

Tradinnova-ACO: Ant Colony Optimization Metaheuristic applied to Stock-Market Investment

Isidoro J. Casanova and José M. Cadenas

Dpto. Ingeniería de la Información y las Comunicaciones.
Facultad de Informática. Universidad de Murcia.
Campus de Espinardo. 30100 Murcia. Spain.
isidoroj@um.es
jcadenas@um.es

Abstract. *A stock market investor buys and sells stocks in order to obtain the best possible profit. This dealing can be depicted on a graph, on which each node represents the stocks purchased. In this way, this paper proposes the use of ant colony optimization for calculating the best sequence when buying and selling, thereby helping when deciding how to invest.*

Key words: Stock-Market Investment, Ant Colony Optimization, Intelligent Systems, Decision Making, Heuristics.

1 Introduction

There are different techniques which allow us to see whether a particular stock is valued as cheap or expensive (fundamental analysis) or whether the price is in an upward or downward trend (technical analysis). However, an investor should not really focus on just the one particular stock, but on all the stocks making up a specific market (in our case the Continuous Spanish Market), or even focus on stocks in other countries, in order to create a stock portfolio.

When creating a portfolio, you can take into account the risk aversion of the investor, who, following the Markowitz theory, will want to obtain maximum profit with minimum risk ([16]). Usually, when a portfolio of stocks is created, the investor hopes to make a profit within a specific period of time, without a dynamic management of the stocks making up the portfolio.

In [1], we proposed a heuristic algorithm called TRADINNOVA, which intelligently simulates the performance of an investor in the Continuous Market applying rules in order to buy and sell stocks. This algorithm is based on the premise that on any given day in the stock market you can select a combination of stocks with which you hope to make a good profit. The heuristic techniques in order to select these stocks can be extremely diverse, from technical or fundamental analysis to artificial intelligence techniques such as neuronal networks, Bayesian networks, news analysis, market climate analysis, etc ([6–13]). Once these stocks have been selected, TRADINNOVA chooses the best stock to buy and its purchase price. As the days go by it decides when to sell and the selling price, what other stocks can be bought and also follows up the buying or selling orders which are not carried out, in order to act upon them.

TRADINNOVA dynamically manages stocks in a period of time, widening the traditional Markowitz model which does not stipulate the moment for buying and selling the stocks, and extending the capacities of selection of the stocks as it is the investor who chooses the heuristic technique to use when selecting stocks, thereby adapting the algorithm to the personal preferences of the investor.

In this paper, we present an extension of TRADINNOVA, called TRADINNOVA-ACO, in which in order to select the stocks to be bought or sold an ant colony optimization algorithm is used.

The paper is organized in the following way: firstly, we introduce the field of stock market investment, commenting on the prediction techniques which are generally used, and we introduce the concept of the optimum sequence in buying and selling. In section 3, we introduce ant colony optimization, seeing how the problem is characterized, in order to then see, in the following section, how this metaheuristic can be used to make stock market predictions. Finally, we finish with conclusions and future research.

2 Stock market investment

2.1 Prediction techniques

Apart from the techniques which try to determine whether a stock is valued as cheap or expensive (fundamental analysis) or whether its price is in an upward or downward trend (technical analysis), there are numerous artificial intelligence studies to predict the value of a specific stock, with neuronal networks being the most frequently used technique ([6-8]). Genetic algorithms are also used to optimize the parameters of these neuronal networks or to optimize the technical analysis of a stock ([9,10]).

There are also studies into how to predict the performance of a stock from news published on the Internet ([12,13]).

If we go by [11], the techniques to apply in finance can be divided into 5 categories (Fig. 1).

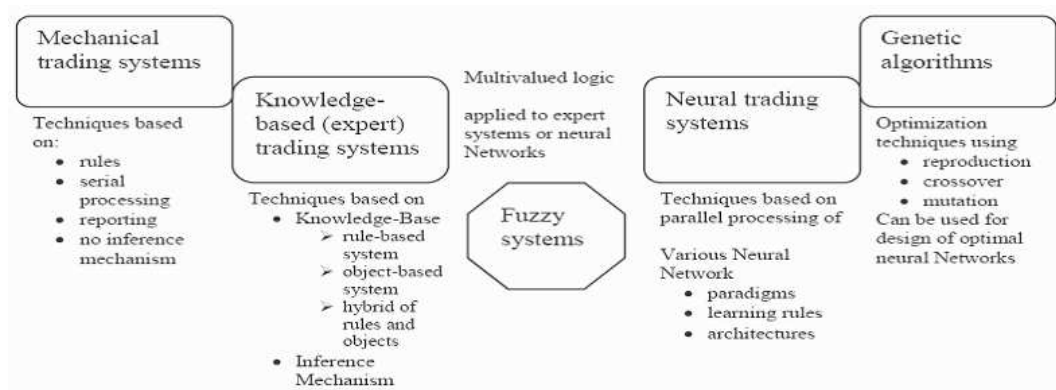


Fig. 1. Classification of the techniques to apply to finance

All techniques advise us on the purchase of a particular stock. However, once the performance of a stock, or a group of stocks, has been predicted, it would be necessary to apply a policy of buying and selling in order to maximize the profit and thereby provide a dynamic management of the stocks, as done in TRADINNOVA.

2.2 Optimum dealing sequence

As the days go by, in the Continuous Market, stocks rise and fall by a certain percentage, and an investor will sell and buy stocks throughout those days. But, what would be the ideal dealing sequence in order to maximize profit?. In principle, the answer would be easy if we already knew how the market was going to perform, as the ideal situation would be to invest, each day, all of the money in the stock in the market which rises the most on that day, as long as the stock rises enough so that commissions don't lose us money in the end.

It is of course, impossible to know which stock this will be. Due to this drawback, what the investor has to do is invest in the stock which has the best chance of rising (statistically, technically, fundamentally, etc). Obviously, a stock has more chance of rising if it is in an upward trend, and supposedly, if a stock is in an upward trend, the purchase of that stock will form part of an "optimum" dealing sequence.

In order to obtain this "optimum" dealing sequence we are going to use the metaheuristic of the ant colony. Here we will briefly describe this metaheuristic.

3 Ant Colony Optimization (ACO)

3.1 Introduction

Ants are social insects which live in colonies and, due to their mutual collaboration, are capable of displaying complex behavior and completing difficult tasks from the point of view of an individual ant.

While they move within the ant hill and the food source, some species of ants deposit a chemical substance known as pheromone (a substance which can be "smelt"). If no trace of pheromone is found, the ants move in a basically random manner, but when deposited pheromone is in existence, they are more likely to follow the trail. Given that ants deposit pheromone on the path they are following, this behavior leads to positive feedback, concluding in the formation of a trail marked out by a higher concentration of pheromone. This behavior allows ants to find the shortest paths between the ant hill and the food source.

3.2 From natural ants to the ant colony optimization metaheuristic

ACO algorithms are directly inspired by the behavior of real ant colonies in order to solve combinatorial optimization problems (NP-Hard problems). The artificial ant is a simple computational agent which tries to find possible solutions to the problem by exploiting the available pheromone trails and heuristic information. ACO algorithms are essentially constructive algorithms: in each repetition of the algorithm, each ant builds a solution to the problem referring to a construction graph.

Characterization of the problems. In general the problems to be resolved by this algorithm have the following characteristics:

1. There is a finite set of components $N = \{n_1, n_2, \dots, n_k\}$, each of these will be represented by a node on a graph.
2. There is a set of constraints *REST* defined by the problem to be solved.
3. The problem presents diverse states which are defined according to components ordered by sequences $EST = \langle n_r, n_s, \dots, n_u, \dots \rangle$ ($\langle r, s, \dots, u, \dots \rangle$ to simplify) on the elements of N . If T is the set of all possible sequences, we will call S the set of possible (sub)sequences which respect the constraints *REST*. The elements in S define the possible status, which correspond with paths on the graph, that is, sequences of nodes or edges.
4. There is a neighborhood structure which defines the edges of the graph. So a_{rs} would be the connection/transition of n_r with its neighbor of n_s if:
 - Both n_r and n_s belong to T .
 - The state n_s can be reached from n_r in a logical step, there should be a valid transition between r and s . The neighborhood reachable by n_r is the set which contains all of the sequences $N_s \in S$.
5. Explicit costs, c_{rs} , associated with each edge must exist.
6. A solution, *SOL*, is an element of S which verifies all of the requisites of the problem.
7. There is a cost, $C(SOL)$, associated with each solution *SOL*.
8. In some cases, it is possible to associate a cost, or an estimation of it can be associated with the states.

Structure of a generic ACO algorithm. Next we will show the original algorithm in pseudocode of the ant colony metaheuristic optimization, as specified in [3].

```

Procedure ACO-Meta_heuristic()
  while (termination_criterion_not_satisfied)
    schedule_activities
      ants_generation_and_activity();
      pheromone_evaporation();
      daemon_actions(); {optional}
    end schedule_activities
  end while
end Procedure

```

```

Procedure ants_generation_and_activity()
  while (available_resources)
    schedule_the_creation_of_a_new_ant();
    new_active_ant();
  end while
end Procedure

```

```

Procedure new_active_ant() {ant lifecycle}
  initialize_ant();
  M = update_ant_memory();
  while (current state <> target state)
    A = read_local_ant-routing_table();
    P=compute transition probabilities(A,M,problem_constraints);
    next_state=apply_ant_decision_policy(P,problem_constraints);
    move_to_next_state(next_state);
    if (online_step-by-step_pheromone_update)
      deposit_pheromone_on_the_visited_arc();
      update_ant_routing_table();
    end if
    M = update_internal_state();
  end while
  if (online_delayed_pheromone_update)
    evaluate_solution();
    deposit_pheromone_on_all_visited_arcs();
    update_ant_routing_table();
  end if
  die();
end Procedure

```

Fig. 2. Original ACO algorithm

How it works. The ants (artificial) in a colony move, concurrently and asynchronously, through the states adjacent to the problem (which can be represented in graph form with weights). This movement is carried out by following a rule of transition based on the local information available in the components (nodes). This local information includes the heuristic and memoristic information (pheromone trail) to guide the search. On moving around the construction graph, the ants build solutions incrementally. Optionally, the ants can deposit pheromone each time they cross an arc (connection) while they build the solution (linear renewal, step by step, of the pheromone trail). Once each ant has generated a solution it is evaluated and can deposit a quantity of pheromone depending on the quality of the solution (linear renewal of the pheromone trails). This information will guide the search of the other ants in the colony in the future.

4 Use of the ant colony for stock market prediction

4.1 Characterization of the problem

What we will attempt to optimize is the daily sequence of buying and selling stocks, starting from their historic price. Therefore each node on the graph will contain a specific stock next to the price registered on a particular day.

Each day it will be possible to either have the money invested in one of these stocks or keep it in cash, and the next day, keep the stock that we had, or sell it in order to recover the money, with which you may (or may not) wish to acquire a new stock.

This deal will produce a profit which accumulates over the days. To select which stock it is best to invest in each day we can use different heuristic techniques, such as technical stock analysis.

To illustrate the problem, we can observe Fig. 3, in which on the first day the investor has a quantity of money which can be invested (or not) in any of the stocks which constitute the market, in this case, A, B or C. For the following day, the investor can either change where they have the stock or cash them in. The days will continue to go by until money has definitely been made and we can check how much money has finally been made with the dealing sequence. What we are trying to discover is the most profitable buying and selling sequence.




Fig. 3. Graph for stock dealing

If we go by the ACO terminology explained in 3.2 each ant would represent an investor buying and selling stocks on the continuous market:

1. For each day there would be different nodes on the graph representing the stocks and their values: $N = \{a_1, b_1, c_1, \dots, n_1, a_2, b_2, c_2, \dots, n_2, a_m, b_m, c_m, \dots, n_m\}$. Where (a, b, c, \dots, n) represent the prices fixed by n different stocks from day 1 until day m .
2. There are constraints to be applied (*REST*) so if each day an ant is in one of the nodes corresponding to that day, the next day it will either be able to continue in that node or jump to any of the other nodes corresponding to the following day.
3. The states of the problems are any combination of the elements of N . We would designate as S the sets of possible (sub)sequences which respect the constraints *REST*. The elements in S would define the buying and selling done by an investor over m days.
4. The neighbourhood structure which defines the edges of the graph will be defined so that if we find ourselves in the node $i_r \in (a_r, b_r, c_r, \dots, n_r)$, we will only be able to move to node j_s as long as $i \neq j$ y $s > r$. In other words, if we find ourselves invested in a

stock i on one day r , the transition which would be considered valid would be to invest in a stock j (different to the previous one) on any other day after s .

5. There are costs or profits which are produced as we move around the edges, depending on whether there is a profit or loss made on the different stock deals.
6. A solution SOL is an element of S that verifies all of the requisites of the problem.
7. There is a profit $B(SOL)$ associated with each solution SOL .
8. When a state has been invested in, we can calculate the profit expected by moving to another state using different heuristic techniques (technical, fundamental, neuronal network analysis).

4.2 Application of the metaheuristic: TRADINNOVA-ACO

The proposal is the application of the following ACO metaheuristic algorithm to stock market investment (modified from the original ACO algorithm specified in [2] and [3]):

```

Procedure Tradinnova-ACO()
  parameter-initialization
  while (termination-criterion-not-satisfied)
    schedule-activities
      investor-generation-and-activity()
      pheromone-evaporation()
      daemon-actions() {optional}
    end schedule-activities
  end while
end Procedure

Procedure investor-generation-and-activity()
  repeat in parallel for k=1 to m (number-of-investors)
    new-investor(k)
  end repeat in parallel
end Procedure

Procedure new-investor(investor-id)
  initialize-investor(investor-id)
  BS = initialize-investor-buy-sell-memory()
  while (current-date <> end-date)
    STC = compute-recommended-stocks ( Heuristic-Type, BS )
    selected-stock = apply-decision-policy ( STC, BS )
    buy-and-sell-to-move-to-next-state( selected-stock )
    if (on-line-step-by-step-pheromone-update)
      deposit-pheromone-on-the-visited-edge()
    end if
    BS = update-buy-sell-internal-state()
  end while

```

```

if (online-delayed-pheromone-update)
  for each visited edge
    deposit-pheromone-on-the-visited-edge()
  end for
end if
release-investor-resources(investor-id)
end Procedure

```

Fig. 4. TRADINNOVA-ACO: Application of the ACO metaheuristic to stock-market investment

In the proposed algorithm, instead of ants there are investors who will take different investment paths, while dealing, from an initial to a final date. The first step includes the initializing of the values of the parameters which are taken into consideration in the algorithm. Amongst others, this means fixing the first pheromone trail associated with each transition, the number of investors in the market (m), and the weights defining the proportion in which they will affect the heuristic and memoristic information in the rule of probabilistic transition.

The main procedure of TRADINNOVA-ACO controls, through the constructor *schedule-activities*, the planning of the following three components:

- generating and bringing into operation the investors,
- the evaporation of pheromone, and
- the daemon actions.

On the other hand, the *initialize-investor-buy-sell-memory()* procedure takes care of specifying which stock is bought initially, as it is considered as the initial state at which the investor begins his or her path as well as storing the corresponding component in the memory of the investor *BS* (the buying-selling sequence to be carried out).

Finally, we must mention that the *compute-recommended-stocks* and *apply-decision-policy* procedures take into consideration the current state of the investor, the current values of the pheromone visible in the said node and the constraints of the problem, in order to establish the probabilistic transition process towards other valid states.

4.3 Making the prediction

The supposition on which we calculate the optimum stock dealing sequence, in other words, that which produces the greatest profit, is that we will be able to focus on where we have invested on the last day, with there being two possible cases: that the money is either invested in a specific stock or cashed in.

If the money is cashed in, this will be due to the fact that it is not currently of interest to invest in a stock, but if the money is invested in a particular stock, we can assume that this latest stock chosen is in an upward trend, meaning that if we maintain this investment, we will achieve the greatest profits from investing in the best possible stock.

In other words, our algorithm proposes buying the last stock found within the optimum stock buying-selling sequence.

5 Preliminary studies and experiments

We have tested TRADINNOVA making a simulation in the Spanish Continuous Market since years 2001 and 2004, having in account operating of the market, applying commissions and giving the orders of buying and selling limited to a price. In the simulation, the most representative stocks (between that they had official quotation from the 02 January 2001 to the 30 December 2004) of the Spanish continuous market are considering. Of the IBEX35, 28 stocks have been chosen and 9 stocks of the remainder of the market. By means of a file .INI the different parameters of input to the algorithm are specified.

The experiments made with TRADINNOVA and the definition of the parameters are shown in [1].

The obtained results have been made with a simulation in different periods, to have more empirical data, obtaining in general quite good results. In the Fig. 5 we show, for the different periods in which the simulation has been made, the revaluation obtained by the IBEX35 and with TRADINNOVA.




Fig. 5. IBEX35 vs TRADINNOVA

As can be seen in [1], the results obtained by this algorithm are extremely encouraging, but a pending subject is to improve the criteria for selection. With TRADINNOVA-ACO we propose a new criteria to select the stocks through the optimum sequence of buying and selling stocks.

This subject is the objective of the TRADINNOVA-ACO algorithm. We are in the phase of modification of the TRADINNOVA algorithm in order to adapt it to the ACO metaheuristic. The preliminary studies augur good results of TRADINNOVA-ACO.

6 Conclusions and future research

In this paper it has been shown that dynamic management of stocks is necessary, because for this TRADINNOVA has been implemented. The results obtained by this algorithm are extremely encouraging and can be seen in [1].

This algorithm leaves the stock selection process up to the investor, who can use any technique. This has made it possible for us to introduce the ant colony optimization metaheuristic as criteria for selection. With this metaheuristic we achieve the optimum stock dealing path which an investor would take between specific dates, something which will help us in the selection process.

Future processes would mean the implementation of modifications in TRADINNOVA in order to adapt it to this metaheuristic and check whether the predictions made surpass the reference index. We also propose the use of other types of metaheuristics for this objective and, therefore, the carrying out of a comparative study of them when achieving the optimum stock dealing path.

Acknowledgements

The authors thank the “Ministerio de Educación y Ciencia” of Spain and the “European Fund for Regional Development” (FEDER) for the partial support given to this work under the project TIN2005-08404-C04-02.

References

1. I. J. Casanova and J. M. Cadenas. TRADINNOVA-Un algoritmo heurístico de compra-venta inteligente de acciones. *CM-II 2006, Congreso Multidisciplinar de Percepción e Inteligencia*, Accepted and pending of publication, 2006.
2. M. Dorigo and G. Di Caro. The Ant Colony Optimization meta-heuristic. In D. Corne, M. Dorigo, and F. Glover, editors, *New Ideas in Optimization*, pages 11–32. McGraw Hill, 1999. London, UK.
3. M. Dorigo, G. Di Caro and L. M. Gambardella. Ant algorithms for discrete optimization. *Artificial Life*, 5 (2): 137–172, 1999.
4. O. Cordón, F. Herrera and T. Stützle. A Review on the Ant Colony Optimization Metaheuristic: Basis, Models and New Trends. *Mathware and Soft Computing* 9(2-3):141–175, 2002.
5. E. P. K. Tsang and S. Martinez-Jaramillo. Computational Finance. Feature Article. *IEEE Computational Intelligence Society*, 2004.
6. K. Swingler. Financial Prediction, some pointers, pitfalls, and common errors. *Neural Computing and Applications*, 4(4):192–197, 1996.
7. M. Zekic. Neural Network Applications in Stock Market - A Methodology Analysis. *9th International Conference on Information and Intelligent Systems '98*, pages 255-263, 1998.
8. M. Januskevicius. Testing Stock Market Efficiency Using Neural Networks Case of Lithuania. *Thesis. SSE Riga Working Papers*, 17(52), 2003. http://www2.sseriga.edu.lv/library/working_papers/FT_2003_17.pdf

9. E. P. K. Tsang and J. Li. Combining Ordinal Financial Predictions with Genetic Programming. *IDEAL-2000, Second International Conference on Intelligent Data Engineering and Automated Learning*, pages 13–15, 2000.
10. A. E. Drake. Genetic Algorithms in Economics and Finance: Forecasting Stock Market Prices and Foreign Exchange - A Review. <http://www.agsm.unsw.edu.au/bobm/papers/drake.pdf>
11. L. Davis. Trading on the Edge: Neural, Genetic, and Fuzzy Systems for Chaotic Financial Markets. In G. J. Deboeck, editor, *PART II: 8. Genetic Algorithms and Financial Applications*. Wiley Finance Edition, 1994.
12. P. Kroha and R. Baeza-Yates. Classification of Stock Exchange News. *Technical Report. Department of Computer Science, Engineering School, Universidad de Chile*, 2004. <http://www.tu-chemnitz.de/informatik/service/if-berichte/pdf/CSR-04-02.pdf>
13. M. A. Mittermayer. Forecasting Intraday Stock Price Trends with Text Mining Techniques. hicss, p. 30064b. *HICSS'04, 37th Annual Hawaii International Conference on System Sciences - Track 3*, 2004.
14. C-C. Tseng. Comparing Artificial Intelligence Systems for Stock Portfolio Selection. *The 9th International Conference of Computing in Economics and Finance*, 2003.
15. G. Kendall and Y. Su. A Particle Swarm Optimization Approach in the Construction of Optimal Risky Portfolios. *The 2005 IASTED International Conference on Artificial Intelligence and Applications*, pages 40–45, 2005.
16. M. Rubinstein. Markowitz's "Portfolio Selection": A Fifty-Year Retrospective. *Journal of Finance*, 57(3):1041–1045, 2002.
17. A. Borodin, R. El-Yaniv and V. Gogan. Can We Learn to Beat the Best Stock. *Journal of Artificial Intelligence Research*, 21:579–594, 2004.