Intraday trading rules based on Self Organizing Maps

Marina Resta

University of Genova

DIEM sez. Matematica Finanziaria, via Vivaldi 2, 16126, Genova, Italy

email: mresta@unige.it

June 29, 2007

Abstract— Working with five minutes data, we have studied a number of trading rules based on the responses of Kohonen's Self Organizing Maps, evaluating the results with both financial and statistical indicators, as well as by comparison with classical buy and hold strategy. At the current stage our major findings may be summarized as follows: a) Kohonen's maps are helpful to localize profitable intraday patterns, and b) they generally make possible to achieve higher performances than common buy and hold strategy.

1 Introduction

Starting from early 90's a plenty of academic work has been spent to analyze the potential of neural networks (NNs since now on) to trade financial markets. During those years almost every aspect of financial markets has been explored by means of artificial neural networks: we have learned that some neural architectures seem to work better than others [17], and that NNs make possible to pick up trading opportunities by monitoring market volatility [8], or that some hybridisation with genetic algorithms can help to achieve better results [2, 9]; additionally, NNs have been proficiently combined with more traditional tools, like moving averages [16], and other popular technical indicators [10, 14]. It is then clear that, at the present state, NNs applications in financial markets have reached a fully developed stage, with smaller space for new contributions to the ongoing debate. However, we have focused on a relatively lesser known topic, that is the profitability of intraday activity driven by neural networks. Previous efforts mainly concentrated in the FOREX (Foreign Exchange) market, with applications of reinforcement learning and genetic algorithms together with a collection of trading indicators [10]. In addition, evolutionary neural trees have been applied to forecast high frequency returns of the Hang-Seng index within the month of December 1998 [4] : in this case the authors focused on the statistical accuracy of their model, comparing its mean squared error, and its mean percentage error to those of two benchmarking series, generated by a random walk and by an AR(1) model.

Respect to the cited works, the novelty of our approach may be briefly condensed as follows.

- Although a number of applications of nearest neighbour methods are known on daily data[1, 6, 12, 13], minor attention has been spent on intraday tradings[7]. With this in mind, we will use high–frequency data to train Kohonen's Self Organizing Maps (SOMs).
- We are interested to study if and how SOMs capabilities can be employed both to exploit relevant temporal patterns at intraday time frequencies, and to support market timing activity, in the specific case to trade the market one time unit in advance.
- Despite from classical approaches that focus on the monovariate series of closing prices (or, sometimes, of closing log–returns), SOMs will be trained with intraday returns bars, i.e. the arrays whose components are the returns on Open, High, Low, and Close (OHLC) prices at the given time frequency. This *modus operandi* is rather unexplored but, in our opinion, very promising, because it allows a more efficient use of the informative content of prices: instead of estimating the embedding dimension on close prices (which is generally employed to provide details about the length of exploitable patterns in monovariate time series), here we gain advantage by working at each time t with the n–uples (the OHLC returns bars) of realizations from four random variables that represent all the available information to traders[1].
- We will introduce a trading system whose rules are based on Self–Organizing Maps responses, and we will examine its capability to provide operators with the timing of their intraday market activity (buy, sell, or standby). In this context, the use of SOMs instead of any other vector quantization method can be motivated since the superior capabilities of SOMs (that we are going to illustrate) for the joined action of features selection and rule extraction. The performances will be then monitored with many statistical and financial indicators, and the results compared to those obtained with the classical buy and hold strategy.

In the framework depicted above, what remains of the paper is organized as follows. In Section 2, after describing the data we have employed, we will give details about the SOMs–based trading rules that we have implemented.

---

[1]If we accept the classical assumption that prices discount all.

Section 3 will show and discuss the results obtained, while Section 4 will conclude.

2 Data and methodology

2.1 A preliminar discussion about input space features

In every market session, once fixed a time unit (say 5 minutes), the unit opening price is called Open, the closing price is termed Close, while the highest and the lowest prices recorded within such unit are known, respectively, as the High (H) and the Low (L); financial practitioners are then used to call *OHLC bar* the array containing the Open, High, Low, and Close prices at the given time unit. Respect to the more usual representation through monovariate data (closing prices), this technique allows to represent every traded asset in the market by four distinct points per time unit: visually this corresponds to plot at each time t a bar (a straight line) whose height depends on the range spanned between Low and High; a horizontal line is put on the left–hand side of the bar, in the proximity of the opening value, and to the right of the bar, in the nearby of the closing value. An example is illustrated in Figure 1, where we show the 5 minutes OHLC bars of the Italian financial index S&P MIB, within a single daily session.
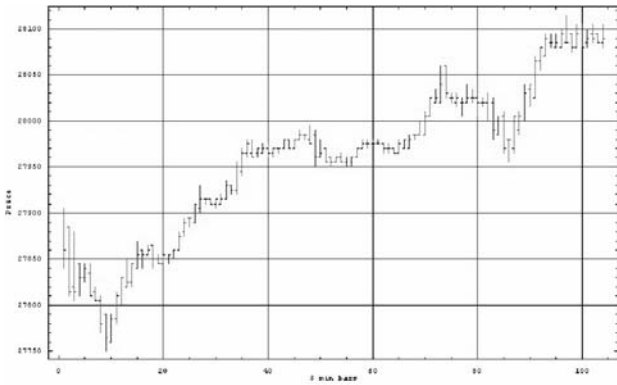

Figure 1: Five minutes bars of the S&P MIB index within the Italian Stock Exchange during the 17 September 2004 session.

In our study we have considered the 5 minutes OHLC bars for the S&P MIB index: this is the reference index for derivatives products traded in the Italian market since 17 September 2004, when the older MIB30[2] has been replaced as benchmark for the Italian stock market. The data were observed from the inception time of the index to 12 May 2006[3], for an overall amount of 43816 records. During this period the S&P MIB was characterized by various

uptrend and downtrend movements: after an upward tendence (from 09/17/2004 to 02/15/2005), the index has gone sideward from 02/15/2005 to 05/17/2005; hence it turned again to bullish (i.e. up, from 05/18/2005 to 09/30/2005). Finally, after a short bearish (down) stage (from 09/30/2005 to 10/28/2005), a new bull market episode characterized the index from 10/28/2005 to 05/12/2006. With a view to the various performances exhibited by the S&P MIB bars, we have divided the raw data into five blocks labelled B01, B02, B03, B04 and B05, corresponding to bullish (B01, B03 and B05), bearish (B04), and sideward (B02) behaviours of the index (see Table 1 for full details).

Name	Type	Starting HLOC bar	Final HLOC bar	Size
B01	Upward	09/17/04 h: 905	02/15/05 h: 1220	10926×4
B02	Sideward	02/15/05 h: 1225	05/17/05 h: 1740	6616×4
B03	Upward	05/18/04 h: 905	09/30/05 h: 1200	10020×4
B04	Downward	09/30/05 h: 1205	10/28/05 h: 1135	2074×4
B05	Upward	10/28/05 h: 1140	05/12/06 h: 1740	14178×4

Table 1: Main details about the blocks into which S&P MIB index bars have been splitted. For each block we have reported the prevailing trend features, together with the initial and final bars within the observation period, and the overall sample length.

Each block in turn, has been splitted into two subsets: earlier 70% of bars has been used to build the in–sample data to train the SOMs; the remaining 30% of every block served to construct the out of sample set. Respect to the original raw OHLC bars, we have done nothing but applying the classical definition of return:

$$r(t) = \log \frac{P(t)}{P(t-1)} \qquad (1)$$

where the argument of $\log$ in Eq.(1) has been replaced with the corresponding Open, High, Low and Close values at time t and $t-1$. In this way, instead of the scalar $r(t)$, at each time t we have got an array, the return bar: $\mathbf{rb}(t) = \left\{ \log \frac{O(t)}{O(t-1)}, \log \frac{H(t)}{H(t-1)}, \log \frac{L(t)}{L(t-1)}, \log \frac{C(t)}{C(t-1)} \right\}$.

To simplify the discussion further, from now on we will indicate the returns bars by different labels, according to the set they belong to, so that we will distinguish returns bars from the neural space ($\mathbf{rb}_N$), obtained through the training of the maps, by those in the input space, belonging to the in–sample ($\mathbf{rb}_{IN}$), or to the test set ($\mathbf{rb}_{OoS}$).

2.2 The SOM based trading rules

In order to give robustness to our study, we have considered for each sub–sample a system of 100 Self Organizing Maps [15]: results are then to be intended on average.

SOMs dimensions have been chosen after a preliminary data snooping procedure, during which we examined the sensitivity of convergence indexes[4] to changes in the number of neurons; Table 2 details the SOMs dimensional features chosen at last for each block of data.

---

[2]MIB30 is the acronym for Milano Indice Borsa: it has been the former leading indicator of the Italian Stock Exchange, built as a weighted sum of price levels from earlier 30 italian stocks by capitalisation.

[3]Source: Datastream.

[4]basically: those described in [5] and in [3].

Block Name	In Set Size	Test Set Size	Nr of Maps	Maps Dimensions
B01	7647×4	3278×4	100	29×15
B02	4630×4	1985×4	100	26×13
B03	7013×4	3006×4	100	28×15
B04	1451×4	623×4	100	20×10
B05	9924×4	4254×4	100	25×11

Table 2: Dimensional features of SOMs employed in our study

After the random initialization, the maps have been trained with bars from the in–sample set. The training stage was carried out for 50 epochs: the number of epochs was chosen as the number at which the already cited convergence indexes exceed the threshold of 60%.

Our trading system is based on a simple sequence of tasks that are illustrated in Figure 2: the flowchart to the left hand side shows the overall procedure, while in the right hand side we detail the organization of branches 1 and 2. Note that the first branch of the procedure is iterated over the input set (whose sample length is assumed equal to z) for a number of epochs depending on the level reached by the convergence indexes: this means that the input patterns are presented to the map more than once. The second part of the procedure, on the other hand, runs once in sequence from $t = 1$ to $t = v - 1$, where v is the length of the out of sample set.
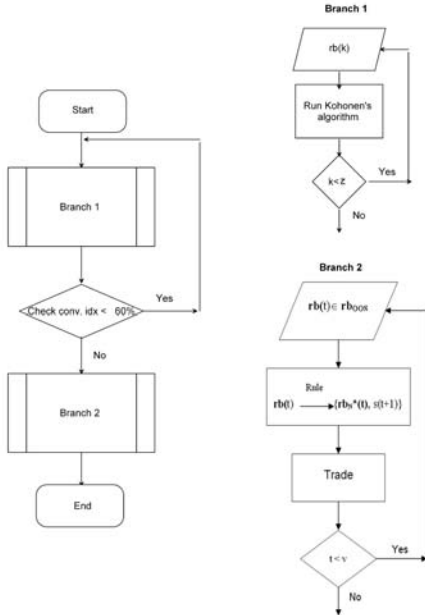

Figure 2: Flowcharts describing how our procedure is organized.

More precisely:

(1) The first branch performs the training of SOMs with the in–sample set. In this way, each returns bar of the input set ($\mathbf{rb_{IN}}$) is associated to a correspond-

ing pattern (the best matching unit) in the neural space ($\mathbf{rb_N}$). Once the 60% convergence threshold is reached, the second part of the procedure activates.

(2) We arrange the $\mathbf{rb_N}$ patterns in the Kohonen's map into three groups of signals: Sell (-1), Buy ($+1$), and Standby (0), according to proper rules that we will describe in a while.

(3) We perform the mapping:

$$\mathbf{rb}_{OoS}(t) \rightarrow \{\mathbf{rb}_N^*(t), s_t(t+1)\}, \quad \forall t = 1, \ldots, v-1$$

where v is the test set length, and $s_t(t+1)$ is the signal for trades at time $t + 1$, obtained in t from the arrays of the neural space. In other words, each $\mathbf{rb}_{OoS}(t)$ is coupled to the best matching unit $\mathbf{rb}_N^*(t)$ in the neural space (according to the Euclidean distance metric), and to the related signal that will be employed to trade one bar in advance.

The way the $\mathbf{rb}_N$ patterns are grouped is important for the overall procedure, since it can affect the signals suggested to operators. In our study, we have considered three different grouping strategies.

– **Rule 1. (R01)** Let us assume to have $\mathbf{rb}(t) \in \mathbf{rb}_{OOS}$. We indicate by $X(t)$ the set:

$$X(t) = \{x : x \in \mathbf{rb}_N^*(t) \wedge x > 0\} \qquad (2)$$

We are now able to evaluate the cardinality of $X(t)$ for the arrays of the map: according to the definition given in Eq.(2), since each bar contains four records, $card(X(t))$ may assume values in the interval $[0, 4]$, where the score 0 corresponds to whole negative OHLC returns bar, and 4 is associated to whole positive OHLC returns bar. The following rule has been then applied:

if $card(X(t)) \geq 3$
    then set signal $s_t(t + 1)$ to buy (+1)
        else if $card(X(t)) \leq 1$
            then set signal $s_t(t + 1)$ to sell (-1)
            else set signal $s_t(t + 1)$ to standby (0)
        end
end
return signal $s_t(t + 1)$

– **Rule 2. (R02)** As a variant of Rule R01, we have also considered a grouping criterion that strongly separates negative by positive patterns. Remembering that $card(X(t)) \in [0, 4]$, the new rule produces active signals (± 1) only in the case of completely positive or negative OHLC return bars:

if $card(X(t)) = 4$
    then set signal $s_t(t + 1)$ to buy (+1)
        else if $card(X(t)) = 0$
            then set signal $s_t(t + 1)$ to sell (-1)
            else set signal $s_t(t + 1)$ to standby (0)

end
end
return signal $s_t(t + 1)$

– **Rule 3 (R03)**. We have tested a rule that behaves assuming the most recent change in price as the best predictor for the future price variation. If we indicate by $\bar{y}(t)$ the last component of each $\mathbf{rb_N}(t)$ bar, i.e. the (neural) closing log–returns we get:

if $\bar{y}(t) > 0$
    then set signal $s_t(t + 1)$ to buy (+1)
        else if $\bar{y}(t) < 0$
            then set signal $s_t(t + 1)$ to sell (-1)
            else set signal $s_t(t + 1)$ to standby (0)
        end
end
return signal $s_t(t + 1)$

In this case, the rule is completely adaptive, and it depends on the behaviour of the (neural) closing fluctuations at the previous step.

From a practical standpoint, the system we have introduced is based on a double clustering action: the first one takes place when we train the map with $\mathbf{rb_{IN}}$ patterns; the final clustering task runs when we transform the 4th dimensional neural space organized by OHLC return bars into a mono–dimensional space, arranged by ± 1 and 0, from which we extract, one bar in advance, our action in the market.

(a) Map   (b) Rule 1
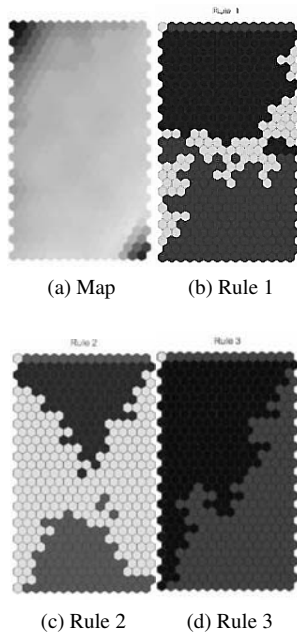
(c) Rule 2   (d) Rule 3

Figure 3: From left upper corner in clockwise sense: the original map, once trained over the $\mathbf{rb_{IN}}$ patterns, and once transformed according to Rules R01, R02, and R03.

Figure 3 visually illustrates how the procedure works through steps (1) and (2). Moving from the upper left side in clockwise sense, the first map is that containing $\mathbf{rb}_N$ arrays: each neuron represents one or more 4th dimensional input samples. In our case, the upper left section and the lower right part of the map (depicted in black and in hard grey) correspond, respectively, to OHLC patterns with highest and lowest fluctuations. Fading grey neurons, on the other hand, represent OHLC returns bars with more reduced fluctuation range. The other maps in Figure 3, show the appearance of the original map when it is organized according to rules R01 to R03: black represents sell positions, shaded grey buy positions, and, finally, white is for standby. The common factor to all rules is that they seem to organize the upper part of the map to sell position (with the exception of neurons in the first row, due to hedge effects), and the lower part to buy actions. It is also notable to observe that buy and sell areas are strongly separated into the map: this is particular evident for the maps derived from Rules R01 and R02, where white coloured cells of the SOM act as a border between buy and sell.

3 Case study

We are now going to evaluate the performances obtained by the trading system described in Sec.2.2, using both statistical and financial indicators. The results will be compared with those of classical buy and hold policy:

$$B\&H(\tau) = \log \frac{C(t + \tau)}{C(t)}, \quad \tau = 1, \dots, v - t$$

where $C(\cdot)$ is the Close value, and v is the out of sample set length.

We have referred to indicators whose use is well documented in the related literature, and mainly the ones employed in [11]. In particular, the statistical significance of the results has been evaluated by means of the indicators given in Table 3; financial relevance of the forecasts has been assessed with the performance measures shown in Table 4. In order to improve the readability of the results, a number of remarks are still noteworthy. All the statistical indicators have been computed using the series of observable closing log–returns:

$$y = \{y(t) : y(t) = \log[C(t)/C(t-1)], t = 2, \dots, v\}$$

opposed to the series of forecasts:

$$\tilde{y} = \{\tilde{y}(t) : \tilde{y}(t) = s_{t-1}(t) \times y(t), t = 2, \dots, v\}.$$

In addition, we have made use, when necessary, of the Dirac's δ function:

$$\delta(t) = \begin{cases} 1, & y(t) \times \tilde{y}(t) > 0; \\ 0, & y(t) \times \tilde{y}(t) \le 0. \end{cases} \quad t = 2, \dots, v$$

Financial performances, in turn, have been evaluated on the series:

$$\hat{\hat{y}} = \{\hat{\hat{y}}(t) : \hat{\hat{y}}(t) = sign[y(t)] \times \tilde{y}(t), t = 2, \dots, v\}.$$

Every time we needed to annualise the results, we have used the annualisation factor $k = 252 \times 104$, in order to take into account of both the daily effect (252 days per year), and the intraday effect (each daily session includes 104 bars at 5 minutes sampling frequency). Finally, we have denoted by $\mu_{\hat{\hat{y}}}$ the sample mean over the series of $\hat{\hat{y}}$'s.

Performance Measure	Description
Mean Absolute Error (MAE)	$MAE = \frac{1}{v-1} \sum_{t=2}^{v} |\tilde{y}(t) - y(t)|$
Root Mean Squared Error (RMSE)	$RMSE = \sqrt{\frac{1}{v-1} \sum_{t=2}^{v} [\tilde{y}(t) - y(t)]^2}$
Theil's inequality coefficient	$U = \dfrac{RMSE}{\sqrt{\frac{1}{v-1}\sum_{t=2}^{v}[\tilde{y}(t)]^2} + \sqrt{\frac{1}{v-1}\sum_{t=2}^{v}[y(t)]^2}}$
Correct Directional Change (CDC)	$CDC = \frac{1}{v-1} \sum_{t=2}^{v} \delta(t)$

Table 3: The statistical indicators used to evaluate the performances of the neural trading system.

Performance Measure	Description
Annualised Return	$AR = \kappa \times \frac{1}{v-1} \sum_{t=2}^{v} \hat{\hat{y}}(t)$
Annualised Volatility	$AV = \sqrt{\kappa} \times \sqrt{\frac{1}{v-2} \sum_{t=2}^{v} [\hat{\hat{y}}(t) - \mu_{\hat{\hat{y}}}]^2}$
Sharpe Ratio	$SR = \frac{AR}{AV}$
Maximum Drawdown	$MDD = \min_{\substack{i=2,\dots,t \\ t=2,\dots,v}} \sum_{i=2}^{t} \hat{\hat{y}}(i)$
Nr. of Up Periods	$NUP = card(\hat{\hat{y}}(t) > 0)$
Nr. of Down Periods	$ND = card(\hat{\hat{y}}(t) < 0)$
Average Gain in Up Periods	$AG = \sum_{t=2}^{v} \delta(t) \times y(t) / NUP$
Average Loss in Down Periods	$AL = \sum_{t=2}^{v} [1 - \delta(t)] \times y(t) / ND$
Average Gain/Loss Ratio	$AGL = AG/AL$

Table 4: Financial performance measures employed in our study.

3.1 Discussion of the results

For every subset (B01 to B05), when this makes any sense, we compare the results of SOMs related trading rules (R01, R02, and R03), to those of standard buy and hold (B&H) strategy. The scores obtained by the statistical indicators are reported in Figure 4, while the financial performances have been rendered in Figure 6, where we plotted for each

test set the Annualised Return (AR), the Sharpe Ratio (SR), the Maximum Drawdown (MDD) and the Average Gain–Loss ratio (AGL), obtained with both the SOMs–based trading rules, and the buy and hold strategy. In the case of statistical indicators, each bar on the horizontal axes corresponds (for rules R01, R02, and R03) to a block of test data (B01 to B05), while the scores obtained are reported on the vertical axes. In the case of financial performance measures, in addition to R01–R03 rules on the x–axes we have also taken into account B&H; on the vertical axes, we have reported the difference between the score of each strategy and that computed on observable log–returns. In the case of SR, such results have also been rescaled in the range [0,1].
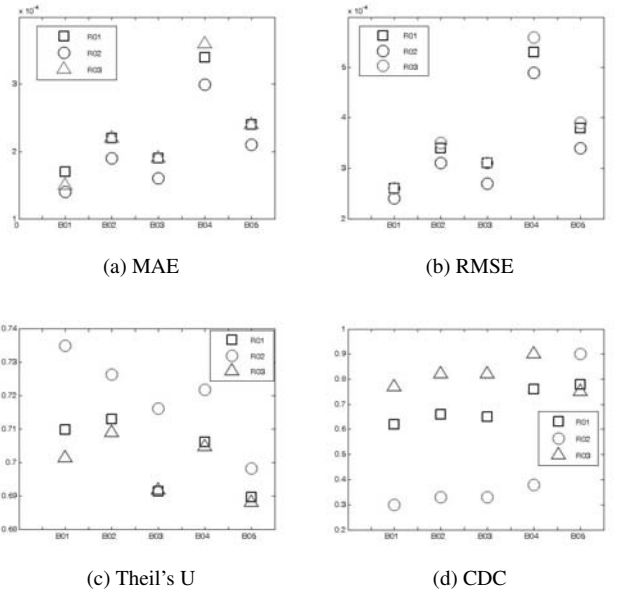

(a) MAE   (b) RMSE

(c) Theil's U   (d) CDC

Figure 4: Statistical performances of SOMs–based trading system. From top to down in clockwise sense: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Theil's inequality coefficient (Theil's U), Correct Directional Change (CDC).

The statistical results provide evidence of a satisfactory capability of the SOMs–based system on intraday data. Assuming 0.01% as the upper bound value, both MAE and RMSE maintain consistently below it, in all the examined data samples. Such performances are counterbalanced by less satisfactory scores achieved by the Theil's U index. At this stage, it could be argued that since U ranges from 0 (complete matching to the reference series) to 1 (maximum distance from the reference series), and the scores we have obtained are sensitively closer to 1 than to 0, the results as a whole are far to be considered encouraging. In order to mitigate such impression, we also studied the U statistics related to the series of cumulative returns $cy = \{cy(t)\}$, being:

$$cy(t) = \sum_{j=2}^{t} y(j), \quad t = 2, \dots, v$$

Numerical results are given in Table 5; the dynamics of observable cumulative returns (equity lines), opposed to those obtained by running Rules R01, R02, and R03 is given in Figure 5.

Name	B01	B02	B03	B04	B05
R01	0.798	0.549	0.674	0.946	0.798
R02	0.760	0.455	0.690	0.834	0.800
R03	0.609	0.752	0.740	0.832	0.813

Table 5: U statistics for the cumulative closing log–returns of the test sets, in the case of different rules R01, R02, and R03.

(a) B01

(b) B02
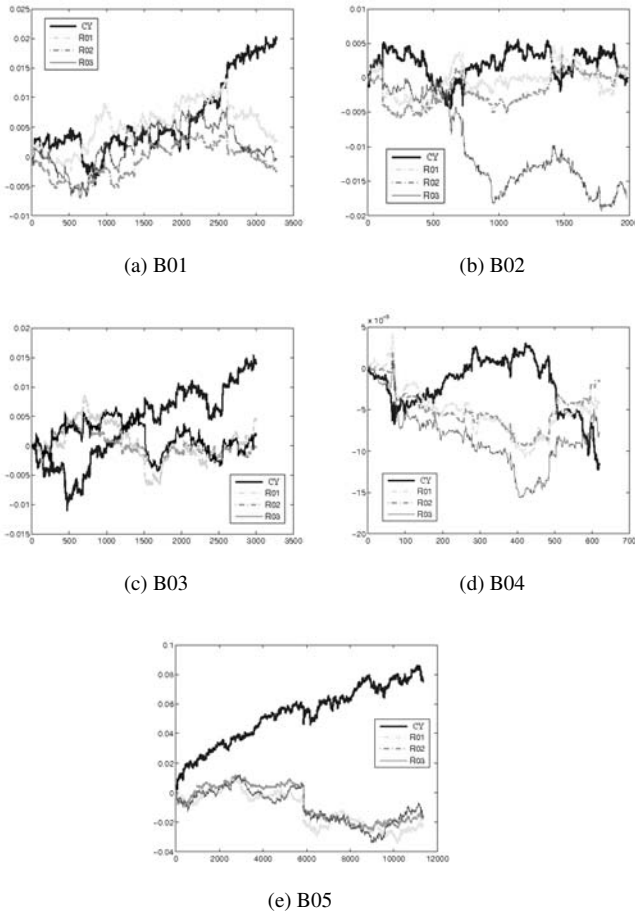
(c) B03

(d) B04

(e) B05

Figure 5: From left to right and from top to bottom: the comparison among equity lines and the curve of cumulative returns as resulting by running Rules R01, R02, and R03, for various test sets.

By comparison of the results in Table 5 with the behaviour of equity lines, it clearly emerges that higher U values are not necessarily related to poor overall performances. More importantly, our system makes possible to capture the market general trend: this is particularly true in uptrend periods B01, B03, and in the sideward phase B02.

Those remarks are somewhat confirmed when we look at the differential scores of financial indicators. In general, the SOM–based system has assured notable annual-
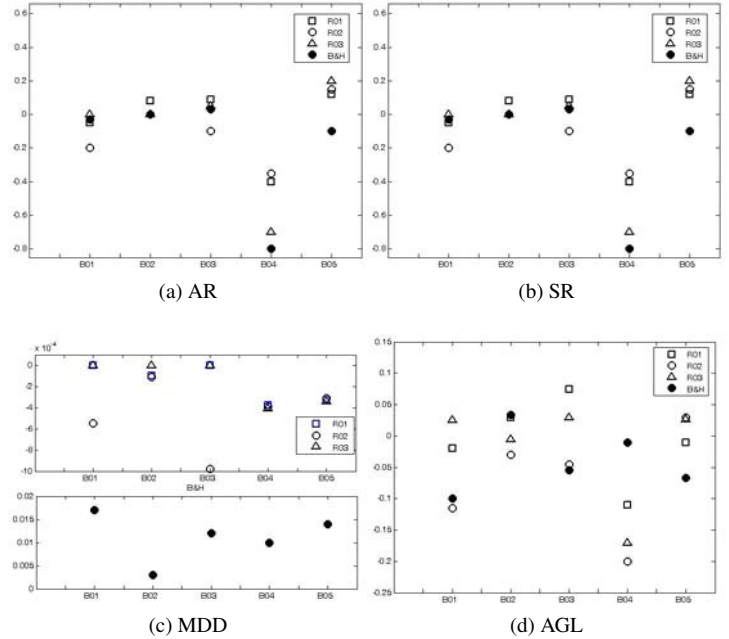

(a) AR

(b) SR

(c) MDD

(d) AGL

Figure 6: Financial Performances of SOM based trading systems. From top to down in clockwise sense: differential values of Annualised Return (AR), Sharpe Ratio (SR), Maximum Drawdown (MDD) and Average Gain–Loss ratio (AGL). The SR scores are scaled within the range [0,1].

ized returns. In particular, rules R01 and R03 have run successfully in three of five data blocks. The performances of rule R02 are apparently penalized by the greater number of standby (0) positions suggested by the rule itself. Note that the buy and hold strategy does not appear to be proficient on intraday data, and, according to the Sharpe Ratio scores that are sensitively lower than in the case of SOM–based rules, it also pays a greater tribute in terms of volatility. Another important evidence comes from the values of MDD and AGL. The SOM–based rules seem to hedge with respect to losses: since we have reported differential values, a negative MDD score means that the neural (R01, R02 or R03) strategy has guaranteed a MDD lower than the one observed on real data; the opposite holds in the case of B&H. Similar considerations may be done for the AGL ratio: non negative values (see for instance R01 and R03 on Figure 6 in four of five blocks) show the capability of our system to gain advantage from intraday fluctuations.

4   Conclusions

We have built a trading system based on Self Organizing Maps (SOMs) to operate on financial markets at intraday trading frequencies. To such aim, we have considered the 5 minutes OHLC bars for the S&P MIB, observed in the interval: 17 September 2004–12 May 2006. Such raw data were splitted into five blocks, according to the prevailing trend (upward, downward or sideward): 70% of samples in each block has been employed to train the sys-

tem, while the remaining 30% has been used as test set. We have then trained SOMs (whose dimensions and parameters were selected through a data snooping procedure) with the in–sample returns bars, and we have associated to each reference vector (i.e. to each neuron) a trading signal, according to three rules that we have introduced and discussed in the previous sections. Finally, we presented the patterns of the test set to the so clustered maps, coupling each returns bar at time t with the signal to trade the index closing price at time $t + 1$. The results were evaluated with the help of a number of statistical and financial indicators. The results obtained are quite promising, since they give evidence (according to both statistical and financial indicators) that SOMs work well when they are called to discover intraday patterns, and such retrieved information can be used to perform market tradings. Some problems still remain open. The first odd is of "technical "nature: we have assumed to stop the learning procedure when a proper threshold value is reached by some convergence indexes, but clearly other stopping solutions could be equally promising. The second problem is the evaluation of the representiveness of the input set respect to the test set: it is evident that strong divergences among such blocks of data can seriously compromise the overall performance of the system. This, for instance, could be the reason for poor performances in the case of the B04 set of data; an alternative explanation, however, could be that SOMs simply run better on uptrend rather than on downtrend patterns. A possible solution (presently not explored in deepest detail) could be that to consider different neural trading systems specialized on various aspects of the market (bullish/bearish), with the possibility of switches from one to another, according to the fluctuations features of the market. Finally, the system we have presented relies on two clustering tasks, the second of which is rule dependent; we have implemented three different rules that seem to work reasonably well, but a straightforward step for future works is the study of more efficient rules.

References

[1] Andrada-Félix J., Fernández-Rodríguez F., García Artiles M. D. and S. Sosvilla-Rivero (2005), *Non-linear trading rules in the New York stock exchange*, Working paper, University of Las Palmas, Spain.

[2] Arifovic, J., and R. Gencay (2001), *Using genetic algorithms to select architecture of a feedforward artificial neural network*, Physica A: Statistical Mechanics and its Applications, 289(3-4), 574-594.

[3] Cattaneo Adorno, M. and M. Resta (2004), *Reliability and convergence on Kohonen maps: an empirical study*, in M. Gh. Negoita, R. J. Howlett and L. C. Jain, Proceedings of Knowledge-Based Intelligent Information and Engineering Systems, 8th International Conference, KES 2004, Wellington, New Zealand, Part I, Lecture Notes in Computer Science, Springer, 426-433.

[4] Chen, S. H., Wang, H. S., and B. T. Zhang (1999), *Forecasting high-frequency financial time series with evolutionary neural trees: the case of Hang-Seng stock index strategies*, Proceedings of the 1999 International Conference on Artificial Intelligence (IC-AI'99), 437-443.

[5] De Bodt, E., Cottrell, M., and M. Verleysen (2002), *Statistical tools to assess the reliability of self-Organizing maps*, Neural Networks, 15, 967-978.

[6] Dablemont, S., Simon, G., Lendasse, A., Ruttiens, A., Blayo, F., and M. Verleysen (2003), *Time series forecasting with SOM and non–linear models–Application to DAX30 index prediction*, Procedings of WSOM 2003, Workshop on Self–Organizing Maps, Hibikino (Japan), 11-14 September 2003, 340-345.

[7] Dablemont, S. and M. Verleysen (2004), *Classification et prdiction fonctionnelles d'actifs boursiers en intraday*, AC-SEG 2004, Connectionist Approaches in Economics and Management Sciences, Lille (France), 18-19 November 2004 (in French), 30-38.

[8] Deboeck, G. J. (1999), *Modelling non-linear market dynamics for intra-day trading*, available on-line at: http://www.dokus.com/PapersontheWeb/intradaymodel.htm

[9] Dempster M. and C. M. Jones (2001), *A real-time adaptive trading system using genetic programming*, Quantitative Finance 1(4), 397-413.

[10] Dempster M., Payne T., Romahi Y., and G. W. P. Thompson (2001), *Computational learning techniques for intraday FX trading using popular technical indicators*, IEEE Transactions on Neural Networks 12, 744-754.

[11] Dunis, C. L. and M. Williams (2002), *Modelling and trading the EUR/USD exchange rate: do neural network models perform better?*, CIBEF working paper, available on-line at: cwis.livjm.ac.uk/bus/cibef/workingpapers/cibef0202.pdf

[12] Fernández-Rodríguez F., Sosvilla-Rivero, S. and M. D. Garca-Artiles (1999), *Dancing with bulls and bears: nearest-neighbour forecasts for the Nikkei index*, Japan and the World Economy 11, 395-413.

[13] Fernández-Rodríguez F., Sosvilla-Rivero, S. and J. Andrada-Félix (1999), *Exchange-rate forecasts with simultaneous nearest-neighbour methods: evidence from the EMS*, International Journal of Forecasting 15, 383-392.

[14] Fernández-Rodríguez F., González-Martel C., and S. Sosvilla-Rivero (2000) *On the profitability of technical trading rules based on artificial neural networks: Evidence from the Madrid stock market*, Economic Letters 69, 89–94.

[15] Kohonen T. (1982), *Self-Organizing Maps*, Springer Verlag, Berlin.

[16] Refenes, A. P. N., Burgess, A. N., and Y. Bentz (1997), *Neural networks in financial engineering: A study in methodology*, IEEE Transactions on Neural Networks, 8(6), 1222-1267.

[17] Refenes, A. P. N., Azema-Barac, M., Chen, J., and S. A. Karoussos (1993), *Currency exchange rate prediction and neural network design strategies*, Neural Computing & Applications, 1(1), 46-58.