

Learning with Imperfections – A Multi-Agent Neural-Genetic Trading System with Differing Levels of Social Learning

Graham Kendall

School of Computer Science and IT
ASAP Research Group
University of Nottingham
Nottingham, NG8 1BB
gxk@cs.nott.ac.uk

Yan Su

School of Computer Science and IT
ASAP Research Group
University of Nottingham
Nottingham, NG8 1BB
yxs@cs.nott.ac.uk

Abstract—Some real life dynamic systems are so large and complex that the individuals inside the system can only partially understand their environment. In other words, the dynamic environment is imperfect to its participants. In this paper, by using the stock market as a test bed, we demonstrate an integrated individual learning and social learning model for optimisation problems in dynamic environments with imperfect information. By applying differing levels of social learning process in an evolutionary simulated stock market, we study the importance of social learning on the adaptability of artificial agents in imperfect environments. Comparisons between the integrated individual and social learning model and other evolutionary approaches for dynamic optimisation problems, particularly the memory-based approaches and multi-population approaches, are also drawn with the emphasis on optimisation problems with imperfect information.

Keywords—*imperfect environments; neural-genetic system; stock market; individual learning; social learning.*

I. INTRODUCTION

Evolutionary algorithms have been widely used for solving optimisation problems in dynamic environments. A dynamic environment is generally characterised by a multimodal and non-static fitness space. Therefore, the aim of an evolutionary algorithm is not only just to find optimal solutions, but also to continuously adapt the solutions to a changing environment [1, 2]. However, one of the important characteristics of dynamic environments has generally been neglected by researchers; that is the imperfectness of dynamic environments. By imperfectness, we mean the individuals in the system can only partially understand their environment due to the complexity of the dynamic system or the inability of the participants to perceive the entire environment. In other words, a different view of the information within the environment is perceived by different individuals. Therefore, the environment consists of a number of different search spaces with different dimensions. These search spaces can overlap, or be completely independent. A good example is the stock market. The theory on the economics of imperfect information points out that different participants in a financial market possess different information from the market [3, 4]. In the stock market, a large number of stock traders and investors

consider different views of the market and use different trading strategies to make trading decisions. Although the market is imperfect to stock traders and investors, due to the limited information they gather from the market, there are still large number of investors who make a profit from the stock market. It is due to the imperfectness and the efficiency of the stock market [5], that some investors can make profits over other investors with inferior information.

In our previous work [6, 7], we developed a multi-agent based simulated stock market model, where artificial stock traders, modelled using Artificial Neural Networks (ANN), co-evolve with each other by the means of an integrated individual and social learning algorithm. The experiments from [6, 7] demonstrated the artificial stock traders learned to develop successful stock trading strategies in an imperfect simulated stock market. In this paper, we centre the study on the integrated individual and social learning algorithm employed in the simulated stock market, from the perspective of learning with imperfect information in dynamic environments, in particular, the impact of differing levels of social learning in an evolutionary system on the adaptability of agents in imperfect dynamic environments. In section II, we discuss the advantages of the integrated individual and social learning algorithm and draw comparisons to other evolutionary approaches for dynamic optimisation problems. Section III describes the simulated stock market and the integrated individual and social learning algorithm. Section IV demonstrates experiments on the simulated stock market with differing levels of social learning. Conclusions and future work is presented in section V.

II. OPTIMISATION IN DYNAMIC ENVIRONMENTS

One problem with applying evolutionary algorithms to dynamic optimisation problems is that they will eventually converge to a local optimum in the search space, and subsequently lose their adaptability when the underlying environment changes. Therefore, most evolutionary approaches endeavour to help the evolutionary algorithms escape from the local optimum and start a new search in the changed fitness space. See [2] and [8] for a good review on the evolutionary approaches to dynamic optimisation problems and [9] for some recent trends.

The integrated individual and social learning algorithm employed in our simulated stock market model [6, 7] essentially mimics human learning behaviours in human societies. Every person in a society attempts to maximise their own utility with the resources available to them by means of individual and social learning. In the face of a changing environment, social learning plays an important role in enabling an individual to adapt to new environments by learning from other better-adapted individuals [10, 11]. Inside our simulated stock market, whilst every artificial stock trader evolves their own individual trading strategy, allowing them to trade more productively utilising a (possibly) unique set of market information, the market itself also serves as a repository of good trading strategies and the knowledge is disseminated among the traders over time. The social learning mechanism enables the traders to explore different search spaces that are constructed on different information sets within the environment and therefore solves the problem of a market which can only be partially understood by its participants due to its imperfectness as a dynamic system.

The two evolutionary approaches discussed in the review [2], an explicit memory-based approach and a multi-population approach, also employ the idea of saving the past experiences in a social memory. Explicit memory-based approaches store good solutions from previous generations in a memory and reintroduce them later into the search population [1, 12, 13]. Thus the diversity of the search population is maintained. Such approaches suit environments where changes occur periodically, i.e., environments whose optimum keeps returning to a previous point in the search space. Bendtsen and Krink [14] took a further step by using a dynamic memory that could be adjusted to the changes in the environment by means of moving externally stored candidate solutions gradually towards the currently nearest best genomes in the search population. Similarly, the multi-population approach [15] created a self-adaptive memory by maintaining small subpopulations for some promising areas in the search space. The main differences and advantages of the integrated individual and social learning algorithm employed in the simulated stock market, compared with the above two approaches, lie in:

- Most importantly, the above two approaches do not answer the question about the imperfectness of a dynamic environment. When a dynamic environment has imperfect information, the environment consists of a number of different search spaces. Each of these search spaces has a different search dimension depending on the different sets of information perceived from the environment. These search spaces may overlap, or may be completely independent from each other. As an example, in the simulated stock market, as is in the real market, every stock trader uses a different information sets from the market for making trading decisions. The above two approaches enable the evolutionary algorithms jump from one area to another, new, promising area in the same search space, but they do not solve the problem of moving from one search space to another search space with different information. The integrated individual and social

learning algorithm in the simulated stock market solves this problem by modelling the market as an imperfect environment and enabling individuals to learn from others who use different information sets.

- Since the environment is non-static, the search spaces in the environment are also non-static. A current search space may disappear from the environment because no one believes that the information is valuable. A new search space may appear because the information has been discovered by someone. As an example, in our simulated stock market, a trader may decide to discard his current strategy, and consider a new set of market indicators from the market for developing new trading strategies.
- The explicit memory-based approach and multi-population approach both seem very similar to the concept of social learning. However, these two approaches are still learning individually from previous experience. For example, in our simulated stock market, the artificial traders are modelled as heterogeneous artificial agents with different minds. Before each trader publishes his trading strategy to the society, he will not only examine his own performance, but also his relative performance compared with other traders who use different information.

In this paper, we change the popularisation of the social learning in the evolutionary market to different extents, by applying different pressures on artificial traders to take part in the social learning. Results from the experiments show that there is a significant impact of the social learning on the adaptability of the artificial stock traders and demonstrates the effectiveness of the integrated individual and social learning algorithm in solving learning problems with imperfect information.

III. THE INDIVIDUAL AND SOCIAL LEARNING ALGORITHM

Our simulated stock market is a neural-genetic stock trading system. Basic market information, such as stock prices and trading volumes, are given extraneously. Artificial stock traders use Artificial Neural Networks (ANN) to detect buy and sell signals from the market, and carry out individual learning by means of a Genetic Algorithm (GA). The following describes the general structure of the trading system:

1. Before trading starts, there are 50 active traders in the simulated stock market. There are 20 technical and market indicators and zero trading strategies in a central pool. The 20 indicators are assigned an equal score of 1. Each trader selects a set of market indicators randomly using roulette wheel selection.
2. With the set of indicators selected, each trader generates ten different models in the form of Artificial Neural Networks as trading strategies. These ten models may have different network architectures, but they use the same set of indicators selected by the trader. The aim is for the trader to evolve better trading models by means of *individual learning*.

3. The time span of the experiment is divided into equal intervals. Each interval contains 125 trading days (6 months trading).
4. Each 125-day trading period is sub-divided into intervals of 5 days. After each 5-day trading period, an individual learning is undertaken by means of a Genetic Algorithm (GA).
5. At the end of each 125-day trading, *social learning* occurs and each trader is given the opportunity to decide whether to look for more successful strategies from the pool or whether to publish his/her successful strategies into the central pool depending on two thresholds θ_1 and θ_2 .
6. After social learning, the system enters the next 125-day trading period and steps 4, 5 and 6 are repeated.

A. Individual Learning

As described above, individual learning occurs every 5 days at which time the trader will calculate the trading model's rate of profit (ROP) by using (1), where W is the trader's current assets (cash + shares). W' is the trader's assets one week before.

$$ROP = \frac{W - W'}{W'} \times 10 \quad (1)$$

The ROP is used to describe a trading model's profitability. A Genetic Algorithm is used to eliminate trading models with poor performance and introduce new models by mutating network connections and changing network architectures. The pseudocode of the GA algorithm is presented as follows.

```

Select a model with the lowest ROP to
  be eliminated;
Select a model to be mutated using
  roulette selection;
Decide number of connections to be
  mutated, m;
i = 0;
While(i < m) {
  Randomly select a connection;
  Weight = weight + Δw;
  i = i + 1; }
With 1/3 probability add a hidden node;
With 1/3 probability delete a hidden
  node;
Replace the model to be eliminated with
  the new mutated model;

```

Where m is a random integer between 0 and the total number of connections in the selected neural network. Δw is a random Gaussian number with a mean of zero and standard deviation of 0.1.

B. Social Learning

Social learning occurs at the end of every 6-month trading period through a self-assessment process. The self-assessment calculates how well the trader has performed in terms of his own profitability and his relative performance to other traders. The self-assessment process uses (2), (3), and (4).

$$S_{peer}^i = \frac{R_i}{49} \quad (2)$$

First, the trader's rate of profit (ROP) for the past six months is calculated by using (1), and the 50 traders are ranked from 0 to 49 (R_i) according to their ROP. Equation (2) gives each trader a score in terms of peer pressure from other traders. In other words, this score shows trader i 's performance compared to other traders.

$$S_{self}^i = \frac{ROP - ROP'}{100} \quad (3)$$

ROP' is the rate of profit for the previous six months. Equation (3) gives the trader's score in terms of his own performance in the past six months compared to the previous six months. Finally, these two types of performance are composed into (4), which gives the final assessment (σ_i) for trader i .

$$\sigma_i = S_{peer}^i + \frac{1}{1 + e^{(1 - S_{self}^i)}} \quad (4)$$

The final assessments for 50 traders are then normalised into the range of [0,1]. In order to study the impact of the social learning on the learning abilities of artificial stock traders with imperfect information, we extend the social learning mechanism with four different parameter settings that force different pressures on artificial traders in taking part in social learning. An arithmetic mean value (Φ) of all 50 traders' normalised final assessments (σ_i) is calculated using (5).

$$\Phi = \frac{1}{N} \sum_{i=1}^N \sigma_i \quad (5)$$

N is the total number of artificial traders, which equals to 50. A trader's social behaviour will now fall into four different categories depending on the values of σ_i and Φ , and two thresholds, θ_1 and θ_2 (see section IV for the description of the thresholds θ_1 and θ_2). The four cases, i.e. four different types of social learning behaviours, are:

CASE1: The trader is successful and he is not using a strategy learned from the market. The trader will publish his novel trading strategy into the central pool and enter the next six months trading using the same strategy.

CASE2: The trader is successful and is using a strategy learned from the market. He will not publish the strategy again, but update this strategy's score in the pool using their six-month ROP (1). The trader will then enter the next six months trading using the same strategy.

CASE3: The trader is not satisfied with his performance in the last six-months trading. The trader will have 0.5 probability of copying a strategy from the central pool, which means the trader will discard whatever model he is using, and select a better trading strategy from the pool using roulette selection, and go into the next six months trading with this copied strategy. Or, with 0.5 probability, the trader will decide to discard whatever strategy he is using, and select another set of indicators as inputs, build 10 new models and go into the next six months trading with the new trading models.

CASE4: The trader is satisfied with his performance in past six months and continues using that strategy.

We only provide a general description of the trading system here to ensure the readability of the paper. For more details please refer to [6, 7] for more details.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

Five stocks from the Hong Kong Stock Exchange and the Tokyo Stock Exchange are selected to be traded in the simulated stock market: CHEUNG KONG (0001.HK), WHARF HOLDINGS (0004.HK), CATHAY PAC AIR (0293.HK), TOYOTA INDUS CORP (6201.JP) and SONY CORP (6758.JP). The simulated stock market was tested on each of the five stocks. The four different settings with differing level of social learning are described as following.

SETTING1 – Social learning is turned off while only individual learning occurs. Each trader in the simulated stock market evolves independently from each other with different sets of information. Each trader can only search in his own search space that is defined by the information he selected from the environment.

SETTING2 – Both individual learning and social learning are turned on. The threshold θ_1 for social learning is set to 1. The threshold θ_2 is set to 0.9. Traders' behaviour during the social learning is described in table I:

TABLE I. TRADERS' ACTIONS DURING THE SOCIAL LEARNING UNDER SETTING 2. SEE SECTION III FOR DESCRIPTION OF CASE 1 TO 4.

Parameters	Trader's Action
$\sigma_i = \theta_1$	CASE1/CASE2
$\sigma_i < \theta_2$	CASE3
$\theta_1 > \sigma_i \geq \theta_2$	CASE4

Setting2 mimics an environment where only the best players are accepted to distribute their knowledge, and strong motivations are forced on individuals to learn from each other.

SETTING3 – Both individual learning and social learning are turned on. The threshold θ_1 for social learning is set to 1. The threshold θ_2 is set to the mean value Φ (See section III). Traders' behaviour during the social learning is described in table II:

TABLE II. TRADERS' ACTIONS DURING THE SOCIAL LEARNING UNDER SETTING 3. SEE SECTION III FOR DESCRIPTION OF CASE 1 TO 4.

Parameters	Trader's Action
$\sigma_i = \theta_1$	CASE1/CASE2
$\sigma_i \leq \theta_2$	CASE3
$\theta_1 > \sigma_i > \theta_2$	CASE4

Setting3 creates an environment where only the best players are accepted but forces less strong motivations on individuals to learn from each other compared with setting1.

SETTING4 – Both individual learning and social learning are turned on. The parameter θ_1 for social learning is set to 0.9. The parameter θ_2 is set to the mean value Φ (See section III). Traders' behaviour during the social learning is described in table III:

TABLE III. TRADERS' ACTIONS DURING THE SOCIAL LEARNING UNDER SETTING 4. SEE SECTION III FOR DESCRIPTION OF CASE 1 TO 4.

Parameters	Trader's Action
$\sigma_i > \theta_1$	CASE1/CASE2
$\sigma_i \leq \theta_2$	CASE3
$\theta_1 \geq \sigma_i > \theta_2$	CASE4

Setting4 mimics an environment where more individuals have the opportunity to distribute their knowledge to the society while the learning atmosphere within the society is moderate.

The experimental results are depicted in Fig. 1 to compare the algorithm where no social learning occurs with the algorithms with differing levels in social learning. All results are taken from a single run on each stock under the four different settings. We compare the traders' performance from each simulation with two benchmarks: bank savings and buy-and-hold strategy. Bank savings means the trader invests the same amount of money in the bank throughout the whole trading period, receiving interest from the bank. Buy-and-hold means the trader keeps his entire asset in a particular stock and holds it until the end of the trading period with the hope to make a profit through the appreciation of the stock.




Figure 1. Comparison between the algorithm without social learning (SETTING1) and the algorithms with different pressures on social learning (SETTING2, 3, 4). On all the X axes, 1 refers to SETTING1, 2 refers to SETTING2, 3 refers to SETTING3, and 4 refers to SETTING4 (see section). (a) – Number of traders outperformed the bank savings. (b) – Number of traders outperformed the buy and hold strategy. (c) – The cumulative total return of the best trader from the 50 traders. (d) – The average cumulative total return of all 50 traders.

Fig. 1, A(a) to E(a) shows the number of traders who outperformed bank savings under the four different settings on a particular stock. A(b) to E(b) shows the number of traders who outperformed the classic buy-and-hold strategy. A(c) to E(c) shows the cumulative total returns of the best traders under four different settings for the particular stock. A(d) to E(d) show the average cumulative total returns of 50 traders. It is clear from Fig. 1 that the agents take no part in social learning, i.e., SETTING 1, generally performed poorly when compared to traders where social learning was allowed. By looking at A(c), B(c), C(c), D(c) and E(c), there is always an individual and social learning setting that helped traders to find better trading strategies, rewarding the trader with higher returns. A reasonable explanation is that the social learning process enabled the traders to escape from non-promising search spaces which are limited by the set of information the trader perceived from the imperfect environment, and enter other promising search spaces which have been well explored by others in the environment, or exploit new set of information which have not been used by other participants in the market. As we discussed in section II, because the stock market is a dynamic environment with imperfect information, the optimisation problem in such an environment is not just to find one good solution at a time, but to also adapt to the changed market environment. Social learning enhances an agent's adaptability by enabling agents to learn successful trading strategies from others who are better adapted by making use of different information. On the other hand, the results also demonstrate the effectiveness of the integrated individual and social learning algorithm in solving learning problems with imperfect information.

Concerning SETTING 2, 3, and 4, i.e., scenarios where the social learning is applied to different extents in the simulated stock market, by examining B, C, D, E from Fig. 1, we can see SETTING 2, where only the best players are accepted to distribute knowledge and strong motivations are forced on individuals to learn from each other, generally recorded better performance cross four criteria a, b, c and d. On the CHEUNG KONG stock, SETTING 4 seems to perform better than SETTING 2 and 3, but SETTING 2 still recorded the highest return from the best trader in A(c). Compared with SETTING 2, SETTING 4 has a moderate pressure on traders in taking part in the social learning but gives more individuals the opportunities to distribute their knowledge to the society. Both SETTING 2 and 4 strengthen the social learning in different ways. It is clear that strengthening the social learning improves the adaptability of artificial traders in the imperfect simulated stock market.

V. CONCLUSIONS AND FUTURE WORK

The integrated individual learning and social learning algorithm addresses the question of searching in non-static search spaces in an imperfect environment where different sets of information are perceived by different individuals, which is generally neglected by other evolutionary approaches in dynamic optimisation problems. The results from the experiments demonstrate that social learning plays an important role in ensuring the adaptability of agents in an

imperfect dynamic environment. The integrated individual and social learning algorithm presents a way of solving learning problems with imperfect information. When two algorithms are compared, the purpose is either to find a better generic algorithm, or to find a better algorithm that is more suitable for a certain problem. From the authors' point of view, we mean the latter. When we compared the integrated individual and social learning algorithm with other evolutionary approaches for dynamic optimisation problems, we stressed the precondition is a dynamic environment with imperfect information. There are many dynamic environment optimisation problems that come with perfect information, and can be handled by other evolutionary approaches. For our future work, we intend to study the impact of the frequency of individual and social learning on the imperfect evolutionary system and the absorption and dissemination of new information from an imperfect environment.

REFERENCES

- [1] J. Branke, "Memory enhanced evolutionary algorithms for changing optimization problems", in *Proc. of Congress on Evolutionary Computation (CEC'03)*, pp. 1875-1882, 1999.
- [2] J. Branke, "Evolutionary approaches to dynamic environments – updated survey", in J. Branke, ed, *GECCO Workshop on Evolutionary Algorithms for Dynamic Optimization Problems*, pp. 27-30, 2001.
- [3] M. Rothschild and J. E. Stiglitz, "Equilibrium in competitive insurance markets: an essay on the economics of imperfect information", *The Quarterly Journal of Economics*, vol. 90, no. 4, pp. 630-649, 1976.
- [4] R. Arnott, B. Greenwald, R. Kanbur, and B. Nalebuff, *Economics for an imperfect world – essays in honor of Joseph E. Stiglitz*, The MIT Press, Cambridge Massachusetts, 2003.
- [5] E. F. Fama, "Efficient capital markets: A review of theory and empirical work", *Journal of Finance*, May: 383-417, 1970.
- [6] G. Kendall and Y. Su, "Co-evolution of successful trading strategies in a multi-agent based simulated stock market", in *Proc. of The 2003 International Conference on Machine Learning and Applications (ICMLA'03)*, pp. 200-206, 2003.
- [7] G. Kendall and Y. Su, "A multi-agent based simulated stock market – testing on different types of stocks", in *Proc. of The Congress on Evolutionary Computation (CEC'03)*, pp. 2298-2305, 2003.
- [8] J. Branke, *Evolutionary Optimization in Dynamic Environments*, Kluwer, 2001.
- [9] J. Branke, "Evolutionary approaches to dynamic optimization problems - introduction and recent trends", in J. Branke, ed, *GECCO Workshop on Evolutionary Algorithms for Dynamic Optimization Problems*, pp. 2-4, 2003.
- [10] A. Bandura, *Social Foundations of Thought and Action: A Social Cognitive Theory*, Englewood Cliffs, NJ: Prentice Hall, 1986.
- [11] S. Mineka and M. Cook, "Social learning and the acquisition of snake fear in monkeys", in T. R. Zentall, et al., eds, *Social Learning: Psychological and Biological Perspectives*. Lawrence Erlbaum Associates, pp. 51-73, 1988.
- [12] S. J. Louis and Z. Xu, "Genetic algorithms for open shop scheduling and re-scheduling", in M. E. Cohen and D. L. Hudson, eds, *ISCA 11th Intl. Conf. on Computers and Their Applications*, pp. 99-102, 1996.
- [13] N. Mori, S. Imanishi, H. Kita, and Y. Nishikawa, "Adaptation to changing environments by means of the memory based thermodynamical genetic algorithm", in *Proc. of The Seventh Intl. Conf. on Genetic Algorithms*, pp. 299-306, 1997.
- [14] C. N. Bendtsen and T. Krink, "Dynamic memory model for non-stationary optimization". in *Proc. of Congress on Evolutionary Computation*, pp. 145-150, 2002.
- [15] J. Branke, T. Kaubler, C. Schmidt, and H. Schmeck, "A multi-population approach to dynamic optimization problems", in *Adaptive Computing in Design and Manufacturing*, Springer, 2000.