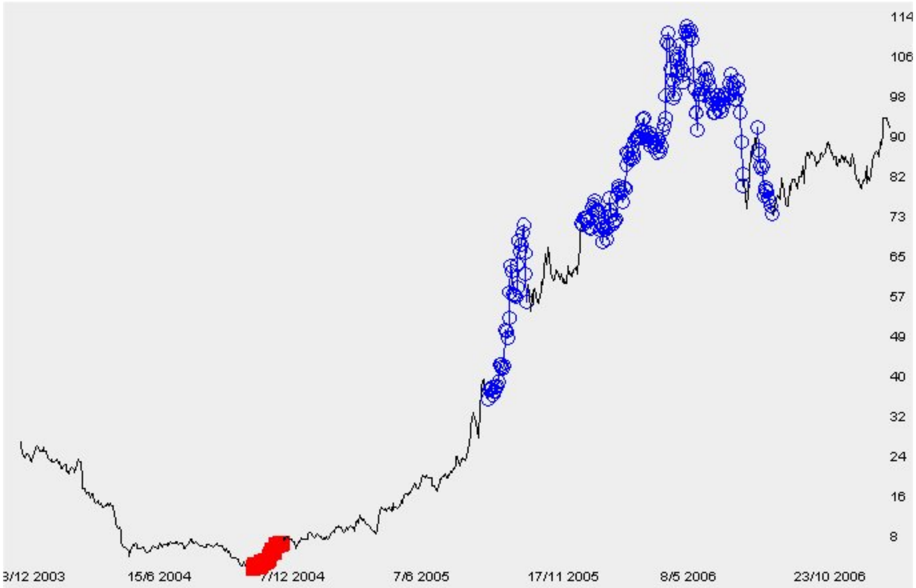

Figure 5.8: Trend Agent signals for NAS stock, 2003-2006

points well. Figure 5.10 illustrates this fact. Most of the buy and sell signals are given near relative tops, but with the overall trend taken into account the agents are performing well.

5.6 The losing stocks

Table 5.4 shows how the different agents would perform when trading in stocks that has showed a negative price developement. This performance is somehow quite similiar to the one for winning stocks and this can be explained by the fact that the nature of the price move is the same, only reversed being down instead of up.

5.6.1 The Optimal signal voting classifier agent

The pleasant performance of the Optimal signal voting classifier agent is one of the most important oberservation to be made. It is a proof of what the overall performance indicated regarding the agent — its ability to catch extreme declines with a negative stock opinion.

5.6.2 The Neural superagent

The good performance of the Neural superagent for losing stocks compared to that for winning ones is an evidence that the finished trained agent has a

Agent Name	Yield	Profitable trades	All-in yield
Buy and hold	-83.0%	0.0% (of 23)	-83.0%
Double Top/Bottom	33.0%	43.1% (of 51)	54.6%
Island Reversal	9.8%	47.5% (of 343)	123.9%
Moving Average 14 (5)	66.2%	41.9% (of 375)	43.5%
Moving Average 40 (10)	35.3%	44.4% (of 133)	10.0%
Moving Average 200 (10)	23.7%	60.0% (of 45)	9.9%
Crossover MA 10+100	47.4%	49.0% (of 2250)	-5.9%
Price Envelope Band	-32.6%	67.6% (of 723)	-51.3%
Breakaway Gap	55.9%	43.3% (of 164)	22.9%
Runaway Gap	51.3%	48.4% (of 128)	91.0%
RSI 9 days	-52.5%	57.4% (of 627)	-71.2%
RSI 14 days	-78.9%	56.2% (of 406)	-87.8%
RSI 21 days	-37.4%	55.7% (of 212)	-69.2%
Head Shoulders	-15.5%	32.8% (of 58)	3.2%
Sell in may	65.9%	58.1% (of 172)	21.1%
Trend	69.4%	81.5% (of 27)	65.3%
Volume	-27.9%	36.8% (of 581)	-29.4%
Hanging man	-73.1%	54.0% (of 454)	-54.2%
Reversal day	-38.0%	42.4% (of 309)	-53.0%
Neural Net	38.5%	53.4% (of 163)	45.4%
ID3 Tree	-27.0%	45.2% (of 5697)	-16.0%
J48 Tree	-57.3%	45.2% (of 4222)	-54.0%
Voting Classifier Buy & sell	-87.2%	43.5% (of 1638)	-32.3%
Voting Classifier Optimal	77.2%	46.4% (of 483)	139.7%
Voting Classifier Extreme	-20.8%	24.4% (of 1746)	52.4%
Neural Superagent	91.6%	50.7% (of 138)	133.8%

Table 5.4: Agent performance for losing stocks

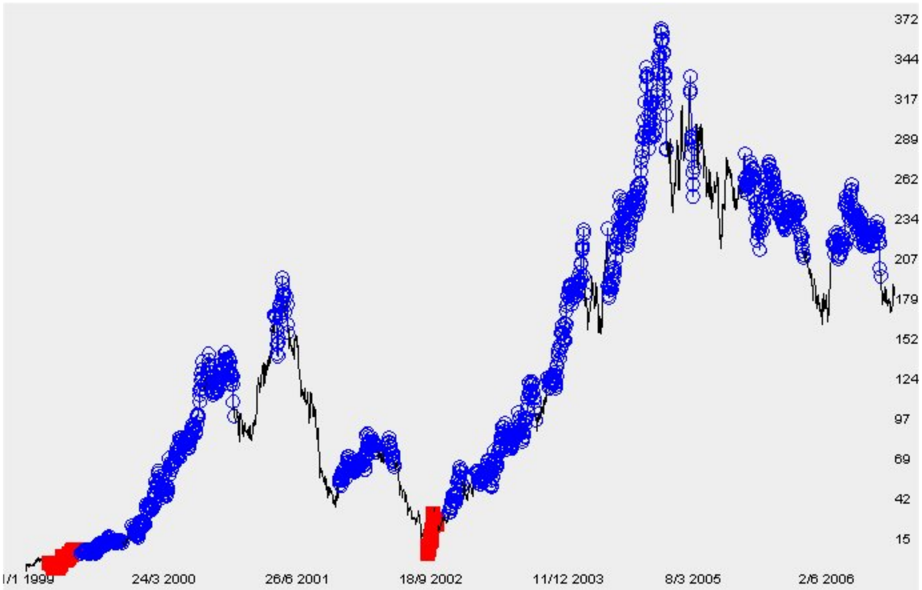

Figure 5.9: Trend Agent signals for FRO stock, 1999-2006

disposition for outputting sell signals more often than buy signals, and is a clear indication that the training has resulted in an agent overweighting negative opinion during all kinds of market trends, impairing its validity as an isolated trading indicator.

5.7 The volatile stocks

Table 5.5 shows the agents performances when trading in volatile stocks - stocks moving in both directions of the scale. Such movements constitute the best trading oppurtunities for a trader able to take buy positions near local bottoms and sell positions near local highs with the continous profit frequent realization of such positions can produce.

5.7.1 The Gap agents

Perhaps the most interesting observation is that the two Gap agents perform different during the volatile markets — the breakaway gap agent earning money and the runaway gap agent losing. An explanation for this behaviour can be found by studying figure 5.11 and 5.12 which shows the two agents signals for the volatile KOG stock in the period 2001-2006. The figures shows how the breakawy gap agent picks up trends more often than the runaway gap agent. The formers signals has several successful trades, as most of the succeding signal shifts occurs on points at which sell signals are above buy signals, and buy signals below sell signals. With the same price movement, trading based on the

Agent Name	Yield	Profitable trades	All-in yield
Buy and hold	40.1%	75.0% (of 24)	40.1%
Double Top/Bottom	42.6%	44.4% (of 54)	13.0%
Island Reversal	167.5%	40.3% (of 258)	49.8%
Moving Average 14 (5)	89.5%	41.3% (of 530)	95.4%
Moving Average 40 (10)	52.6%	37.0% (of 219)	22.5%
Moving Average 200 (10)	-44.1%	32.9% (of 70)	-37.6%
Crossover MA 10+100	130.1%	48.5% (of 2063)	47.4%
Price Envelope Band	-60.9%	64.9% (of 878)	-50.5%
Breakaway Gap	95.5%	43.0% (of 237)	335.8%
Runaway Gap	-10.8%	35.6% (of 202)	-36.7%
RSI 9 days	-42.4%	59.2% (of 901)	-44.5%
RSI 14 days	-41.1%	58.1% (of 566)	-76.8%
RSI 21 days	-82.6%	55.8% (of 312)	-73.9%
Head Shoulders	-38.8%	38.3% (of 60)	-29.4%
Sell in may	76.8%	54.1% (of 231)	67.4%
Trend	-124.3%	22.6% (of 53)	-76.6%
Volume	155.4%	42.6% (of 702)	208.4%
Hanging man	-86.5%	52.8% (of 614)	-53.7%
Reversal day	-78.4%	40.7% (of 236)	-25.2%
Neural Net	65.8%	46.8% (of 205)	47.0%
ID3 Tree	38.6%	47.9% (of 8181)	5.9%
J48 Tree	101.9%	47.1% (of 5601)	279.6%
Voting Classifier Buy & sell	25.6%	43.7% (of 2711)	13.3%
Voting Classifier Optimal	169.1%	44.9% (of 532)	771.9%
Voting Classifier Extreme	105.3%	27.3% (of 2543)	189.7%
Neural Superagent	-64.4%	34.5% (of 197)	-77.4%

Table 5.5: Agent performance for volatile stocks

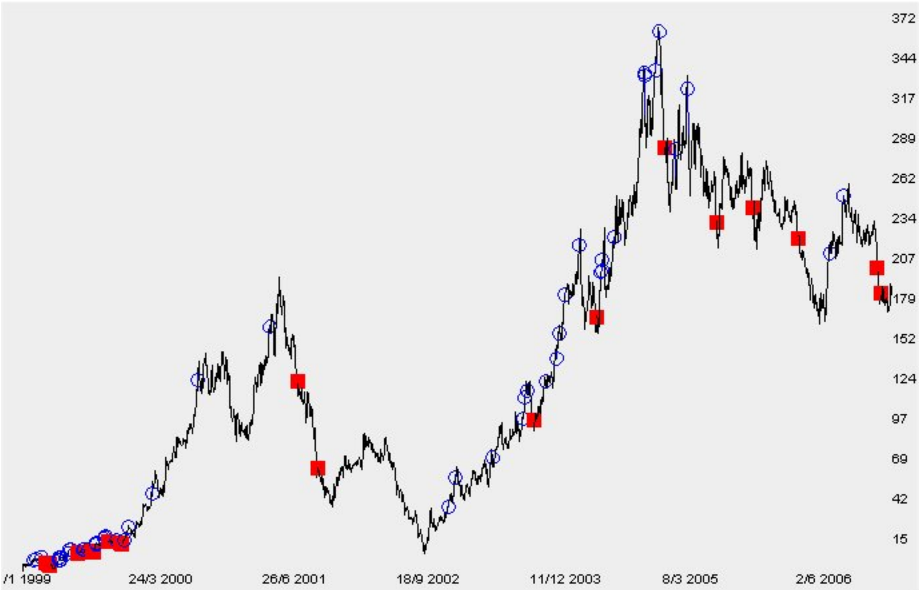

Figure 5.10: Gap agent signals for FRO stock, 1999-2006

Breakaway gap agent would in comparison result in three transactions being made, all losing. This is because the agent misses some of the more important up- and down moves with its strict rule set.

5.7.2 The volume agent

The volume agent is one of the best performing agents during the volatile periods. This indicates that sudden market interest in particular volatile stocks is a good pointer of when their trend is about to reverse.

5.7.3 The RSI and Price envelope band agents

The performance of the RSI and Price envelope band agents for the volatile stocks is quite opposite to what is expected from them. Waving formations are the perfect medium for successfull range trading. The fact that they yield negative results indicates that the overbought- and oversold threshold values probably are not strict enough: the volative movements are determined overbought and -sold at to early moments during the most profit-yielding fluctuations. The profitable trade ratio of 64.8% for the Price envelope band agent and close to 60% for the RSI agents does however suggest that most of the wave entry and exits are timed correctly according to theory.

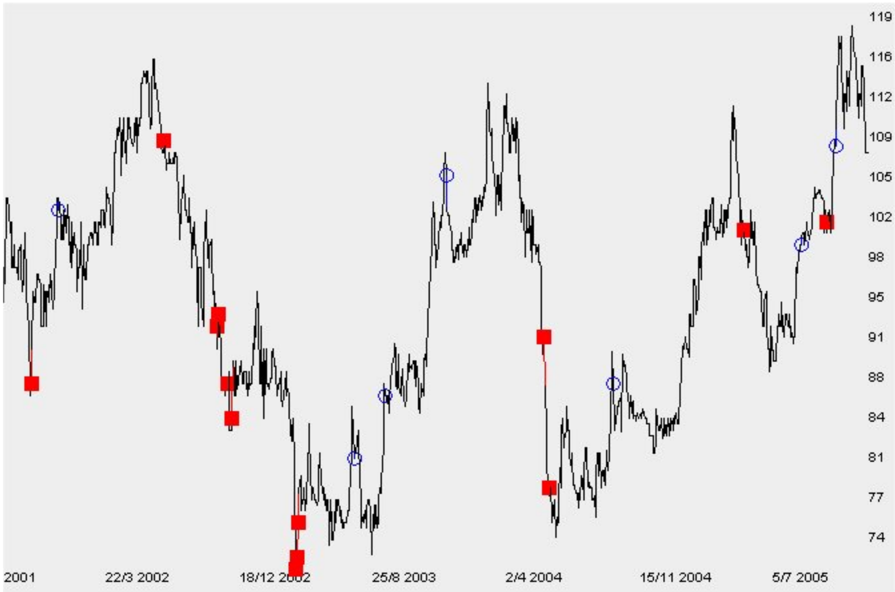

Figure 5.11: Normal GAP agent signals for KOG-stock 2001-2006

5.7.4 Optimal signal voting classifier

Clearly outperforming all other agents with 771.9% all-in yield, the previously indicated strengt of the Optimal signal voting classifier agent is confirmed. During volatile periods, it is able to shift its positions at advantegeous points, realizing several winning trades.

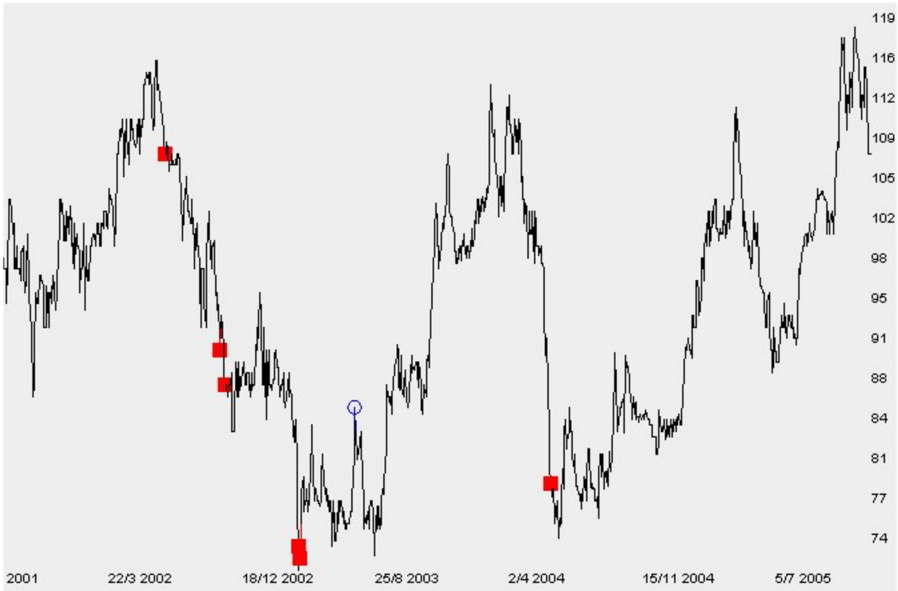

Figure 5.12: Breakaway GAP agent signals for KOG-stock 2001-2006

Chapter 6

Conclusion

Despite the positive outcome during testing, the implemented agents general usage potenial in real-world situations are far from proven. They do lay a good foundation for further improvement and testing of the implemented system, which can serve as an important tool for an investor able to use it with caution.

6.1 Future work

Observations from the test results and factors not taken into account upon developement gives the implemented program a large range of improvement potential.

6.1.1 Adaptive parameters

The test results indicates that, for many of the implemented trading agents, the parameters for technical patterns used has been either too loosely or too strictly set. Agents with abilities to dynamically learn from their mistakes and use input from other relevant agents to shift parameters may be able to adapt better to unexpected market fluctiations. Examples of such a feature is an RSI agent gradually raising the overbought treshold value for each consecutive day it is penetrated.

6.1.2 Using stop loss points and protecting profits

Success of trading is dependent on effective control of losses. It is not necessary to be always be right as long as the winning trades are substansial enough to return an overall profit. Automatic selling of bought stocks when trend lines are penetrated, trading ranges are exceeded, or prices has retracted back to the last relative high or low are simple approaches for effectively minimizing the amount of big losses. Stop loss points are not only used for avoiding losses, it can also be applied for profit protection. During up trends, stop points can for instance be raised intermittently as the market rises. When such a stop point is penetrated, profits can be secured by stopping positions when a reversed trend starts. Figure 6.1 illustrates how such an approach could work in practice. In this illustration, the first stop point is set 10% lower than the initial buy, and is gradually risen to 10% below the last relative high when prices increase more than 30% from the last stop point.

Figure 6.1: Illustrative use of stop points: TAT stock 2003-2006

6.1.3 New patterns detection

Before a head and shoulders formation is finalized, a large price decline has already been realized during the finalization of the last shoulder pattern. Being able to trade in a stock at earlier points of formations could result in far more profitable pattern recognizing agents than those implemented. Such obersvations can lead to detection of more profitable technical formations than those now mainly in use. By extracting and defining new movement patterns from price increases or declines, the system can get an edge over those only based on traditional, well-known definitions.

6.1.4 Genetically evolved decision trees

Decision trees have a big advantage over the voting classifiers, regarding the ability to express specified relations between a subset of the agents, e.g give weights to RSI agents only when the Trend agent indicates a sideways moving market.

The creation of the trees based on the Weka machine learning framework was proven insufficient for the test data. Introduing genetic algorithms with its ability to measure a complete solution over the training set during evolving, one might be able to generate decision trees with overall much better performance than those implemented.

6.1.5 Using fundamental factors

Technical analysis cannot tell wheter a company is performing well - it can only give a hint of how the stock market players' general psychology towards the particular stock as an asset or the stock market as a whole is. The greatest investor of our time, Warren Buffet, has over the whole of his lifetime emphasized his love for fundamental factors and avoided companies where profits not easily can be measured in cash [25, page 240]. Buffet lost out on the incredible bubble pricing of the dotcom period at the end of the 1990s and beginning of 2000s because of his aversion of speculative technology stocks without proven income [25, page 12]. Using Buffet's fundamentally based scepticism during the the dotcom years combined with the signals from the Optimal signal voting classifier agent would have resulted in little profit during the up cycle but still extremely good results for a short selling investor when the market psychology turned and both the fundamental analysis and technical signals forecasted price declines History has provided several such contradictory fundamental and technical investement cases with varying outcome, but with a longer perspective, it seems clear that combining the two ideologies gives a clear edge on the strictly pavlovian approach. Fundamental analysis is therefore without a doubt a natural part of a general stock picking system, and anything from simple divident yield or earnings growth calculations to commodity performace, peer stocks evaluation and general sector comparisons seems like good supplements to technical factors in a trading system.

6.1.6 Commision

Trading at a stock exchange must be done through a broker which carries orders out for the investor. Each transaction is subject to a fee to be paid to the broker for the accomplishment of the trade, named brokerage commission. With too many transactions being carried out in order to achieve it, earnings can vanish through these fees. Keeping the number of transactions at a minimum when investing in the market is therefore necessary to maximize profits. Awarding agents producing few trades seems logical to achieve more dependable training and test results.

6.2 Reliability

The Optimal signal voting classifier agent turned out be profitable. The insecurity of the stock market pricing makes historical achievements no safety in prediciting future results but the agents positive performance in both major down and up cycles and the clear overweight of extremely high yielding trades compared to extremely low makes it a good bet to catch the big winnings the stock market can yield.

The testing has been done using a diversified portfolio consisting of all avaible stocks above a minor level of daily trades on the Oslo Børs stock exchange. This has given the agent a proven effect when an investor is able to spread its money. With only limited funds available, the natural consequence is that trading takes place in only a few stocks. In this case, results have the potential to vary in great degree on both the up- and downside. A more cautious trading approach, using technical signals combined with stop loss points for stock trades already founded on fundamental factors, would probably be the most logical application of the implemented agent. In such cases, one would be able to avoid trading against the market psychology in strong trends, but still be certain that ones investements are fundamentally well grounded and not pure gambling.

Bibliography

[1] Jeremy J. Siegel. *Stocks for the long run*. McGraw-Hill, third edition, 2002.

[2] Geir Linlokken and Steffen Frohlich. *Technical Stock Analysis - for reduced risks and increased returns*. Investtech.com, 2004.

[3] Erstattes av børsroboter. http://www.dn.no/forsiden/borsMarked/article945925.ece. Accessed 01.02.2007.

[4] Jack D. Schwager. *Getting started in Technical Analysis*. John Wiley and Sons Inc, first edition, 1999.

[5] Support and resistance, trend line and channels. http://www.2bull.com/learn/sp.htm.

[6] John J. Murphy. *Technical Analysis of the Financial markets: A comprehensive guide to trading methods and applications*. Prentice Hall Press, 1999.

[7] Relative strength index. http://www.gannalyst.com/Gannalyst_Professional/Gannalyst_Indicators_Relative_Strength.shtml.

[8] Tradingday.com. Technical analysis tutorial. http://www.tradingday.com/c/tatuto/thegapprimer.html. Accessed 23.11.2006.

[9] Island reversal. http://www.shareselect.com.au/showpage.asp?ID=Technical+Analysis+-+Gaps+-+Island+Reversal. Accessed 07.02.2007.

[10] Leigh Stevens. *Essential Technical Analysis: Tools and Techniques to Spot Market Trends*. John Wiley and Sons, inc., 2002.

[11] David Chapman. Sell in may and go away! http://www.gold-eagle.com/editorials_04/chapmand042904.html. Accessed 13.02.2007.

[12] Yahoo Finance. Sp 500 index: Historical prices. http://finance.yahoo.com/q/hp?s=%5EGSPC.

[13] Mark W. Riepe. Sell in may and go away? not so fast. http://www.fpanet.org/journal/articles/2003_Issues/jfp0703-art3.cfm. Accessed 26.01.2007.

[14] Tror på ti prosents børsoppgang. http://www.dn.no/forsiden/article111637.ece. Accessed 26.01.2007.

[15] Chew Lim Tan Jingtao Yao and Hean-Lee Poh. Nural networks for technical analysis: A study on klci. *International Journal of Theoretical and Applied Finance Vol. 2, No. 2 221-241*, 1999.

[16] The r project for statistical computing. http://www.r-project.org. Accessed 08.02.2007.

[17] William J. O'Neil. *How to Make Money in Stocks: a winning system in good times or bad*. McGraw-Hill, 2nd edition, 1995.

[18] Jgap: Java genetic algorithms package. http://jgap.sourceforge.net.

[19] Melanie Mitchell. *An introduction to genetic algorithms*. MIT Press, first edition, 1998.

[20] Ian H. Witten and Eibe Frank. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, San Francisco, 2nd edition, 2005.

[21] Kursoversikt. http://hopey.netfonds.no/kurs.php. Accessed 09.02.2007.

[22] Slutt for totalindeksen. http://www.dn.no/arkiv/article19012.ece. Accessed 08.02.2007.

[23] Historical index data. http://www.oslobors.no/ob/index_history. Accessed 08.02.2007.

[24] Oslo børs benchmark index (osebx). http://hopey.netfonds.no/index-history.php?index=OSEBX.OSE. Accessed 08.02.2007.

[25] Warren Boroson. *J.K. Lasser's Pick stocks like Warren Buffet*. John Wiley and Sons, Inc., 2001.

Appendix A

Test cases

A.1 Training set

To train the agents, american stocks from the New York Stock Exchange in the period between 1970 towards end 1998 has been used:

ALU 93-98, BOL december 83-june 91 and 92-99, C april 87-july 93, CAT 87-august 94, COP 86-98, DE 82-november 95, DIS february 73-february 86, F 88-may 94, FL 71-86, FPL march 85-99, GLW march 92-march 96, GT 80-may93, HPQ 85-95, IBM aug79-97, JNJ 71-81 and 82-89, MUR april 83-99, PFE 84-91 and may 91-may 95, RAD 92-98, RSH 82-august 97, S 90-99, SLB may 83-may97, TIN april 90-99, WEN june-99, XOM 88-97.

When not otherwise stated, the stock has been used from and till january the given year.

A.2 Tickers used

Ticker symbols for stock histories used in the test set:

ASC, ACY, ACTA, AKER, AKASA, AKD, AKFP, AKVER, AKS, AKY, AIK, ALX, APL, APP, AWO, ASD, GAS, BIOTEC, BIRD, BJORGE, BLO, BLU, BON, CECO, CNR, CAPTU, CEQ, COV, COR, CRU, DAT, DEEP, DESS, DIAG, DNBNOR, DNO, DOF, DOM, DYNA, EDRILL, EDBASA, EIOF, EKO, ELT, EME, EXPERT, FARA, FAR, FAST, FOE, FRO, FUNCOM, GRO, GEO, GGS, GOL, GOGL, GRR, GGG, HNA, HAVI, HEX, IGNIS, IMAREX, INM, IGE, ITE, JIN, KIT, KOM, KOG, KVE, LSG, MEC, NTL, NAVA, NOD, NEC, NSG, NHY, NAS, NUT, OCR, ODF, ODIM, OPERA, ORK, OTR, PAN, PAR, JACK, PHO, POLI, PRS, PSI, QFR, QEC, REVUS, RGT, RCL, SAS-NOK, SCI, SCH, SCORE, SDRL, SEVAN, SIOFF, SIT, SIN, SOI, SOFF, SONG, SRI, STL, STP, SNI, STB, SUB, SUO, TAA, TAD, TST, TAT, TCO, TEL, TGS,

TOM, TREF, VEI, WWI, YAR

A.2.1 Subset: winning stocks

AIK, ACTA, AKER, AKVER, AKY, ASC, BJORGE, BON, CRU, DNBNOR, DNO, DOF, EXPERT, FAR, FOE, FRO, GOGL, GRO, GRR, HEX, HNA, IGE, KOM, LSG, NAS, NHY, ODF, ODIM, ORK, PAN, PAR, PRS, SCH, SCI, SDRL, SEVAN, SIOFF, SIT, SOFF, SRI, STL, SUB, SUO, TEL, TGS, TST, VEI, WWI, YAR

A.2.2 Subset: losing stocks

ACY, ALX, DOM, EME, KVE, NUT, PHO, TAD, TCO, ASD, CNR, COR, EDBASA, ELT, IGNIS, ITE, OPC, OTR, PAN, SCH, SOI, STP, TOM

A.2.3 Subset: volatile stocks

APP, BIRD, CECO, CNR, FARA, FAST, GOL, KIT, KOG, MEC, NEC, NOD, NSG, NTL, OTR, PSI, QFR, RCL, REVUS, SAS-NOK, SNI, STB, TAT, TAA

A.3 Tickers not used

Ticker symbols for stocks that has been excluded from the test set

A.3.1 Illiquidity

Stock exluded from training set because of several sets of to few consecutive trade days:

AGR, AAV, AFG, AFK, AGI, AURG, BHOC, BEL, CNS, EID, EXE, FSL, FIRM, FOS, RISH, GOD, GYL, HELG, HJE, HOLG, HSPG, IMSK, ISSG, LUX, MAMUT, MRG, MEDI, MELG, NAM, NESG, NORD, NORGAN, NORMAN, NOV, NVF, OILRIG, ODFB, OLT, PDR, POWEL, PRO, RIE, RING, RVSBG, SADG, SANG, SKI, SOLV, MING, MORG, NONG, PLUG, ROGG, SVEG, VSBG, SPOG, STA, SST, SFM, TECO, TOTG, TTS, WILS, VIZ, VVL

A.3.2 Too late listing

Ticker symbols for stocks that got listed for trade at the stock exchange to late to produce enough historical data:

AKFP, AUSS, BWG, BOR, BWO, CLAVIS, COD, CMI, DESSC, DOLP, ECHEM, EMS, FAIR, FAKTOR, HRG, IOX, KOA, MAFA, NAUR, NPRO, PERTRA, PBG, REC, REPANT, RXT, SBX, SPITS, TPO, TELIO, TIDE, TROLL

Appendix B

The program

The program requires java runtime environment version 5.0 or later and needs to be started using the following syntax from the command line in the root folder of it:

java -Xmx128m -cp .;jgap.jar no/ntnu/stockmeks/StockTerminal <input paramters>

By means of input parameters, the implemented program invokes different functionality:

- **gui** or no parameters — Starts the graphical user interface
- **tree** — writes a decision tree training set to the file "train.arff"
- **evaluate** — starts evaluation of all agents and writes them to the file "evaluation-table.txt" and a summary of all transactions performed to the file "transactionsummary.txt." All transactions for all stock tickers in the training set are logged in the subfolder "transactionlogs"
- **evaluatem** — starts evaluation of the voting classifier performance over each month in the training set and outputs the results to the command line. Tables showing monthly performance is written to the subfolder "monthly"
- **geneticsbuysell <population size> <number of evolutions>** — starts evolving a population of the given size over the number of evolutions using the Buy & Sell fitness function. The best chromosome of each evolution is written to the file "genetics.txt"
- **geneticsoptimal <population size> <number of evolutions>** — starts evolving a population of the given size over the number of evolutions using the Optimal signal fitness function. The best chromosome of each evolution is written to the file "genetics.txt"

- **geneticsextreme <population size> <number of evolutions>** — starts evolving a population of the given size over the number of evolutions using the Extreme signal fitness function. The best chromosome of each evolution is written to the file "genetics.txt"