Neural-Evolutionary Learning in a Bounded Rationality Scenario

Ricardo Matsumura de Araújo and Luís C. Lamb

Institute of Informatics
Federal University of Rio Grande do Sul
Porto Alegre, 91501-970, Brazil
{rmaraujo,lamb}@inf.ufrgs.br

Abstract. This paper presents a neural-evolutionary framework for the simulation of market models in a bounded rationality scenario. Each agent involved in the scenario make use of a population of neural networks in order to make a decision, while inductive learning is performed by means of an evolutionary algorithm. We show that good convergence to the game-theoretic equilibrium is reached within certain parameters.

1 Introduction

Classical economics makes the assumption that economic agents involved in any system are perfectly rational. This usually means that each agent has knowledge of all relevant aspects of the environment, a logical and coherent preference, and enough computational power to process all this information in order to choose the best course of action to attain the highest, optimal point in his or her preference scale [13]. This assumption facilitates the use of analysis tools within game theory to predict the outcome of interactions of multiple agents. However, it is not the case that real agents behave in a perfect rational sense: they are usually endowed with *bounded rationality* [12].

This term, coined by Herbert Simon in the 1950s, refers to an approach now widely used for modeling reasoning in economic scenarios. Systems endowed with bounded rational agents might offer quite different behavior from ones with rational agents. Some systems, on the other hand, seem to allow analysis *as if* agents were rational [11], even though this may not be the case. Such systems are of particular interest, since they allow the use of traditional, tractable techniques for behavioral analysis [6].

Arthur [1] has proposed *the El Farol problem* in order to provide insights in systems of interacting bounded rational agents in a simplified market model. Since then, this model has been widely discussed together with other evolutionary game scenarios, such as the Minority Game [16]. The El Farol problem is as follows. *There are N agents; each agent has to decide whether or not to go to the El Farol Bar at some week; an agent will go if he or she expects that at most aN agents are going, where $a \in [0, 1]$; otherwise the bar would be overcrowded and the agent will not go. The only source of information available to the agents is a global history of past weeks attendance, and no explicit communication is allowed between the agents*[1]. The interest in modeling this

[1] Some papers report experiments with explicit communication, see [14] for an example.

kind of problem is the assumption that the utility gain of each agent depends directly on the decision of all other agents, leading to an interesting paradox: if many agents believe that El Farol will be overcrowded, few will go; if they all believe nobody will go, all will go.

In [1], an experiment was set by allowing each agent to choose over an internal set of fixed strategies in order to predict the next week attendance. Each strategy used only a history window to make the prediction (e.g. the same as the previous week, the average over 5 last weeks, a fixed number etc.) and the most accurate predictor was chosen at each time step. With a total of 100 agents and setting the bar to be overcrowded if 60 or more agents attended ($a = 0.6$), the simulations showed a convergence to a mean attendance of 60. Game theory tells us that a Nash equilibrium exists using mixed-strategy and 60 is the expected mean attendance [1], the same results observed in Arthur's simulations.

In response to this experiment, [6] suggested that if agents could be *creative*, in the sense that they could come up with *new* strategies other than those pre-specified (as was the case in Arthur's simulations), the system would not show a convergence, but rather it would behave in a more chaotic manner, thus showing that game-theoretic expectation would be of little use to predict the behavior of such systems. Learning was conducted using an evolutionary paradigm, where each strategy is represented by an autoregressive equation, as a function of past weeks attendances, which parameters were mutated to give birth to new, possibly better, generations of predictors. In the present paper, we refer to "dynamic learning" as the model of learning that allows the creation of new models, in contrast to "static learning" where learning is only made on choices over a fixed set of pre-specified models. Even though in [6] dynamic learning is explicitly included in the model, the predictions made are only able to capture linear patterns, which is a rather strong assumption on the system.

We believe that it is plausible that the study of other classes of machine learning algorithms applied to the same problem may show a different behavior, possibly closer to the real world market behavior that the problem tries to model. In this paper we present empirical results from the simulation of the El Farol problem using agents capable of dynamic learning through a population of neural networks, evolved by an evolutionary algorithm. By doing this, we aim to further understand the role of learning in the dynamics of economic scenarios and multi-agent systems in general. We then show that the system shows a better convergence to the game-theoretic equilibrium when compared to the proposed setup in [6].

Section 2 provides the architecture used to model the agents; Section 3 presents detailed results of the experiments; Section 4 the concludes the paper and discusses directions of future work.

2 The Hybrid Agent Architecture

An agent is defined as a system that receives a vector $\mathcal{A}$ of length M, representing a history window of past weeks attendance, and outputs a single bit of information, namely "0" if it will not attend to the bar in the current week and "1" if it will. The vector $\mathcal{A}$ is here considered external to the agent since it is meant to be a perfect information available to every agent. We do not consider the case where agents have different perceptions on this information. Next, we describe the composition of an agent.

In multi-agent decision problems it is common the use of the concept of "mental models" [1, 4, 3], where each agent has a population of predictors and one is chosen to make the decision at each simulation step. This concept is used in this experiment as well. Each agent is internally equipped with a population of K models, represented as neural networks, chosen due to its ability to capture non-linear patterns [7]. Neural networks have been widely used in economic modeling (see e.g. [8]). In [9] neural networks were applied to the Minority Game, a variant of the El Farol Problem, but agents were composed of a single neural network, differing from the approach taken here, and [2] has presented an application of genetic learning over neural networks within an alternative economic context.

We use multi-layer perceptrons (MLP) [7] composed of an input layer, one hidden layer and one single output unit. The input layer is composed of M input units, each receiving one unique value from $\mathcal{A}$. This way the number of input units effectively represents the agents memory size. The hidden layer is composed of H units. The value of H roughly controls the capacity of the neural network to process the information received [15] and is the same to every neural network in the system. The ouput of the output unit is taken as a prediction of attendance based on the input. All nodes in a layer are fully connected to the next layer and use a sigmoidal activation function. Fig. 1 shows the topology adopted.

All networks have their prediction accuracy evaluated, at each simulation step, through a *fitness function* and the best performing network is chosen by each agent to make its current week's prediction. The agent's decision is then: Output "0" if predicted attendance is greater than aN; Output "1" otherwise.

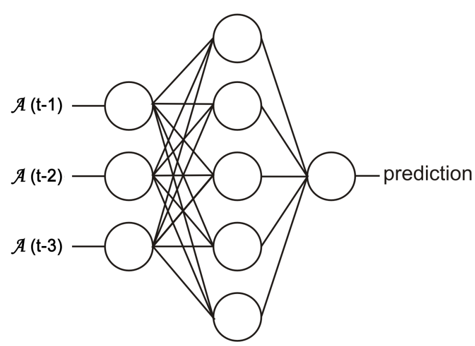

Fig. 1. Example topology for $M = 3, H = 5$

On top of the population of predictors runs an evolutionary algorithm. The model of learning used here closely follows the one described by Fogel in [5]. For every agent, the algorithm generates one offspring for each neural network by adding a zero-mean random gaussian value $\sigma \in [-1, 1]$ to each weight i:

$$w^i_{offspring} = w^i_{parent} + \sigma$$

The operation results in a total of $2K$ neural networks. All of them are evaluated with the fitness function and the best K replace the current population. Through this proceeding, it is ensured that new strategies are created all the time and put in competition with previous ones, generating possibly better ones.

3 Simulation Setup and Results

The simulations described here were made using the following parameters: $N = 100$, $K = 10$, $M = 12$, $H = 50$, $a = 0.6$. Except for H, which was chosen for having presented good results in trials, all other parameters follow from those used in [1] and [6]. The fitness function of the system was taken as being the sum of squared errors of tests through 10 different history windows.

Figure 2 (a) shows a typical attendance over 500 weeks *without* the evolutionary learning applied. Thus, each agent can only choose among the fixed set of models created at the start of the simulation. It is interesting to note that randomly initialized neural networks do not contain any explicit strategies as was the case in Arthur's model. Despite this, the system's behavior is very close to that of [1], showing a mean convergence near Nash equilibrium of 60 with minor fluctuations, with a mean of 59.02 and standard deviation of 4.13. These values were calculated in the "steady-state" region, where the transient fluctuations due the random initialization are surpassed, which we take as starting around week 50.

By allowing the agents to learn using the evolutionary algorithm we reach the typical results depicted in Fig. 2 (b). Despite a small increase to 4.25 in standard deviation, convergence is still observed. The mean attendance, after the transient stage, is 58.98. We show that agents are effectively learning in Fig. 3, where the average fitness of the agents is shown over 500 weeks of a typical simulation.

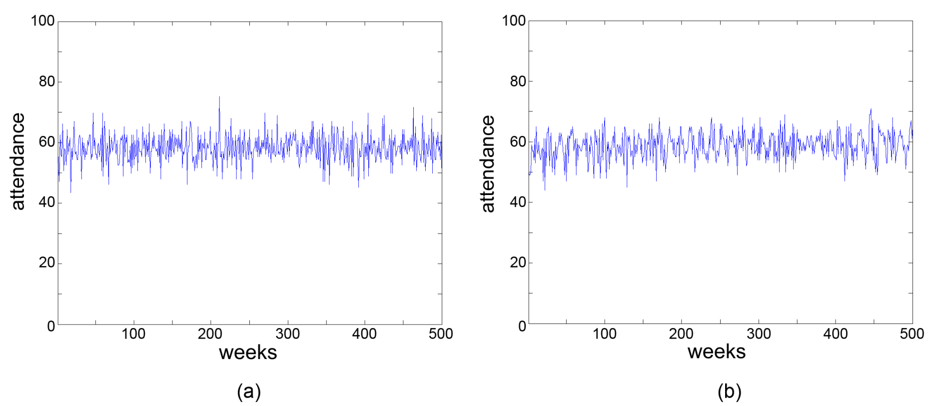

Fig. 2. Typical weekly attendance: (a) without evolutionary learning; (b) with evolutionary learning.

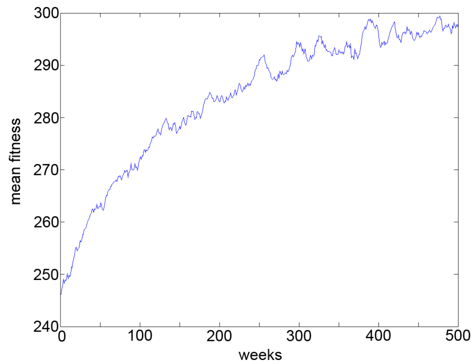

Fig. 3. Mean fitness averaged over all 100 agents at each week.

Even though the results with static learning are very similar to Arthur's in [1], the results with the evolutionary (dynamic) learning applied differ qualitatively from those presented in [6]. The mean attendance gets closer to the game-theoretic expectation and standard deviation in a typical trial is much smaller. In fact, although dynamic learning is clearly taking place, the system presents almost identical behaviour compared to the case where only static learning is acting. Thus, the overall behavior is predicted in both cases, to some extent, by classical game-theory. This better converging behavior might be explained by the use of neural networks, which are known to be able to capture non-linear patterns, in contrast to the linear predictors used in [6]. It is known that the behaviour of this kind of evolutionary game is very dependent on the memory size of the agents [10], being important to notice that this comparisons were made using the same memory size.

4 Conclusions and Future Work

In this paper we proposed a neural-evolutionary model for learning in a bounded rationality scenario. We illustrated the use of our approach by means of an application to the well-known El Farol Problem, where each agent is equipped with a population of neural networks as predictors, which learns by induction through an evolutionary algorithm. Empirical results showed a good convergence to the game-theoretic expected equilibrium.

The results presented here show that, in spite of the use of dynamic learning, the problem does not necessarily present chaotic behavior as suggested in [6]. This corroborates the hypothesis that the underlying learning paradigm seems to play a substantial role in the rate and stability of convergence in our case study and, possibly, in evolutionary games in general. Although we make no claims that the behavior shown here represents real market behavior, these results are important to better understand the role of learning in bounded rationality scenarios and multi-agent systems in general.

Our experiments have also suggested that the stability (mean and standard deviation) is highly influenced by the *complexity* of the neural networks i.e. simulations using networks of different sizes (number of nodes and layers) presented qualitatively different behaviors. As future work we plan to study such variations by analysing the role of complexity and computational power in the emergent properties of evolutionary games.

Acknowledgements: This work has been partly supported by CNPq and FAPERGS.

References

1. W. Brian Arthur. Inductive reasoning and bounded rationality. *American Economic Review (Papers and Proceedings)*, 84:406–411, 1994.
2. Gianluca Baldassarre. Neural networks and genetic algorithms for the simulation models of bounded rationality theory - an application to oligopolistic markets. *Rivista di Politica Economica*, 12:107–146, 1997.
3. Bruce Edmonds. Modelling bounded rationality in agent-based simulations using the evolution of mental models. In T. Brenner, editor, *Computational Techniques for Modelling Learning in Economics*, pages 305–332. Kluwer, 1999.
4. Bruce Edmonds and Scott Moss. Modelling bounded rationality using evolutionary techniques. *Proceedings AISB'97 workshop on Evolutionary Computation*, pages 31–42, 1997.
5. David Fogel. *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*. IEEE Press, 2a. edition, 2000.
6. David Fogel, Kumar Chellapilla, and Peter Angeline. Inductive reasoning and bounded rationality reconsidered. *IEEE Transactions on Evolutionary Computation*, 3(2):142–146, July 1999.
7. Simon Haykin. *Neural Networks: A Comprehensive Foundation*. Prentice Hall, 2nd edition, 1998.
8. Ralf Herbrich, Max Keilbach, Thore Graepel, Peter Bollmann-Sdorra, and Klaus Obermayer. Neural networks in economics. *Advances in Computational Economics*, (11):169–196, 1999.
9. W. Kinzel, R. Metzler, and I. Kanter. Dynamics of interacting neural networks. *Physica A*, 33(14):141–147, April 2000.
10. Esteban Moro. The minority game: an introductory guide. In Elka Korutcheva and Rodolfo Cuerno, editors, *Advances in Condensed Matter and Statistical Physics*. Nova Science Publishers, Inc., 2004.
11. Scott Moss. Boundedly versus procedurally rational expectations. In H. Hallet and P. McAdam, editors, *New Directions in Macro Economic Modelling*. Kluwer, 1997.
12. Ariel Rubinstein. *Modeling Bounded Rationality*. Zeuthen Lecture Book Series. The MIT Press, Cambridge, Massachussets, 1998.
13. Herbert Simon. A behavioral model of rational choice. *The Quarterly Journal of Economics*, LXIX, February 1955.
14. Frantisek Slanina. Social organization in the minority game model. *Physica A*, (286):367–376, 2000.
15. Joseph Wakeling and Per Bak. Intelligent systems in the context of surrounding environment. *Physica Review E*, 64(051920), October 2001.
16. Yi-Cheng Zhang. Modeling market mechanism with evolutionary games. *Europhysics News*, March/April 1998.