The TNT Financial Trading System: a midterm report

Luis Torgo

LIACC-FEP, University of Porto
R. de Ceuta, 118, 6., 4050-190 Porto, Portugal
email: ltorgo@liacc.up.pt,
Home page: http://www.liacc.up.pt/~ltorgo

Abstract. This paper briefly describes a commercial project whose main goal was to develop a system that should be able to autonomously trade in financial markets. The system should base its actions on prediction models obtained using data mining techniques. This paper describes the current results of the project after the implementation of the first version and after one full year of real time simulation of its trading activities. In terms of data mining the system can be regarded as a good case study of a full data mining project as it involves all major steps: data collection; data pre-processing; model construction; predictions; and decision making based on predictions. In terms of techniques the system uses a wide range of modelling approaches that together compete for the final trading decision at each time step (the system makes decisions every 15 minutes). Decisions are based on an analysis of the recent performance of the models and thus performing some kind of dynamic model averaging. The results of this first year of simulations are quite encouraging, though highlighting some problems of the approach.

1 Introduction

In this paper we briefly describe a commercial project carried out with the goal of developing an autonomous financial trading system. Financial trading is an almost ideal challenge for data mining systems. In effect, this is an application scenario where data abounds and where we aim at forecasting the evolution of a quite complex dynamic system that is influenced by too many factors for allowing any attempt to model it analytically. Although many claim that this is an unpredictable system (usually known as the Efficient Markets Hypothesis), several researchers have found windows of inefficiency that if correctly anticipated can allow for profit to be achieved. Other researchers have also found out interesting herding behaviors of traders that in the presence of similar market features tend to behave in the same way. This lays ground for a model to be successful in identifying these regularities and thus being able to anticipate them.

The main motivations for this project were the facts that: it is now possible to collect an incredible amount of information on financial markets, most of the times with no cost at all; recent developments on online mining methods provide encouraging results in terms of modelling this type of data streams; and the price of hardware keeps dropping while its power has increased to a level where it is possible to access significant computing power with low costs. All these allowed us to start a project with a very limited budget but with a very high potential in terms of ROI.

The first version of the TNT system has these main characteristics:

- Intra day trading with DJ EuroStoxx50 and S&P 500.
  - Ready for use with any other active but these were selected based on costs of the intra day quotes data.
- Working in real time since Jul/Aug of 2003 till Aug 2004 with decisions made every 15 minutes.
- Completely autonomous system.
  - From data collection to decision communication there is no human intervention.
  - Daily reporting (summary of decisions and results).
  - Real time information of decisions.

The results of this first year of real time simulation with the system have shown that it is able to generally achieve an interesting rentability with a manageable risk exposition. Moreover, our analysis of results shows a large margin of improvement (which is being done on the currently under development version 2).

2 Main Characteristics of the TNT System

The main goal of the TNT system is to trade actives in such a way that the positions that are held over time are profitable. Moreover, another relevant issue is that the risk of these positions should be low. Informally, this means that during the period a position is held it is not only important that the net result is positive, but it is also important that during the holding period the position does not go over periods of large losses.

In our system we are assuming that we are trading with futures which means that we have two types of positions: long and short positions. Long positions are the more "normal" positions where we buy an active at time t and sell it later at time $t + h$. A short position is a notion that may feel awkward for non-trading specialists. In this type of positions we open the position buy selling the active, and close it at a later time by buying. This means that when we open a short position we are in effect making a promise that we will buy the active in the future, which we obviously expect to happen at a lower price than the current price at which we are selling, i.e. we are betting on a downwards trend of the prices, while when we open a long position we are betting on an upwards trend. Figure 1 illustrates these two types of positions.

Every 15 minutes the TNT system will make a decision concerning the position it should take. Positions are taken by issuing orders: buy, sell, or hold (do nothing). The system makes this decision looking at several factors:

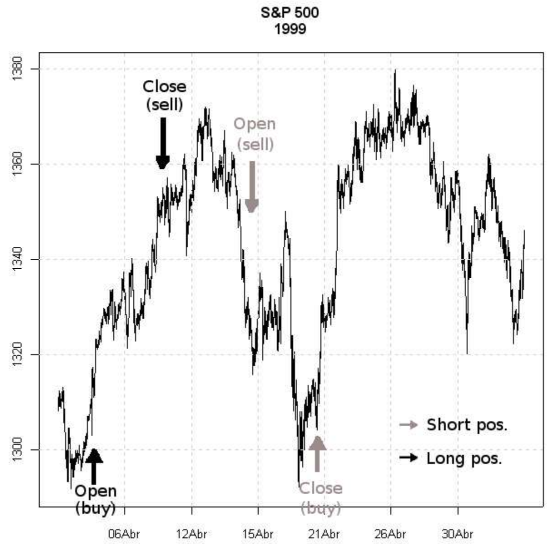

Fig. 1. The decision problem in financial trading.

- The forecast of the future prices of the active according to the predictive models used by the system.
- The current position the system is holding for the active (are we currently short, long or out of the market?).
- The risk handling strategy of the system: this involves looking at the rentability of the current position (if any), and checking it against the risk preference criteria of the system.

It is the interaction between these factors that leads to a decision, which means that not only the predictions of the data mining models have impact on the decisions.

2.1 The modelling tasks

As we have already mentioned the TNT system uses a large variety of models for obtaining predictions of the future evolution of the prices. This variety has the objective of fighting the complexity of the modelling tasks by developing models with quite different biases that may be more adequate to different dynamic regimes of the prices time series we are forecasting [2].

The modelling task we are facing poses several important question that need to be addressed, namely:

- How to handle the huge and constantly growing data that we have available?

- Forget old data? When? How to solve computational efficiency problems?

– Which information to use ?
  - Quote prices only? News? Prices of related assets (e.g. gold, oil, etc.) ?

– Which predictor variables should be used in model construction?

– What should be the target variable?

– How to move from predictions into decisions?

All these questions have many possible alternatives and there is not necessarily a correct answer. Again, this creates several possible sources of variability among different models.

Instead of searching for the "ideal" answer to these questions, which would assume there there exists such answer, we have followed a different path. We have generated a quite large set of alternative models (that we will call artifical traders as explained in Section 2.2), and then we have used a dynamic online selection process based on the current compettitiveness of the alternatives. This means that at any moment the system checks the recent performance of the different artificial traders involved and weights their "opinions" based on the results of this performance analysis. This provides means for copying with non-stationarity effects that are known to exist in financial time series data. This means that in different periods of time we may have different answers to the questions outlined above as being the current "best".

With respect to the alternatives that were used to generate the different variants of artificial traders we have used different ways of handling the growth of the data (sliding windows of different sizes, growing windows, etc.); we have used basically information contained only in the quotes although with several variables created based on this price information (e.g. technical indicators that try to capture different dynamic properties of the prices time series); we have used different alternative target variables (e.g. future returns with different future lags); we have used different risk profiles and risk handling strategies as well as different ways of interpreting the predictions (i.e how to move from a numerical prediction into a trading signal).

Finally, we have also used several modelling techniques like for instance neural networks (with different parameterizations), support vector machines and so on.

2.2 Ensembles of artificial trading agents

TNT uses a large set of internal artificial trading agents that fight together for influencing the final decisions of the system. Each of these artificial traders has several components, each of them with several possible alternative choices in terms of different parameterizations. Figure 2 gives an overall picture of the architecture of each artificial trader in TNT, as well as the main parameters of each component.

Each of the artificial traders "inside" TNT runs in continuous in a kind of simulated trading environment, whose goal is to collect statistics of its performance. Based on these performance values TNT achieves an ensemble prediction of the order / position it should hold in future. This means that TNT uses a kind

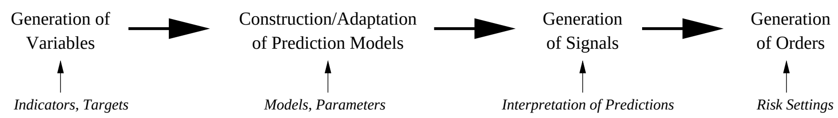

Fig. 2. The architecture of an artificial trader.

of dynamic model averaging [1], as these statistics are constantly being updated as time goes by. In this first version we have tried several alternative ways of forming these kind of committees of the "best" artificial traders based on their recent performance. These alternatives have to do with how many traders enter the committees, how their performance is evaluated and so on.

3 Results of the Real Time Simulation

As mentioned before we let the system run for around one year in real time and registered its trading performance. In this section we present and discuss some of the results.

3.1 S&P 500 results

Table 1 shows the overall performance statistics of the best 3 committees for the simulation period. Figure 3 shows the monthly rentabilities of these 3 committees.

Table 1. Overall results of best 3 committees on S&P 500. *rate* - overall rentability of the strategy; *rate.bh* - the rentability of the market during the simulation period; *over.bh* - gains over the market; *max.drawdown* - maximum successive loss; *sharpe* - Sharpe Ratio; *sterling* - Sterling Ratio.

	rate	rate.bh	over.bh	max.drawdown	sharpe	sterling
best.r2	27.49	13.92	13.57	5.6	1.815	4.409
best	24.86	13.92	10.94	5.9	1.701	3.497
best.r1	15.12	13.92	1.20	9.8	1.045	1.259

The overall performance of the committees is quite satisfactory, namely the two best committees, both in terms of rentability (even taking into account that this was a period where the market was positive, 13.92%), as well as in terms of risk exposure. Still, a deeper analysis of the results has shown that the dynamic weighing schema that was used is loosing some of the potential of the individual traders as the best artificial traders achieved the results shown on Table 2 during the same period. Nevertheless, the larger profit achieved by these traders comes

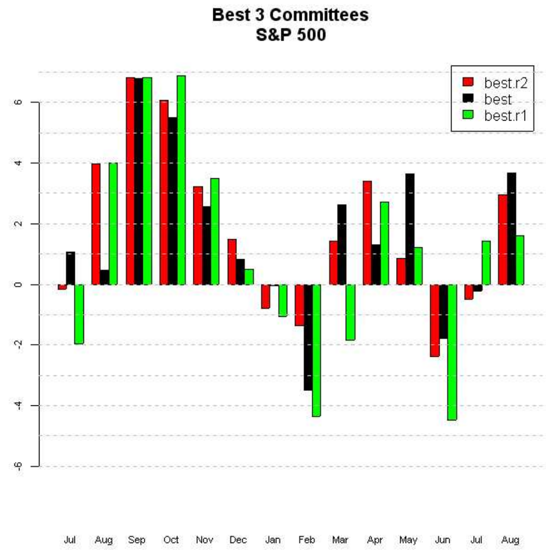

Fig. 3. The best 3 committees for S&P 500.

at the cost of some larger risk exposure, which confirms a kind of averaging effect so typical of ensembles' performance.

Table 2. Overall results of best 3 artificial traders on S&P 500.

	rate	rate.bh	over.bh	max.drawdown	sharpe	sterling
AT1	34.72	13.92	20.80	8.1	1.974	3.404
AT2	29.01	13.92	15.09	7.8	1.743	3.058
AT3	24.08	13.92	10.16	10.5	1.448	1.795

3.2 DJ EuroStoxx50 results

The overall results for DJ EuroStoxx50 are presented in Table 3. As we can observe the results are not so interesting particularly when compared to the market. Moreover, there are large drawdowns during this period which is not very promising in terms of real trading.

However, a deeper analysis shows that in this case it is more evident that the weighing schema we have used is not the best, as the results of the best 3 artificial traders shown in Table 4 indicate.

Table 3. Overall results of the 3 best committees on DJ EuroStoxx50.

	rate	rate.bh	over.bh	max.drawdown	sharpe	sterling
wavg3	10.46	11.99	-1.53	14.2	0.501	0.667
wavg3.r2	9.83	11.99	-2.16	18.2	0.464	0.462
wavg3.r1	4.72	11.99	-7.27	18.8	0.224	0.215

Table 4. Overall results of best 3 artificial traders on DJ EuroStoxx 50.

	rate	rate.bh	over.bh	max.drawdown	sharpe	sterling
AT1	27.86	10.6	17.25	8.2	1.431	2.711
AT2	23.53	10.6	12.93	10.3	1.120	1.820
AT3	23.53	10.6	12.93	10.3	1.120	1.820

4 Conclusions

We have briefly presented the midterm results of a commercial project that has the goal of developing an autonomous artificial trader. We have sketched the main components of the system as well as some of its main ideas.

The system was actively tested during one year of real time simulation. The results of this simulation indicate some potential of the system and provide relevant queues on how to improve its performance.

We are currently finishing the second version of the system according to our previous results analysis. Soon another period of active simulation will start, and hopefully that will lead to the active use of the system in a real world trading environment.

References

1. J.Z. Kolter and M.A. Maloof. Dynamic weighted majority: A new ensemble method for tracking concept drift. In *Proceedings of the Third IEEE International Conference on Data Mining*, Los Alamitos, CA, 2003. IEEE Press.
2. Hui Zou and Yuhong Yang. Combining time series models for forecasting. *International Journal of Forecasting*, 20:69–84, 2004.