

ASTA

a Test Bench and Development Tools for Trading Algorithms

Thomas Hellström¹
Department of Computing Science
Umeå University
S-901 70 Umea, Sweden

PRELIMINARY REPORT

Presented at EUFIT-98 Aachen, Germany. Sep 10, 1998

Abstract

This paper describes the principles behind and the implementation of **ASTA**, an Artificial Stock Trading Agent written in The Matlab language. The primary purpose of the project is to provide an easy-to-use environment for developing multi-stock trading algorithms. A key ingredient in the system is a stable and realistic test bench, highlighting and helping to avoid the huge risks of data snooping involved in this kind of financial prediction. The behavior of the agent is controlled by buy and sell rules, which can be composed interactively or written as user defined functions. Various types of combinations of rules and data screening can be easily performed, all within the Matlab language syntax.

In addition to being a Windows-based developing tool and test bench, the system can be run in batch mode, where a supplied objective function maps a trading strategy to a profit measure. This can be used to tune parameters or automate the development of trading strategies, for example with genetic methods. Examples of tuning parameters in standard technical indicators are presented. A modified Sharpe Ratio is used as performance measure.

To improve the performance of a given algorithm, the data from the simulated trades can be output and post-processed by classification methods, such as artificial neural networks or fuzzy-rule bases.

The **ASTA** system has been applied successfully to historical stock data, and results covering 11 years of the Swedish stock market are presented.

Keywords: trading, artificial trader, technical analysis, data mining, prediction, stock returns, Sharpe ratio, parameter tuning.

¹thomash@cs.umu.se, <http://www.cs.umu.se/thomash/>

1 Introduction

The idea of expressing stock prediction algorithms in the form of trading rules has gained considerable attention in academic research in the last years. The international conference NNCM-96 devoted a whole section in the proceedings to "Decision Technologies". Bengio [3] writes about the importance of training artificial neural networks with a financial criterion rather than a prediction criterion. Moody and Wu [9] use reinforcement learning to train a trading system with objective functions such as profit, economic utility and Sharpe ratio. Atiya [2] describes a trading system based on time-variable stop-losses and profit objectives.

This paper describes the system ASTA, which is an implementation of an Artificial Stock Trading Agent. With ASTA, trading-rule based prediction algorithms are easily evaluated using historical data. The system is also a high level development tool where trading rules can be developed, combined and tuned.

The program is written in the Matlab programming language and is used either as an ordinary objective function called from a user's program, or as an Windows based tool for making benchmarks and development of trading algorithms.

1.1 Objectives

The development of ASTA was initiated by a need for good working tools for the following research tasks:

1. A test bench for trading algorithms.

Many "technical indicators" for stock prediction are accepted and widely used without having ever been subject to an objective scientific analysis with historical stock data. It is true that many commercial software packages for technical analysis offer both a comprehensive programming language and a simulation mode where the performance can be computed. However, most available products do not take this task very seriously and real trading simulation with a multi stock portfolio is seldom possible. Furthermore, often the performance measures are not sufficient for a serious evaluation of the behavior of an algorithm for a longer period of time. Therefore, there is a need for a scientific test bench for the methods and algorithms already developed and in common use.

The need for proper evaluation of *new* trading algorithms is of course the same as for existing ones. ASTA is developed in Matlab and therefore is suitable for tests of algorithms developed in the same language, but Matlab can communicate with other languages.

2. An interactive development tool for trading rules.

There are reasons to believe that a successful trading system consists of many disjunct parts where a buy signal can be, for example, "screened" by looking at the traded volume. A buy signal issued with a low traded volume may then be rejected. Other composite rules include looking at the general trend of the stock before accepting a signal from the system. ASTA provides the possibility to test such composite rules easily. The included function library and the possibility to define a trading strategy interactively makes implementation and evaluation of many trading strategies possible "without programming".

3. A non-interactive development tool for trading rules.

Furthermore, it is possible and maybe also fruitful to automate the development of trading rules. Since ASTA defines the trading strategy as symbolic **Buy rules** and **Sell rules** given as arguments to the system, it would be perfectly possible to construct buy and sell rules, in a genetic framework, for example.

Even if the general look of the algorithm is fixed, there are often a lot of tunable parameters that affect the trading performance. Examples are filter coefficients, order of polynomials and levels above or below which an entity should pass in order to generate a trading signal. Since we believe that the actual behavior during a realistic trading situation is essential for proper selection and optimization of an algorithm, there is a need for an objective function that can be included in an optimization phase for parameter tuning.

4. A data generation tool for post processing

The comprehensive and user-friendly macro language in ASTA makes it a very suitable tool for extracting data for further analysis, such as classification with neural networks or fuzzy rule bases. A raw selection of trading situations is first set up with simple trading rules. Data ("features") for these situations are then automatically written to a file. A "target" value for each trade is also supplied. Thereafter it will be a classification task to find out how to distinguish between a profitable trade and a non-profitable one, based on the given features.

2 Basic approaches to Stock predictions

Prediction algorithms for stock prices can be categorized in a number of ways. One categorization focuses on the way the points to predict are selected. Two broad classes can be identified; "The Time Series Approach" and "The Trading Rule Approach"

2.1 The Time Series Approach

The traditional way to define a stock prediction problem is to view the stock returns as a time series $y(t)$. For example, one-day stock returns are defined as

$$y(t) = \frac{Close(t) - Close(t-1)}{Close(t-1)}. \quad (1)$$

To predict the return h days in the future, $y(t+h)$ is assumed to be a function g of the previous (lagged) values

$$y(t+h) = g(y(t), y(t-1), \dots, y(t-k)). \quad (2)$$

The task for the learning or modeling process is to find the function g that best approximates a given set of measured data. The unknown function g can be chosen in many ways. Common choices are:

- A linear AR model:

$$g(t) = \sum_{m=0}^d a_m y(t-k) + a_{-1} \quad (3)$$

- A general nonlinear model implemented by a multi-layer-feed-forward neural network. A 1-hidden-layer net defines g as

$$g(t) = \sum_{j=0} \left(w_j h \left(\sum_{m=0} w_{m,j} y(t-k) + w_{-1,j} \right) + w_{-1} \right) \quad (4)$$

where h is a nonlinear function such as the sigmoid function $h(x) = \frac{1}{1+e^{-x}}$.

The unknown parameters (a_m for the AR model and $w_j, w_{m,j}$ for the neural network) are normally computed by the learning algorithm to minimize the root mean square prediction error:

$$RMSE = \sqrt{\frac{1}{N} \sum_{t=1}^N (g(t) - y(t+h))^2}. \quad (5)$$

It is most common to let the minimized RMSE measure (5) be the end point in the prediction task. However, in order to utilize the predictions, a decision-taking rule has to be created. A simple rule often used is

$$D(t) = \begin{cases} \text{Buy} & : \text{if } g(t) > \alpha \\ \text{Sell} & : \text{if } g(t) < -\beta \\ \text{Do nothing} & : \text{otherwise} \end{cases}, \quad (6)$$

where α and β are threshold parameters for buy and sell actions depending on the predicted change in the stock price.

The time series formulation based on the minimized RMSE measure (5) is not always ideal for useful predictions of financial time series. Some reasons for that are:

1. The fixed prediction horizon h does not reflect the way in which financial predictions are being used. The ability of a model to predict should not be evaluated at one single fixed point in the future. A big increase in a stock value 14 days into the future is as good as the same increase 15 days into the future!
2. The characteristics of the underlying processes that generate the time series are dramatically different in different parts in the series.
3. The RMSE measure (5) treats all predictions, small and large, as equal. This is not always appropriate. Prediction points that would never be used for actual trading (i.e. price changes too small to be interesting) may cause higher residuals at the other points of more interest, in order to minimize the global RMSE.
4. A small predicted change in price, followed by a large real change in the same direction, will be penalized by the RMSE measure. A small predicted change in price, which actually resulted in a much higher value, but in the same direction, will be penalized by the RMSE measure. A trader is normally happy in this case, at least if, say, the small positive prediction was large enough to give a buy signal.
5. Several papers report a poor correlation between the RMSE measure and the profit made by applying a prediction algorithm, e.g. Leitch & Tanner [5] and Bengio [3]. A strategy that separates the modeling from the decision-taking rule, such as the one in 6, is less optimal than modeling the decision-taking directly [8]. Arguments 3 and 4 both give some explanations to these results.

2.2 The Trading Rule Approach

Instead of separating the prediction task and the decision task as was done in the "Time Series Approach", algorithms can be constructed to recognize situations where one should buy and sell stocks respectively. A trading rule can be described as a time series $T(t)$ defined as

$$T(t) = \begin{cases} \text{Buy} & : \text{if } g(\mathbf{X}(t)) > 0 \\ \text{Sell} & : \text{if } g(\mathbf{X}(t)) < 0 \\ \text{Do nothing} & : \text{otherwise} \end{cases}. \quad (7)$$

The unspecified function g determines the type of the trading rule.

The argument $\mathbf{X}(t)$ has the form

$$\mathbf{X}(t) = (R_1(t), \dots, R_N(t)), \quad (8)$$

where each $R_k(t)$ is an observable feature at time t . In the case of stock predictions it may be for example the k -day returns defined as

$$R_k(t) = 100 \cdot \frac{\text{Close}(t) - \text{Close}(t - k)}{\text{Close}(t - k)} \quad (9)$$

or standard technical indicators such as the Stochastic Oscillator, the Relative Strength Index (RSI) or the Moving Average Convergence/Divergence (MACD) [1].

The task for the learning process in The Trading Rule Approach is to find the function g so as to maximize the profit, when applying the rule on real data. Various ways to measure the profit

are discussed in section 4.1. Note the difference between this and The Time Series Approach, where the learning task is to find a function g that minimized the $RMSE$ error (5) for the entire time series.

The Trading Rule Approach avoids many of the problems previously described of the Time Series Approach but does indeed have problems of its own:

1. Statistical significance. The trading rule $T(t)$ normally gives Buy or Sell signal only for a minor part of the points in the time series. While being one of the big advantages, it also presents serious statistical problems when computing levels of significance for the produced performance. It's no problem to find a rule $T(t)$ that historically outperforms any stock index as long as it doesn't have to produce more than a few signals. It's clear that a rule $T(t)$ that produces many signals for a given time interval, is more reliable than one that just produces a few, even if the profit is the same. However, it's *not* clear how this tradeoff between profit and number of signals should be settled in an optimal way. The selection bias when choosing "good" trading rules makes a complexity argument worthless since even a very simple trading rule can produce very high profit if it just has to signal a few times. The general problem remains unsolved in the present work, with the single but important observation that the number of produced signals should be as high as possible in order to render the trading rule credible.
2. Commonly used methods for modeling or inductive learning, such as Regression models and Artificial Neural Networks cannot be applied in a normal fashion since a set of training data with Patterns and Targets are not available in the same way as in the Time Series Approach. However, a related modeling problem can be formulated as an extension to a system based on trading rules. The points where $T(t)$ signal *buy* or *sell* can be used to extract examples that can be input to a modeling or classification routine. The trading rule T may be refined in this way to produce a higher hit rate and most probably a higher overall profit. Therefore, a data extraction function has been built into the **ASTA** system.

The ASTA system is constructed according to the "Trading Rule Approach". In the next section the evaluation of a prediction system based on these principles will be investigated.

3 What Is Good Performance

A way to measure the performance of an algorithm is needed in two phases of the trading system's development cycle. Firstly, during the modelling phase where optimal settings on unknown parameters in the algorithm has to be decided. Secondly when the complete algorithm is put to test on historical data to see if it serves the original purposes of the development project. As was mentioned in section 2.1, there are good reasons to use the same performance measure in these two phases. The final result may otherwise be sub optimal.

Evaluation of prediction performance is an important, difficult, and often overlooked stage in the development of prediction algorithms for financial data. In the case of an artificial trader based on the "Trading Rule Approach", one problem is that the produced trades are comparatively few, which gives a somewhat weak statistical basis for our performance measures. Apart from this we have the problem with overtraining. By tuning the parameters in the artificial trader to maximize the performance on historical data, we always run a risk of fitting the algorithm too closely to the data set. Even if the trader's behavior is a rather "small" model (i.e. it has relatively few degrees of freedom), the problem must not be overlooked. Illustrative examples can be found in [7].

A third problem is the fact that the process is non-stationary, i.e. the performance of a trading behavior varies over time due to varying global conditions affecting the stock market as a whole. One strategy may work fine in a trending market, while another works best in a non-trending market.

Before starting to optimize the behavior of our artificial trader we must decide three things:

- What performance measure is relevant? Several options exist.

- What is the best way to measure the selected performance measure?
- With what do we compare our artificial trader's performance when we say that it is good?
We need a good benchmark.

We start by investigating alternative performance measures.

3.1 Performance measures

3.1.1 Total Profit.

An intuitively appealing and very common measure is the total wealth achieved by the trader when simulating trading over the available training data period. The wealth is often presented as a function of time in a so called *equity diagram*. For comparison, a global stock index is often presented in the same diagram. The curves are scaled so the leftmost point has a wealth of 1 for both the trader and the index. The values for other points along the date axis can be then interpreted as wealth relative to the one at the starting point. A value 2.10 will mean, for example, that the initial capital has increased to 210% of the start capital. Therefore, the final value on the very right of the diagram will be the net result after completed trading for the entire time period. One example of an equity curve can be seen on the top diagram in figure 1. The major drawback of this method is that the trades in the beginning of the time period affect the end result more than the trades in the end of the time period. This is a consequence of the cumulative nature of the simulation. The profits in the beginning of the time period are being reinvested and therefore will appear "several times" in the total wealth resulting from trading during the entire time period. This is a nice effect if you want to show off your trading and stun people with an exponential increase in wealth. Its commonly used when the financial institutes present results for their mutual funds over the years. However, the index curve is often not included in these cases since they would reveal that the mutual fund actually hasn't performed any better than the market in general. One has to realize that a constant economical growth of say 5% over a period of 10 years gives rise to a seemingly impressive exponential wealth curve for the average stock. With or without clever trading systems.

3.1.2 Average Profit Per Day

To get around the cumulative effect in the total-profit measure, one could compute the average profit per day (percentage) and use that as a measure of the trading performance. By this we lose the possibility to evaluate the trader's behavior during the varying global market conditions that arise during any longer simulation period.

3.1.3 Profit Per Year.

A compromise between using the total profit and the average profit per day would be offered by computing the annual profit achieved by applying the trading algorithm. The mean annual profit could be used as total measure for the entire time period. However, it's also very important to pay attention to the performance in each individual year.

3.1.4 Statistical Significance

As discussed in section 2.2, the profit from a Trading Rule based system has a very weak statistical significance if the number of trades produced during the simulation is too low. The number of trades are therefore of central importance and should be presented along with the trading results.

3.2 What is the best way to measure the selected performance measure?

Its important to realize that the profit measure we select for optimization and evaluation is a stochastic variable which has not got any single "value" that can be measured. Instead it has a

probability distribution which can be described by its statistical moments such as the *expectation value* and the *variance*. What we get when we compute e.g. the annual profit for a year, is merely *one* sample from this distribution (assuming it is stationary over the years). The task of "maximizing" the profit by changing the parameters in the model now becomes less well defined. Obviously, the common way of using the *mean* value is just *one* option. Other choices could be the lower limit of a confidence interval, the *median value* or the *Sharpe Ratio*. The issue is further discussed in section 8 where the profit is examined as a function of various parameters in the model.

3.3 Benchmarks

Apart from a performance measure, we also need something to which to compare the performance. What is good and bad performance depends on the alternatives. The natural benchmark is to use some kind of stock index, which is also the most common method used by professional broker companies. A broker who outperforms the index gets bonuses and the one who doesn't might risk his job. Mutual funds in particular are governed by these ways of thinking. Even if the individual traders won't lose their jobs, the managers know that their funds will lose customers if they start doing worse than the global stock index. This is one of the reasons that so many mutual funds acting on one particular market perform almost identically.

For our purposes however, it's a reasonable benchmark since it compares the performance to a very available alternative: that of buying a mutual fund instead of doing the trading ourselves.

The following two kinds of indexes have been tried:

- The Swedish index *Generalindex*.
- An artificial index with equal parts for all the stocks available for the artificial trader.

No significant differences affecting the performance evaluation have been found between the two kinds of indexes and the official *Generalindex* has therefore been chosen as benchmark.

3.4 Conclusions Regarding Performance Evaluation

The artificial trader's result is presented as annual profits together with the increase in index. The mean difference between these two figures constitutes the net performance for the trader. The performance is displayed in both tabular and graphical formats as shown in the table part of figure 3 and in figure 1. The second last line of the table presents the annual profit above the index. The mean value of this entity is presented in the second rightmost column and represents a one-figure performance measure. However, the individual figures for each year should also be considered. The question whether the optimization and selection of trading algorithms should be based on the mean value or, for example, on the worst value for the investigated years, remains open. The Sharpe Ratio gives a risk dependent measure and is discussed further in section 8.

4 Design of the Artificial Trader

In this section we discuss the design and implementation of the ASTA system. The task of the Artificial Trader is to act on an artificial market with a large number of available stocks that vary in prices over time. The Trader has to execute the trading rule $T(t)$ at every time step and decide whether to buy or sell stocks. The sole aim of the Trader is to produce as high a profit as possible. Various aspects of the calculation of the performance were discussed in the previous section 4.1. The presentation and evaluation of the trading results are a major part of the ASTA system.




Figure 1: Presentation of performance for a (successful) trading strategy

4.1 Performance Evaluation

The result of the artificial trader is presented as annual profits together with the increase in index. The mean difference between these two figures constitutes the net performance for the trader. The performance is displayed in both tabular and graphical formats as shown in the table part of Figure 4 and in Figure 5. The second last line of the table presents the annual profit above the index. The mean value of this entity is presented in the second rightmost column and represents a one-figure performance measure. However, the individual figures for each year should also be considered.

We now turn to the general architecture of the developed system.

4.2 Basic Architecture

The architecture of ASTA is based on an object-oriented approach with two major objects; the Market and the Trader. The two objects have a number of attributes and operations that can be applied on the objects. The basic components and their relations are presented in Figure 2.

4.2.1 The Market Object

The Market Object consists essentially of the total number of stocks that should participate in the trading simulation. A stock is defined by four time series; **Close**, **High**, **Low** and **Volume**. Each of these time series has a numeric value (or NaN in the case of a unavailable value) for each date in the time period of interest. The basic operation on the Market Object is the simulation of changing prices as the date moves from start date to end date. The attribute **T** is updated by the **Step in time** operation applied to the Market Object.

4.2.2 The Trader Object

The Trader Object is more complex than the Market Object, as far as both attributes and allowed operations are concerned. The **Portfolio** attribute keeps track of the possession of stocks during the simulation. The **Cash** attribute is initialized to a certain value and thereafter modified automatically as stocks are bought and sold. The **Buy rule** and **Sell rule** are the main

attributes that affect the behavior of the Trader. They are expressed in a high-level language and may include calls to a large number of predefined Matlab functions that access the stock data in the Market Object. User-defined functions can also be called directly. The values of the **Buy rule** and **Sell rule** attributes are set interactively to enable fast experimentation when developing trading algorithms.

The basic operations on the Trader Object are **Buy_recommendations** that evaluates the **Buy rule** and **Sell_recommendations** that evaluates the **Sell rule**. The result is vectors with buy and sell recommendations for all relevant stocks. These recommendations are then carried out by the **Buy** and **Sell** operation. The behavior of the Trader is modified by a number of **Other parameters**, e.g. minimum and maximum values for one individual trade.

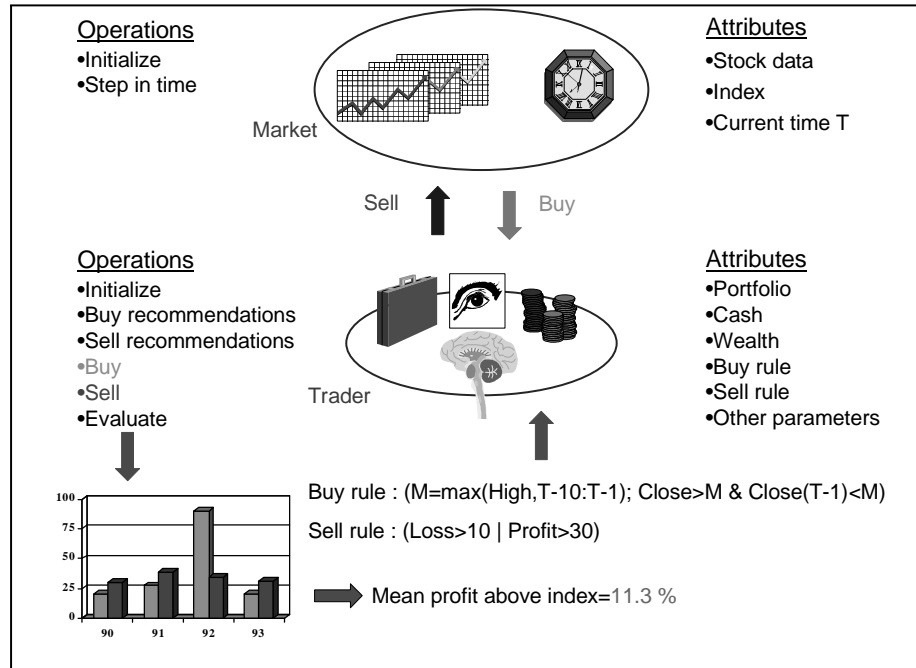


Figure 2: Basic Components of the ASTA System

4.2.3 Other Parts of the System

The Market and Trader Objects have to be controlled by a support system that takes care of the following "meta" operations:

- Simulation.

The **Step in time** operation has to be applied to the **Market** Object in a loop for the selected time period. For each time step, the **Trader** Object should also be activated. The following pseudo code describes the full ASTA system:

```

Trader.Initialize
Market.Initialize
loop until Market.T ≥ EndDate
    s = Trader.Sell_Recommendations
    Trader.Sell(s) % Sell all stocks of type s
    s = Trader.Buy_Recommendations
    n = T.available_cash / length(s)
    Trader.Buy(s, n) % Buy n stocks of type s
    Market.Step in time
end loop
Trader.Evaluate

```

- User interface.

The end user assigns values to parameters such as the **Buy rule** and **Sell rule** of the Trader and the desired time period for simulation. After the simulation the computed performance of suggested trades is presented and optionally printed to different media.

The "silent mode" makes it possible to use the system as an objective function in a parameter optimization. The Artificial Trader is then called as a standard Matlab function from the optimization program code, returning a function value for each set of input parameters.

ASTA is available in two versions:

- As a Matlab function, `fasta.m`, that takes a **Buy rule** and **Sell rule** as in parameters and returns the performance for a selected time period.
- As a Windows application to be run under Matlab. The **Buy rule** and **Sell rule** and all other settings are controlled interactively and the results are presented on the screen.

4.3 Predefined ASTA Functions

ASTA has a large number of predefined functions that make it possible to express compound trading rules interactively (as **Buy rules** and **Sell rules**). They also provide the developer of new algorithms with basic database access functions as well as some useful high-level functions.

A complete description of the predefined functions in ASTA can be found in section 6.4.

4.4 Other Design Issues

We conclude the general description of the design with some specific issues that must to be considered in the design of an artificial trader.

4.4.1 When Is Data Available?

In a real trading situation on day T , the Close, High, and Low for day T are not available. However, it is common that prediction algorithms assume that this data is available and can be included in the prediction of tomorrow's prices. Furthermore, it is often assumed that the Buy and Sell recommendations from the trading system can be executed using the close prices for day T . Of course, in reality this is hard to achieve, since the stock market is closed by the time the data for the current day becomes available, not to mention it having been transferred to a computerized trading system. ASTA can be configured to execute the simulated trades using prices from either day T or day $T+1$ for buy and sell. The data for day T is assumed to be available *for analysis* on day T .

4.4.2 How Many Stocks to Buy?

The trading rule $T(s)$ does not contain any guide to how many stocks or how big a portion of the cash should be invested in the particular stock that gets a *buy* signal. It is obvious that this affects the performance of a trading system and it is also clear that a *risk* estimate coupled with the trading signal could be of help. The present version of **ASTA** does not offer any sophisticated procedures for the selection of the number of stocks to buy. The agent simply divides all available cash equally among the stocks that get a *buy* signal (taking into account the upper and lower limits for one individual trade). When a stock gets a *sell* signal, all shares are sold.

4.4.3 Reinvesting Money

The Artificial Trader buys and sells stocks and hopefully increases the wealth during the simulated time period. In such a strategy there will be a cumulative effect, since old profits will be reinvested in future trades.

4.4.4 Interest on Cash

The money not invested in stocks is assumed to yield interest in a bank account. The level of interest varies over time and in our test database follows the daily Swedish three-month-bond interest rate (configurable).

4.4.5 Unused buy and sell signals

One problem of the basic approach of evaluation by simulated trading is that a large portion of the buy signals will be neglected due to lack of money for the artificial trader. Even sell signals will be neglected simply because the stock that issues the signal is not in the trader's portfolio. Tests show that more than 90% of the signals often are ignored in this way. This seriously affects the statistical basis of a performance analysis for the training algorithm. A possible way to attack the problem is a top level loop with multiple runs and randomization of entities like:

- The starting date of trading.
- The acceptance of buy signals.
- The set of available stocks for trading.

The present version of ASTA does not include any of the suggested improvements. Since the running time of one simulation of a 10-year trading takes in the order of 0.5-3 minutes on a 266 MHz Pentium PC, the problem needs proper analysis before any further development along these lines is conducted.

5 Using ASTA

ASTA is available in two versions:

- As a Matlab function, `fasta.m`, that takes a **Buy rule** and **Sell rule** as in parameters and returns the performance for a selected time period.
- As a windows application to be run under Matlab. The **Buy rule** and **Sell rule** and all other settings are controlled interactively and the results are presented on the screen.

5.1 fasta

The purpose of the "batch" version of the program is to provide an objective function that can be used for parameter tuning or automated generation of trading rules. This report does not describe that part of the system more than by showing the following example of how it can be used:

Example 1 :

```
% load the workspace with stock data:
fasta('astasxg')
% Simulate trading with specified Buy and Sell rules for 1990-1995.
% The performance is returned in the p variable.
p = fasta('Close(T)>Close(T-1)', 'Profit>20 | Loss>10', [90 95])
```

5.2 wasta

Asta can also be run as an interactive windows application in the Matlab system. In this section examples from the interactive version are demonstrated.

To start ASTA, run the **wasta.m** function from the matlab environment. The following screen should be displayed:

ASTA

From date: 87 To date: 97 Courtage (%): 0.15 Min Courtage: 90 Min buy (% per trade): 5 Max buy (% per trade): 20 Initial Cash: 100000

Buy rule: `Close(T) > High(T-1)`
 Sell rule: `Loss>10 | Profit>20`

Run Save graph Optimize

Market: 32 stocks. Dates: 820104-980409 (4071 days)
 Load: astasxg
 Generate: SIG
 Save:

Dump Trades: ☐ To window ☒ Today's ☐ Tomorrow's
☐ To file ☐ Diagram

Performance:

Annual profits:	87	88	89	90	91	92	93	94	95	96	97	Mean	Total
Strategy profit	-5.9	26.3	-9.5	-34.3	1.0	-5.7	33.6	13.9	39.8	36.6	30.2	11.4	154.4
Index profit	-7.9	51.9	22.9	-29.7	5.4	-0.0	52.1	4.6	18.3	38.2	23.8	16.3	310.3
Difference profit	2.0	-25.6	-32.4	-4.6	-4.4	-5.7	-18.6	9.3	21.5	-1.6	6.4	-4.9	-155.8
Number of trades	81	43	50	85	65	69	92	49	37	50	59	62	680

Switch to graph window for performance plots

Help End

Figure 3: Screen layout for the windows version of ASTA

In this picture, the user can set up parameters for a simulated trading. The stocks of interest are picked from a large database in Metastock format. Stocks from the Swedish stock market have been used in the presented runs. 32 major stocks with active trading for the years 1987-1997 have been selected. Other databases could easily be interfaced to the system.

The most interesting items are the lines "Buy rule" and "Sell Rule". This is the place where the trading algorithm is decided. The rules follow Matlab syntax and can include any of the large number of predefined functions or the users own functions with new algorithms. The simulation of trading starts by clicking on the "Run" button. The results are presented in tabular form and in graphs such as figure ??.

The command button "Optimize" initiates a whole series of simulations with different values on symbolic parameters in the Buy and Sell rules. The name of the parameters is given in the "Parameter" text box and the values to be included in the multiple simulations are given in the "Values" text box. The results are presented in four graphs which are also written as eps files. Examples can be found in the section 8.

6 Developing Trading Algorithms with ASTA

One of the key goals of using ASTA has been to provide a solid and easy-to-use basis for development of trading algorithms. The "Trading rule" approach has been taken as described in Section 2.2. The algorithm should be formulated in a day-by-day structure, where buy and sell recommendations are produced daily by the algorithm. The decisions should be based on past stock data only. Fundamental data about the companies cannot be used in the present version of the system. It should also be emphasized that the "Trading rule" approach implies a decision-making system and not a modeling system. However, implemented algorithms may use modeling "internally" and base the buy and sell recommendations upon it.

The following goals and guidelines have been kept in mind while designing the development interface in ASTA:

- Fool proofness

An "off-by-one-error" can seldom result in more drastic consequences than when developing

prediction algorithms. Looking into the future, which is normally pretty hard to do, is in the case of prediction algorithms *surprisingly* easy to do, even unintentionally. In the case of indexing time series vectors however, it should not come as a surprise that finding an error in code segments, like the following one, is indeed difficult.

$$i = k - h + 1$$

$$x(m) = y(k + 2 * (j - j) - 1)...$$

Therefore, the access functions implemented in **ASTA** have been designed with a check for arguments that peep into the future. The global variable **T** is automatically updated by the simulation mechanism to point at the "current day". Attempts to address the stock time series before this point will automatically issue an error message. Example:

$y = \text{Close}(T - 4 : T)$ assigns the close prices for the current day and the 4 previous days to the variable y .

$y = \text{Close}(T : T + 1)$ issues an error message when executed during the simulation.

The same check routine has been implemented in the higher level routines as well, thereby issuing clearer error messages pointing to the first place, where an illegal time reference is found.

- Ease of use

Existing ASTA functions can be used interactively to create **Buy rules** and **Sell rules**. When the user has implemented a new trading algorithm the Matlab function can also be used in the **Buy rules** and **Sell rules**. It can easily be tried out interactively with different settings on parameters etc.

- High level functionality

Predefined ASTA functions can perform operations such as moving averages and detection of crossings between time series. Other predefined functions compute derived entities from the stock data time series; gaussian volume, volatility, trend and ranks. These functions can be included in the user-defined algorithms and will simplify the coding considerably.

ASTA is equipped with a large number of predefined functions to be used interactively in **Buy rules** and **Sell rules**. They may also be used in user-defined functions. Before going into details of the available functions, some basic concepts will be described.

6.1 Global Variables

The variables listed in table 1 are dynamically set by the system and should normally not be changed by any user routines.

Name	Type	Size	Description
T	numeric	1	Row pointer to "current" day
Stocks	numeric	$1 \times N$	Column pointers to allowed stocks. Stocks is set to all stocks in market when evaluating a Buy rule Stocks is set to all stocks in portfolio when evaluating a Sell rule
Market	struct	1	The Market object
Trader	struct	1	The Trader object
ST	struct	1	Parameters controlling the Trader's behavior

Table 1: Global variables in ASTA

The variable **T** is central both when defining **Buy rules** and **Sell rules** and also when writing new functions for prediction. The idea is that the user never should have to worry about the time aspect, since it is automatically taken care of by the simulation framework. By using **T** as index, the "current" day will always be referenced. By using **T-1**, yesterday will be referenced etc. Attempts to use **T+1** will issue an error message, thus prohibiting the program from peeping into the future.

It is seldom necessary to use the variable **Stocks** when writing user-defined functions. However, it should be made clear that the variable is set to all stocks in the market when the system looks for stocks to buy (i.e. when the **Buy rule** is evaluated) and to all stocks in the Trader's portfolio when the system tries to find stocks to sell (i.e. when the **Sell rule** is evaluated). The function call **Close(T)** for example, returns a row vector with the close values for all stocks in the **Stocks** variable for the day **T** (the "current" day in the simulation).

6.2 Time Series Matrices

Most time series data within the ASTA system are enclosed in objects denoted *Time Series Matrices* and are abbreviated *TSM* in this documentation. They are ordinary 2-dimensional Matlab matrices where each row represents one date and each column represents one stock. The actual dates and actual stocks in a particular matrix vary, but the last row most often corresponds to the "current" day **T**. The columns normally correspond to the stocks in global **Stocks**.

The available data is the stock time series enclosed in the **Market** Object; Close, High, Low and Volume. These are the only features available when creating **Buy rules** and **Sell rules**, and also when creating new prediction algorithms. **Close**, **High**, **Low**, and **Volume** are also the names of the access functions used throughout the system.

Example 2 : *The call **Close(T-5:T)** returns a 6 row matrix where each column is assigned the close value for a stock in the global **Stocks** variable. Row 1 corresponds to day **T-5**, row 2 to day **T-4**, down to row 6, which corresponds to day **T**.*

A whole range of functions that accept TSM:s as arguments and/or generate it as output are included in the system. This design principle turns out to be convenient, since the matrix operations within Matlab can be used and makes the code very readable and effective. The Matlab function OHIGH in the following example detects Close values that cross a *k* day maximum of High from below.

Example 3 :

```
function b = OHIGH(k)
    C = Close(T);
    Cm1 = Close(T-1);
    M = max( High( T-k:T-1 )); % The call to High returns a k
                                % rows long Time Series Matrix
    b = C>M & Cm1<=M;
    return
```

The function returns a binary row vector with a "1" in those columns where the corresponding stock fulfilled the condition. It is an example of the data type: *Indicator Vector*.

6.3 Indicator Vectors

An Indicator Vector is a single-row TSM. Sell and Buy rules should evaluate to binary Indicator Vectors that are used directly as sell and buy recommendations. Other Indicator Vectors such as the return value from the function calls **Profit** and **Trend1** contain real numbers as values.

In the previous section the function OHIGH returned an Indicator Vector with binary values, "0" or "1". A number of predefined functions return Indicator Vectors, and a user-defined trading algorithm or component should normally return a vector of this type. This enables such compositions of **Buy** rules as:

$$\text{OHIGH}(10) \ \& \ (\ \text{Trend1}>1 \ | \ \text{Trend3}>0.5 \) \quad (10)$$

The **Buy** rule evaluates to a binary Indicator Vector with "1" for stocks where OHIGH(10) returned "1" and either the short trend is greater than 1%/day or the long trend is greater than 0.5%/day. Functions **Trend1** and **Trend3** are described below.

6.4 Predefined ASTA Functions

This section provides a list of all predefined functions that have been included to make it possible to express compound trading rules interactively (as **Buy rules** and **Sell rules**) and also to provide the developer of new algorithms with basic access functions as well as some useful high level functions.

The functions are divided into four categories:

1. Feature Functions.

Feature Functions are primarily access functions to stock prices (**Close**, **High**, **Low**) and traded **Volume**. Numerous functions for commonly used derived entities, such as trend, gaussian volume, and volatility, are provided as well. A Feature Function always returns a TSM.

2. Indicator Functions.

Indicator functions return an Indicator Vector most often used in the logical expressions that define the **Buy rules** and **Sell rules**. Examples are standard technical indicators, such as MACD and Stochastics. Other Indicator Functions implement a dynamic stop-loss handling and achieved profit or loss since a stock was bought.

3. Operator Functions.

Operator functions perform a computation on one or several TSM:s and return another one. Examples are computation of moving averages, minimum values in time windows etc.

4. Utility Functions.

Some functions are support functions that may be of use primarily when debugging new algorithms. Functions that extract stock names and actual dates are in this category.

IMPORTANT:

All functions of type 1, 2, and 3 return a TSM. Each column represents one stock in the global variable **Stocks**. **Stocks** is automatically set to all stocks in the market when the system looks for stocks to buy (i.e. when the **Buy rule** is evaluated) and it is set to all stocks in the Trader's portfolio when the system tries to find stocks to sell (i.e. when the **Sell rule** is evaluated). It is seldom necessary to use the **Stocks** variable explicitly in **Buy rules** or **Sell rules** or when writing user-defined functions.

6.5 General Parameters

This section describes some parameters, which are common to many of the predefined functions.

6.5.1 days

This parameter tells for which day or days the function should return values. In the case of **days** being a vector, the result is a TSM with the same number of rows as elements in the **days** parameter. Each row contains values as if computed on the corresponding day.

The default value for the **days** parameter is the global **T**, i.e. the current day.

Example 4 : *$Close(T)$ returns a single row TSM with close prices for all the stocks in the global variable **Stocks**.*

Example 5 : *$Trend3(T-9:T)$ returns a ten-row TSM with 5-day trends for all the stocks in the global variable **Stocks**.*

To simplify the description of functions below, the **days** parameter is not covered. The **Description** field in the tables refers to *each* row.

6.5.2 plot

Many functions have the optional **plot** argument. By setting this parameter to the global constant **PLOT**, the function is plotted versus time for all the selected stocks selected in a row. The exact contents of the plot depends on the function. User-defined functions can easily incorporate the plotting functionality, which is extremely useful during development and debugging.

Only one function at a time can have the **plot** parameter set to a non-zero value (the global constant **PLOT** equals to 1). The default value for the **plot** parameter is 0.

6.6 Feature Functions

Feature Functions are primarily access functions for obtaining stock prices (**Close**, **High**, **Low**) and traded **Volume**. Derived entities such as ranks and trends are also provided.

All the Feature Functions return a TSM with one row for each day in the **days** parameter and one column for each stock in the global **Stocks** variable. Table 2 lists all predefined Feature functions in **ASTA**. The **plot** argument is described in section 6.5.2.

Function	Parameters	Description
Close	days, plot	The close value
High	days, plot	The highest traded price during day day
Low	days, plot	The lowest traded price during the day
Volume	days, plot	The traded volume value
Rank1	days, plot	The 1-day rank value
Rank2	days, plot	The 2-day rank value
Rank3	days, plot	The 5-day rank value
Rank4	days, plot	The 20-day rank value
Trend1	days, plot	The 1-day trend (%/day)
Trend2	days, plot	The 2-day trend (%/day)
Trend3	days, plot	The 5-day trend (%/day)
Trend4	days, plot	The 20-day trend (%/day)
Gvol1	days, plot	The 2-day gaussian volume
Gvol2	days, plot	The 5-day gaussian volume
Gvol3	days, plot	The 10-day gaussian volume
Gvol4	days, plot	The 20-day gaussian volume
Volatility1	days, plot	The 2-day relative volatility
Volatility2	days, plot	The 5-day relative volatility
Volatility3	days, plot	The 10-day relative volatility
Volatility4	days, plot	The 20-day relative volatility

Table 2: Predefined feature functions

Some special concepts need to be further explained:

6.6.1 k-step Return

The k -step return $R_k^m(t)$ of a stock m with a price time series $Close^m(t)$ is defined as

$$R_k^m(t) = 100 \cdot \frac{Close^m(t) - Close^m(t - k)}{Close^m(t - k)}. \quad (11)$$

By setting k at different numbers we get measures indicating how much the stock has increased since its value k days ago. The Return values are not pre computed in **ASTA**. The related **Trend** functions described in the next section can often be used instead.

6.6.2 k-step Trend

The k -step trend $T_k(t)$ of a stock m is defined as

$$T_k^m(t) = \frac{1}{k} \cdot R_k^m(t). \quad (12)$$

By setting k at different numbers we get measures indicating how much the stock has increased per day since its value k days ago.

The four functions **Trend1**, **Trend2**, **Trend3**, and **Trend4** can be used to access four return measures for $k = 1, 2, 5$ and 20 . They are computed and stored during market initialization to speed up the use during the simulation phase.

6.6.3 Rank

The k -step Ranks A_k^m for stocks $\{s_1, \dots, s_N\}$ are computed by ranking the N stocks by the k -step returns R_k^m . The ranking orders are then normalized so the stock with the lowest R_k gets rank -0.5 and the stock with the highest R_k gets rank 0.5 . The definition of the k -step Rank A_k^m for a stock s_m , belonging to a set of stocks. $\{s_1, \dots, s_N\}$ can therefore be written as

$$A_k^m(t) = \frac{\text{order}(R_k^m(t), \{R_k^1(t), \dots, R_k^N(t)\}) - 1}{N - 1} - 0.5 \quad (13)$$

where the *order* function returns the ranking order of the first argument in the second argument, which is an ordered list. R_k^m is the k -step return (definition 11) computed for stock m .

The Rank for a stock is a measure seldom or never seen as input to Automatic Trading Systems. Since the ultimate goal is to "beat the market", i.e. to do better than the average stock, it seems reasonable to have them at least available as potential inputs. The rank concept is further analyzed in xxx.

The four functions **Rank1**, **Rank2**, **Rank3**, and **Rank4** can be used to access four rank measures for $k = 1, 2, 5$ and 20 . They are computed and stored during market initialization to speed up the use during the simulation phase.

6.6.4 Gaussian Volume

The gaussian volume $V_n(t)$ is defined as

$$V_n(t) = (V(t) - m_V(t)) / \sigma_V(t) \quad (14)$$

where $m_V(t)$ and $\sigma_V(t)$ are computed in a running window of length n . I.e.:

$$m_V(t) = \frac{1}{n} \sum_{i=1}^n V(t-i) \text{ and} \quad (15)$$

$$\sigma_V = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (V(t-i) - m_V(t))^2}. \quad (16)$$

V_n expresses by how many standard deviations the volume differs from its running mean. Therefore, it can be used as an indication of "abnormal" values caused by a sudden change in market interest for a particular stock. That's the reason why the sums in definition 15 and 16 don't include today's volume, i.e. $V(t)$. The four functions **Gvol1**, **Gvol2**, **Gvol3**, and **Gvol4** can be used to access four Gaussian volumes V_n for $n = 2, 5, 10$ and 20 . They are computed and stored during market initialization to speed up the use during the simulation phase.

6.6.5 Relative Volatility

The Relative volatility $W_n(t)$ is here defined as

$$W_n(t) = 100 \cdot \frac{\sigma_C(t)}{m_C(t)} \quad (17)$$

where $m_C(t)$ and $\sigma_C(t)$ are computed on the time series `Close` in a running window of length n . I.e.:

$$m_C(t) = \frac{1}{n} \sum_{i=0}^{n-1} \text{Close}(t-i) \text{ and} \quad (18)$$

$$\sigma_C(t) = \sqrt{\frac{1}{n-1} \sum_{i=0}^{n-1} (\text{Close}(t-i) - m_C(t))^2}. \quad (19)$$

Definition 17 is recognized as the "coefficient of variation" or "relative standard deviation" in statistics. Relative volatility W_n is preferred for ordinary volatility estimation by the standard deviation, because it provides a measure that can be compared over longer time periods and also between stocks with totally different price levels.

The four functions **Volatility1**, **Volatility2**, **Volatility3**, and **Volatility4** can be used to access four volatilities W_n for $n = 2, 5, 10$ and 20 . They are computed and stored during market initialization to speed up the use during the simulation phase.

6.7 Indicator Functions

These functions return an Indicator Vector, i.e. a TSM with one row and one column for each stock in the global variable **Stocks**. Each column can be used in logical expressions to compose **Buy rules** and **Sell rules**. They can also be used with user-defined functions to create compound trading rules. Table 3 lists all predefined Indicator functions in **ASTA**. The **plot** argument is described in section 6.5.2.

Function name	Parameters	Description
Profit	-	% increase in Close since buy of stock
Loss	-	% decrease in Close since buy of stock
Stoploss	S1, S2, P1	Implements a dynamic stop-loss function: Stop loss is initially set to S1 If Profit > P1 , change the stop loss to S2 Stoploss returns "1" iff the Loss is less than the dynamic set stop loss. This function keeps track of each stock separately.
Underlow	L	"1" iff Close crosses a L day min of Low from above
Overhigh	L,plot	"1" iff Close crosses a L day max of High from below and High(T-1) is less than the same max
Days	dys	"1" iff day number is in vector dys (1:Monday...7:Sunday)
Stoch	K,Ks,S,plot	Implementation of the technical indicator Stochastics.
Macd	K,D,S,plot	Implementation of the technical indicator MACD

Table 3: Predefined indicator functions

Some Indicator Functions require further explanation:

6.7.1 Profit and Loss

The system automatically keeps track of the change in price of each stock since it has been bought. The **Profit** is defined as

$$\mathbf{Profit}(t) = (Close(t) / \text{BuyPrice} - 1) * 100. \quad (20)$$

The **Loss** is sometimes a more convenient measure and is defined as

$$\mathbf{Loss}(t) = (1 - Close(t) / \text{BuyPrice}) * 100. \quad (21)$$

BuyPrice is the average buy price of each stock in a portfolio. **Profit** and **Loss** are not defined for stocks outside the portfolio.

The functions are of great use when defining trading rules. For example, a simple **Sell rule** may be defined as

$$\mathbf{Loss} > 10 \mid \mathbf{Profit} > 20 \quad (22)$$

The **Sell rule** evaluates to an Indicator Vector with "1" for stocks where either **Loss** returned a value less than 10 or **Profit** returned a value greater than 20. A more complex function based on the profit and loss concept is **Stoploss**.

6.7.2 Stoploss

An often mentioned "trading rule to live by" is that of never allowing a profitable trade to turn into a loss. This idea is implemented in the function **Stoploss**. Each stock gets a stop-loss parameter of its own. This parameter is initialized to the value **S1** when a stock is bought. It is changed to **S2** when the **Profit** function registers more than **P1%** profit. Typical value for the parameters **P1**, **S1** and **S2** are 20%, 10%, and -5% respectively. These values gives the following behavior during trading:

1. The stock will be sold immediately if the price drops by more than 20% below purchase price, thus cutting the maximum loss.
2. If the price ever rises by 10% above the purchase price, a consecutive drop in price to 5% above purchase price will issue a sell signal. In this way, a profitable trade will never turn into a loss.

6.8 Operator Functions

These functions perform an operation on one or several TSM:s and return another TSM with transformed values.

All Operator Functions return a TSM with one row for each day in the **days** parameter and one column for each stock in the global **Stocks** variable. The parameter **x** (same applies to **x1** and **x2**) can be either a TSM or a string with the name of a function returning a TSM.

The parameter **days**:

- If **x** is a TSM, the **days** parameter should be row numbers in the matrix. The default value for **days** is the last row in the matrix.
- If **x** is a string, the **days** parameter should be a vector with day numbers such as **T**, **T-1** etc. The default value for **days** is **T**.

The Operator Function performs its operation on each of the lines in the **x** matrix and returns a TSM with the result for each row. In the following description of Operator Functions, the **d** variable is implicitly assumed to run through the full range defined by the **days** parameter. Each value on **d** results in one row in the output TSM.

Table 4 lists all predefined Operator functions in **ASTA**.

Function name	Parameters	Description
Minn	x, L1, days, plot	The L1 -day-minimum value in the rows of x
Maxx	x, L1, days, plot	The L1 -day-maximum value in x
Mav	x, L1, days, plot	The L1 -day-moving av. for columns in x
Stdev	x, L1, days, plot	The L1 -day-st.deviation for columns in x
Mavx	x1,L1,x2,L2,A, days,plot	"1" iff Mav(x1,L1) crosses Mav(x2,L2) from below. The A argument (optional) is the min required angle for the crossing.

Table 4: Predefined operator functions

6.8.1 Minn

Computes the minimum value in a L1 days long window backwards.

If x is a string, Minn is defined as: $Minn = \min(x[d - L1 + 1 : d])$ where the square brackets denote function invocation of the function **x**.

Example 6 $Minn('Close', 4, T)$. Will generate a 1 row TSM: $\min(Close(T - 3 : T))$.

Example 7 $Minn('Close', 4, T - 1 : T)$. Will generate a 2 row TSM: $\begin{pmatrix} \min(Close(T - 4 : T - 1)) \\ \min(Close(T - 3 : T)) \end{pmatrix}$.

If x is a not a string, Minn is defined as: $Minn = \min(x(d - L1 + 1 : d, :))$.

Example 8 $Minn(Z, 4, 10)$ where Z is a 10 rows long TSM. Will generate a 1 row TSM: $\min(Z(7 : 10, :))$.

Example 9 $Minn(Z, 4, 9 : 10)$ where Z is a 10 rows long TSM. Will generate a 2 row long TSM: $\begin{pmatrix} \min(Z(10 - 4 + 1 : 10, :)) \\ \min(Z(10 - 4 + 1 : 10, :)) \end{pmatrix}$.

6.8.2 Maxx, Mav and Stdev

These functions have the same parameters and behavior as **Minn** above. Just substitute *min* for *max*, *mean*, and *std*.

6.8.3 Mavx

Crossings between moving averages of various signals are often used in the definition of traditional technical indicators. The **Mavx** function supplies a convenient way to implement the detection of such crossings. The optional argument **A** makes it possible to specify the minimum angle for the crossing. Crossings at lower angles than the specified one will not generate a logical "1" at that point. **Mavx** can be used to detect both crossings from below and above by changing the order of **L1** and **L2** arguments.

Example 10 $Mavx('Close', 2, 'Close', 20, 45)$ returns a binary Indicator Vector with "1" iff a 2 day moving average of **Close** crosses a 20 day moving average of **Close** from below in an intersection angle no less than 45 degrees. The computation is performed "today" since the **days** parameter defaults to **T**.

6.9 Utility Functions

These functions are support functions that may be of use primarily when debugging new algorithms. They can not be included in **Buy rules** and **Sell rules** since they don't return Time Series Matrices. Table 5 lists all predefined Indicator functions in **ASTA**.

Date	days	Vector with the real date (yymmdd) for days
Stockname	s	String with the name of stock s .

Table 5: Predefined utility functions

7 Examples

In this section one example from the Windows version of ASTA is presented. 32 major stocks with active trading from the Swedish stock market for the years 1987-1997 have been selected for analysis. The main ASTA screen is shown in Figure 4. The most interesting items are the lines "Buy rule" and "Sell Rule". This is the place where the trading algorithm is defined. The rules follow the Matlab syntax and can include the predefined functions or the users' own functions with new algorithms.

The example shows a test of the Stochastics indicator $Stoch[t, K, KS, D, Buylevel, Sellevel]$, common in technical trading, and here defined as

$$\begin{aligned}
 Stoch(t) &= mav(100 * (Close(t) - L) / (H - L), D) \text{ with} \\
 L &= mav(\min(Low(T - K : T)), KS) \text{ and} \\
 H &= mav(\max(High(T - K : T)), KS).
 \end{aligned} \tag{23}$$

The parameter K is the length of the window, and KS and D the length of the moving average function mav . A Buy signal is issued when $Stoch(t) > Buylevel$ and a Sell signal when $Stoch(t) < Sellevel$. The parameters $K, KS, D, Sellevel$ and $Buylevel$ control the performance of a trading strategy based on the indicator and will be subject to analysis in the next section.

In Figure 4, the "standard" values 30, 3, 3, 20, 80 are used for the parameters. The results shown in Figure 4 and Figure 5 are quite stunning, with an average annual profit of 53.4% compared to the 16.3% achieved by the index. It is noteworthy that the only negative result is for the year 1997, where the strategy only made 1.7% whereas the index increased by 23.8%.

The screenshot shows the ASTA software interface. At the top, the title "ASTA" is centered. Below it, there are input fields for "From date" (87), "To date" (97), "Courtage (%)" (0.15), "Min Courtage" (90), "Min buy (%)" (5), "Max buy (%)" (20), and "Initial Cash" (100000). The "Buy rule" is set to "Stoch(30,3,3,20,80)>0" and the "Sell rule" is "Stoch(30,3,3,20,80)<0". There are buttons for "Run", "Save graph", and "Optimize". Below these, the "Market" is set to "32 stocks. Dates: 820104-980409 (4071 days)". There are buttons for "Load" (astasxg), "Generate" (SIX), and "Save". To the right, there are checkboxes for "Dump Trades" (To window, To file, Diagram) and radio buttons for "Buy price" and "Sell price" (Today's, Tomorrow's). At the bottom, there is a "Performance:" section with a table of annual profits and a "Switch to graph window for performance plots" button.

	87	88	89	90	91	92	93	94	95	96	97	Mean	Total
Annual profits:													
Strategy profit	27.0	53.2	73.6	-9.3	9.6	32.4	313.6	22.5	25.1	38.3	1.7	53.4	3863.5
Index profit	-7.9	51.9	22.9	-29.7	5.4	-0.0	52.1	4.6	18.3	38.2	23.8	16.3	310.3
Difference profit	34.8	1.3	50.6	20.4	4.2	32.4	261.5	18.0	6.9	0.1	-22.1	37.1	3553.2
Number of trades	35	36	38	72	71	87	55	73	49	50	64	57	630

Figure 4: ASTA command window with Stochastics buy and sell rules




Figure 5: Performance of the Stochastics indicator.

8 Viewing the Trader as an Objective Function

The obvious wish to maximize the profit of a trading system can be tackled in two ways:

1. Parameterizing the **Buy rule** and **Sell rule**, i.e. introducing parameters within the rules.

Example:

Buy rule = $'Close(T) > Maxx('High', Nhigh, T - 1)'$

Sell rule = $'Loss > L \mid (Profit > P \ \& \ Close(T) < Close(T - 1))'$

The function *Maxx* determines the maximum time series value in the interval $[T - Nhigh, T - 1]$. It is now possible to optimize the profit *P* with respect to the parameters *Nhigh*, *L* and *P*.

2. Viewing the **Buy rule** and **Sell rule** as symbolic expressions. The optimization then turns into a search problem, most naturally implemented in a genetic framework or using Inductive Logic Programming.

An optimization of one of the parameters in the Stochastics indicator will serve as an example for approach 1. We set up an optimization with buy and sell rules according to:

Buy rule = $'Stoch(30, 3, 3, Sellevel, 80)'$

Sell rule = $'Stoch(30, 3, 3, Sellevel, 80)'$

The excess profit *P* can now be optimized with respect to *Sellevel*. The graphs in Figure 9 are automatically generated by the "Optimize" function in the ASTA system. The name of the parameter is given in the "Parameter" text box and the range in the "Values" text box.

It must be emphasized that optimization of the performance is multi-dimensional parameter estimation problem. The graphs presented show the profit *P* as a function of *one* of the involved parameters whereas the rest of the parameters are fixed. The main purpose is to illustrate the possibilities and problems involved, even in a one-dimensional optimization.

From the summary graph (top left) of the entire training period we can deduce that the highest profit is achieved for a *Sellevel* somewhere around 35. However, looking at the data at a higher resolution reveals a more complicated situation. In the bottom left diagram the same relation is plotted with one curve for each year in the training data set. It is clear that the mean profit is totally dominated by the results from *one* of the years (1993).

From these curves we can learn at least two important points:

1. The spread between individual years is very high.
2. The location of the maximum is not obvious.

This behavior of data is typical for most variables. The profit cannot be easily described as a function of measurable variables without introducing a dominant noise term in the function. Let us view the annual excess profit as a stochastic variable $P(\theta)$, where θ stands for one particular setting of the parameters that affect the profit. In the shown example, θ is the *Sellevel* parameter. $P(\theta)$ has been sampled once per year during the eleven years in the data set

$$\{P_1(\theta), P_2(\theta), P_3(\theta), P_4(\theta), P_5(\theta), P_6(\theta), P_7(\theta), P_8(\theta), P_9(\theta), P_{10}(\theta), P_{11}(\theta)\}. \quad (24)$$

Viewed this way, the task of tuning the parameter θ to find the "maximum" profit P is not well defined. P is a stochastic function and consequently has no "maximum". It has a probability distribution with an expectation value E and a variance V . It's important to realize that tuning θ in order to maximize $E[P(\theta)]$ is just one of the available options. Maximizing $E[P(\theta)]$ provides the highest *mean* performance. Another possibility is to maximize the lower limit of a confidence interval. Since the risk factor is always a major concern in investments, and since the spread between individual years obviously can be very high, this sounds like a promising idea. A lower limit P_{low} for a confidence interval could be defined as

$$P_{low} = E[P(\theta)] - \sqrt{V[P(\theta)]}, \quad (25)$$

where $V[P(\theta)]$ is the variance of the stochastic variable $P(\theta)$. Yet another possibility is to use the Sharpe Ratio SR , which expresses the excess return in units of its standard deviation as

$$SR = \frac{E[P(\theta)]}{\sqrt{V[P(\theta)]}}. \quad (26)$$

The Sharpe Ratio is normally used to *evaluate* the performance of a trading strategy (Sharpe [10, 11]). However, Choey and Weigend [4] suggest to use the Sharpe Ratio as an objective function in portfolio optimization and derive a learning algorithm for artificial neural networks. The Sharpe Ratio should be as high as possible for an optimal trading algorithm.

Due to the high noise level in the data we have added a modified Sharpe Ratio where the outliers have been removed. The largest and smallest $P_i(\theta)$ in expression 24 have been removed before taking the expected value and standard deviation for each θ . The unmodified Sharpe Ratio is shown in the top right diagram and the one with outliers removed in the mid right diagram. As we can see, the one with the outliers removed clearly reveals a maximum at around 20-25 followed by a clear decline. The unmodified Sharpe Ratio shows no such pattern.

9 Other ASTA features

By adding a PLOT parameter in the call to many of the predefined ASTA functions, illuminating graphs will be generated. The call to the previously used Stoch function may for example look like

$$\text{Stoch}(30,3,3,20,80,\text{PLOT})>0. \quad (27)$$

The enabled PLOT option will generate separate graphs with the signals from the Stoch function. In figure 7 the components of the Stochastics indicator are displayed. The two horizontal lines mark the buy level (80%) and the sell level (20%). When the "Oscillator %K" value passes these lines, a buy or sell signal is issued.

To analyze the generated trades in even more detail, the "dump trades to window" button has been checked. All generated trades will be dumped in ascii form in the Matlab command window. The result is presented in figure 8. By "Dump trades to file", a data file for post-processing with classification tools will be written to disk. It can be used to improve the performance for the used trading algorithm.

ASTA

From date To date Courtage (%) Min Courtage Min buy (%) per trade Max buy (%) per trade Initial Cash

93 93 0.15 90 5 20 100000

Buy rule

Sell rule

Run Save graph Optimize

Market: 32 stocks. Dates: 820104-980409 (4071 days)

Load

Generate

Save

Dump Trades Buy price: Sell price:

☒ To window ☒ Today's ☒ Today's

☐ To file ☐ Tomorrow's ☐ Tomorrow's

☐ Diagram

Performance:

Annual profits:			
	93	Mean	Total
Strategy profit	: 352.1	352.1	352.1
Index profit	: 52.1	52.1	52.1
Difference profit	: 300.0	300.0	300.0
Number of trades	: 41	41	41

Switch to graph window for performance plots

Help End

Figure 6: ASTA command window with Stochastics buy and sell rules. The PLOT option in the Buy rule generates separate plots with signals for each stock. The "Dump trades to window" check box generates a list of all trades in the Matlab command window.

10 Results and Further Development

The presented system provides a powerful tool for the development and evaluation of trading algorithms. Parameter settings can be tested and data screening can be easily performed interactively. As was mentioned in section 1.1, one of the reasons of the development of the ASTA system was to use it as an objective function when tuning model parameters and to find the general structure of trading rules, for example within a generic framework. The preliminary results show both possibilities and difficulties. The track can be examined considerably.

The dangers with "data snooping" got highlighted by breaking down the performance measures into shorter intervals. However, the inherent uncertainty resulting from the noisy processes involved calls for more computer-intensive simulation schemes in order to achieve statistically significant performance measures. Methods such as those described in section 4.4.5 will be implemented and evaluated.

11 Acknowledgments

I would like to thank Kenneth Holmström for his enthusiastic feedback and ever stimulating conversations. This work was sponsored in part by department of Mathematics and Physics at Mälardalen University, Västerås Sweden.




Figure 7: Stochastics buy and sell signals for the ABB stock. Upgoing bar denotes a buy signal and downgoing bar denotes a sell signal.

Trade	Date	No.	Stock	Price	Total Profit (%)	Wealth	Cash
1	930105	77	21 Skandia	85.27	6566.05		93343.95
2	930105	79	27 Sydkraft C	82.60	6525.22		86728.73
3	930105	1535	31 Allgon B	4.28	6574.72	99730.00	80064.01
4	930107	375	20 SHB A	26.34	9875.96		70119.50
5	930107	329	32 Nokia A	30.00	9870.00	99585.62	60159.50
6	930108	349	10 Hennes & Mauritz	28.20	9841.80		50260.08
7	930108	323	25 SSAB A	30.50	9851.50	99860.44	40318.58
8	930115	431	7 Avesta Sheffield	22.68	9776.63		30528.47
9	930115	351	20 SHB A	27.84	9771.77	100374.14	20666.71
10	930121	258	2 ABB A	38.40	9907.20		10702.32
11	930121	64	14 Kinnevik B	81.00	5184.00	NaN	5428.32
12	930127	54	21 Skandia	97.46	5262.68	NaN	84.48
13	930302	-64	14 Kinnevik B	79.00	-5056.00	-2.47	128010.13
14	930322	-131	21 Skandia	93.71	-12275.62	3.78	131216.32
15	930323	-431	7 Avesta Sheffield	23.16	-9980.37	2.08	NaN
16	930325	-323	25 SSAB A	33.75	-10901.25	10.66	132663.45
17	930329	4897	31 Allgon B	5.45	26675.18		132626.85
18	930406	2031	31 Allgon B	5.51	11186.34		132246.33
19	930519	-258	2 ABB A	43.70	-11274.60	13.80	164709.29
20	930525	113	22 Skanska B	97.50	11017.50	NaN	94.54
21	930608	-329	32 Nokia A	46.25	-15216.25	54.17	187009.69
22	930610	154	27 Sydkraft C	97.95	15083.58		184210.74
23	930621	-349	10 Hennes & Mauritz	38.00	-13262.00	34.75	13226.55
24	930621	132	22 Skanska B	99.00	13068.00		185878.28
25	930716	-233	27 Sydkraft C	93.56	-21799.15	0.88	197758.60
26	930719	308	7 Avesta Sheffield	34.97	10770.39		197540.39
27	930720	235	2 ABB A	46.00	10810.00		199141.88
28	930906	-235	2 ABB A	47.20	-11092.00	2.61	273799.16
29	930916	-308	7 Avesta Sheffield	36.86	-11352.70	5.41	272928.45
30	930920	227	9 Ericsson B	97.83	22207.73		272118.39
31	930921	-245	22 Skanska B	138.00	-33810.00	40.37	270682.53
32	930922	-726	20 SHB A	95.00	-68970.00	251.03	269365.98
33	930923	672	32 Nokia A	80.00	53760.00		270719.82
34	930927	562	10 Hennes & Mauritz	43.20	24278.40		282121.66
35	931004	495	2 ABB A	49.20	24354.00		309655.25
36	931108	-227	9 Ericsson B	102.87	-23350.81	5.15	NaN
37	931111	136	21 Skandia	170.55	23194.39		429195.43
38	931126	-672	32 Nokia A	100.50	-67536.00	25.62	399726.79
39	931203	764	30 Volvo B	88.20	67384.19		387109.87
40	931215	-136	21 Skandia	159.30	-21664.94	-6.59	434649.03

Figure 8: Dump of generated trades from Stochastics trading rules for the year 1993.

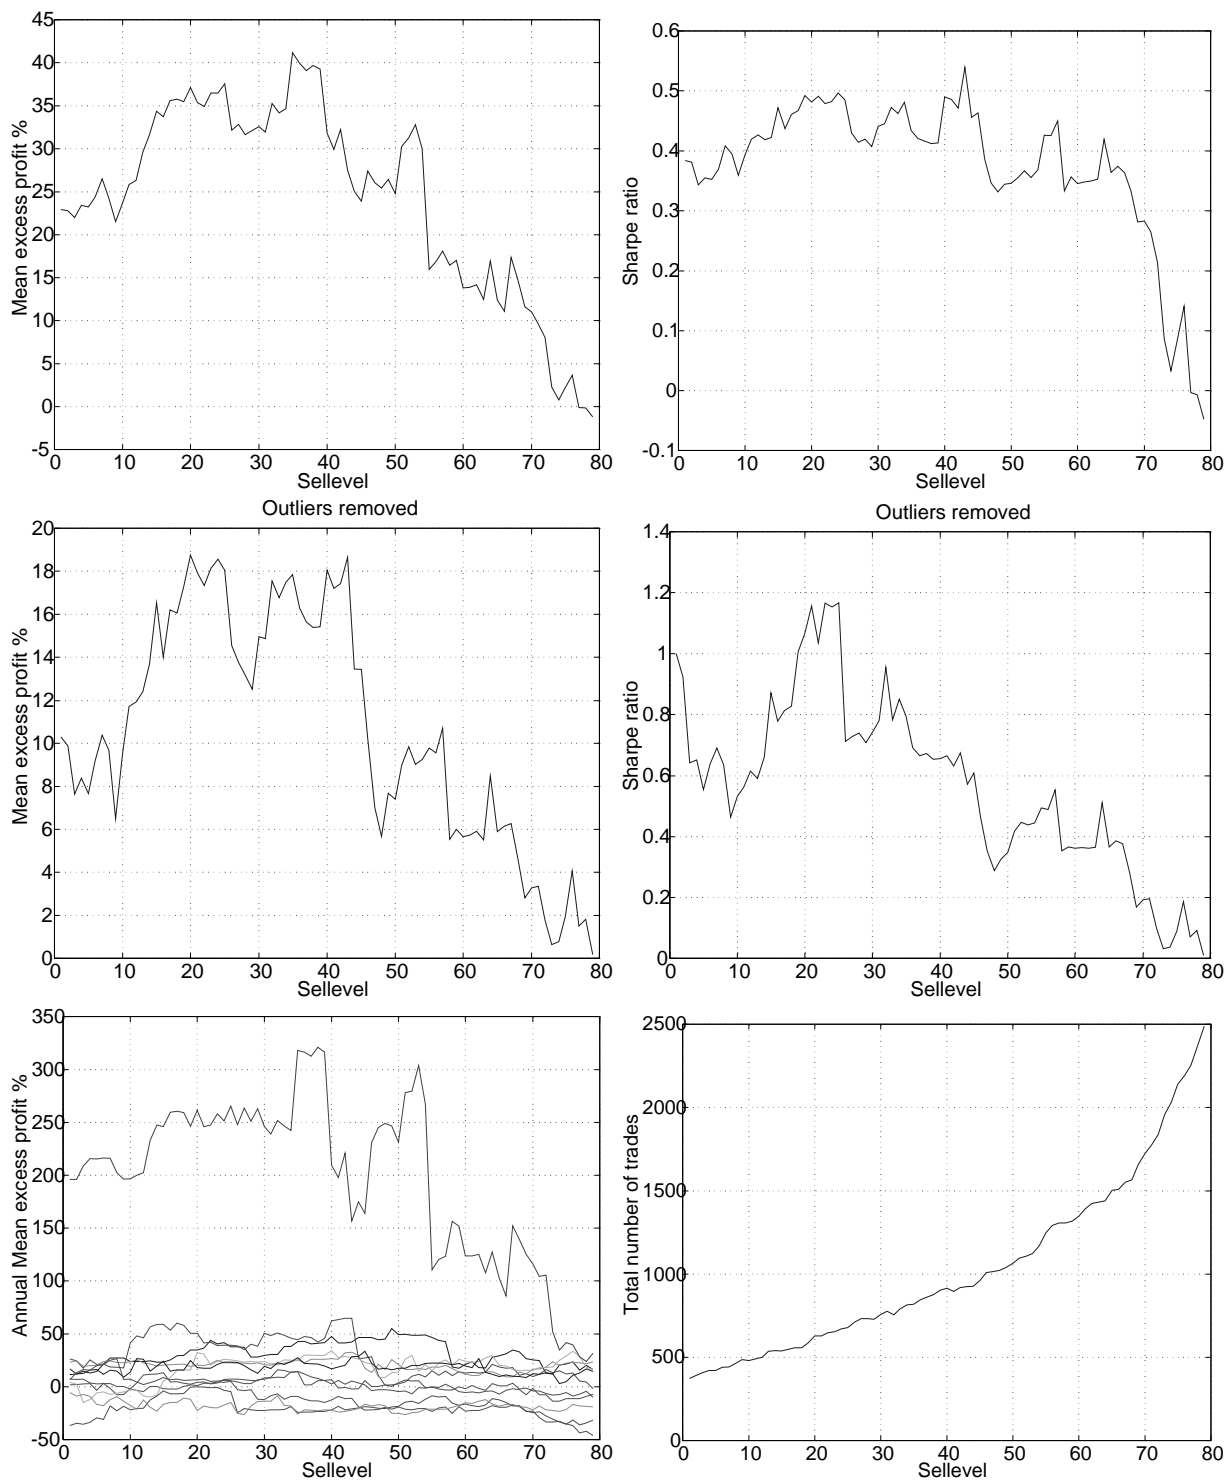


Figure 9: Trading results as a function of the Sellevel parameter. Stocks: SXG. Years:87-97.
 Buy rule: $\text{Stoch}(30,3,3,\text{Sellevel},80) > 0$. Sell rule: $\text{Stoch}(30,3,3,\text{Sellevel},80) < 0$

References

- [1] S. B. Achelis. *Technical Analysis from A to Z*. Irwin Professional Publishing, Chicago, 2nd edition, 1995.
- [2] A. Atiya. Design of time-variable stop losses and profit objectives using neural networks. In A. S. Weigend, Y. S. Abu-Mostafa, and A.-P. N. Refenes, editors, *Decision Technologies for Financial Engineering (Proceedings of the Fourth International Conference on Neural Networks in the Capital Markets, NNCM-96)*, pages 76–83, Singapore, 1997. World Scientific.
- [3] Y. Bengio. Training a neural network with a financial criterion rather than a prediction criterion. In A. S. Weigend, Y. S. Abu-Mostafa, and A.-P. N. Refenes, editors, *Decision Technologies for Financial Engineering (Proceedings of the Fourth International Conference on Neural Networks in the Capital Markets, NNCM-96)*, pages 36–48, Singapore, 1997. World Scientific.
- [4] M. Choey and A. S. Weigend. Nonlinear trading models through Sharpe Ratio maximization. In A. S. Weigend, Y. S. Abu-Mostafa, and A.-P. N. Refenes, editors, *Decision Technologies for Financial Engineering (Proceedings of the Fourth International Conference on Neural Networks in the Capital Markets, NNCM-96)*, pages 3–22, Singapore, 1997. World Scientific.
- [5] G. Leitch and J. Tanner. Economic forecast evaluation: Profit versus the conventional error measures. *The American Economic Review*, pages 580–590, 1991.
- [6] T. Hellström. Asta - a test bench and development tool for trading algorithms. Technical Report IMA-TOM-1998-11, 1998.
- [7] T. Hellström and K. Holmström. Predicting the Stock Market. Technical Report IMA-TOM-1997-07, Department of Mathematics and Physics, Mälardalen University, Sweden, 1997.
- [8] J. E. Moody. Shooting craps in search of an optimal strategy for training connectionist pattern classifiers. In J. E. Moody, S. J. Hanson, and R. P. Lippmann, editors, *Advances in Neural Information Processing Systems 4, Proceedings of the 1991 NIPS Conference*, pages 847–854, San Mateo, CA, 1992. Morgan Kaufmann Publishers.
- [9] J. E. Moody and L. Z. Wu. Optimization of trading systems and portfolios. In A. S. Weigend, Y. S. Abu-Mostafa, and A.-P. N. Refenes, editors, *Decision Technologies for Financial Engineering (Proceedings of the Fourth International Conference on Neural Networks in the Capital Markets, NNCM-96)*, pages 23–35, Singapore, 1997. World Scientific.
- [10] W. F. Sharpe. Mutual fund performance. *Journal of Business*, pages 119–138, January 1966.
- [11] W. F. Sharpe. The Sharpe Ratio. *Journal of Portfolio Management*, 21:49–58, 1994.