

Spatial evolutionary game-theoretic perspective on agent-based complex negotiations

Siqi Chen,¹ Jianye Hao,² Gerhard Weiss,³ Karl Tuyls⁴ and Ho-fung Leung⁵

Abstract. The complexity of automated negotiation in a multi-issue, incomplete-information and continuous-time environment poses severe challenges, and in recent years many strategies have been proposed in response to this challenge. For the traditional evolution, strategies are studied in games assuming that “globally” negotiates with all other participants. This evaluation, however, is not suited for negotiation settings that are primarily characterized by “local” interactions among the participating agents, that is, settings in which each of possibly many participating agents negotiates only with its local neighbors rather than all other agents. A new class of negotiation games is therefore introduced that take negotiation locality (hence spatial information about the agents) into consideration. It is shown how spatial evolutionary game theory can be used to interpret bilateral negotiation results among state-of-the-art strategies.

1 Introduction

As one of the most fundamental and powerful mechanisms for managing inter-agent dependencies, automated negotiation is central for resolving distributed conflicts between two or multiple parties. Recent years have witnessed an increasing interest in developing negotiation models and strategies for a variety of problems, for example, its deployment in business process management, electronic commerce and markets, task and service allocation, etc. As a result, automated negotiation brings together research topics of artificial intelligence, machine learning, game theory and economics.

Many novel strategies for complex negotiations have been proposed, but they are primarily evaluated in terms of their performance in fixed tournaments, where agents remain their strategies unchanged through tournaments, and which opponents an agent needs to interact with and when they encounter are both fixed. Even although some recent work [1, 2] employ empirical game theory (EGT) to investigate the fitness of the strategies in more open settings where agents are allowed to deviate to different strategies, it still suffers from the small number of possible involved players, and more importantly the limitation of not considering the location of individuals. Against this background, this paper investigates strategy performance in a more interesting but complex environment in which the number of players may be very large and the interaction range of each involved agent is locally limited. Specifically, we consider negotiation settings in which the location of players may affect other agents’ choices of

new strategies. This allows to better understand the the impact of different settings on negotiation strategies’ fitness in a very dynamic environment.

2 Evolutionary game-theoretic analysis of repeated negotiation games

The EGT analysis used in previous work is based on the assumption that each player interacts with all other involved players, that is, global interaction is assumed (e.g., [1, 3]), this however does not hold in many real-life cases. Locality thus is an important factor in negotiation that has not been well studied so far. Moreover, the number of possible players is rather limited; otherwise the resulting profiles/nodes would be extremely large to be well analyzed. There naturally arises another question how the fitness of the strategies changes when the player size dramatically grows? For these reasons, we investigate how a population of players behave by changing their negotiation strategies in the case of local and global interaction ranges. In contrast to global interaction where a player negotiates with all other players, local interaction takes into account the agents’ local neighborhood. Toward this end, evolutionary game theory, more precisely spatial evolutionary game theory [5, 6], is applied to the tournament results. It helps to analyze the impact on fitness (i.e., how well an individual is adapted to a dynamic environment) of each species (strategy) competing with others locally.

In the context of this research, an individual is located at a certain environmental position (cell) and its fitness is determined by the average payoff of its strategy playing against its neighbors⁶. We assume that there is a population of players using the strategy set consisting of the nine top negotiation strategies, with a payoff matrix suggesting utilities of any pair of strategies. Initially, every strategy has a equal population of 100 players randomly distributed over a 30×30 two-dimensional hexagon lattice Λ . Each cell is occupied by a strategy and bordered with other six cells, that is to say, every single cell has six neighbors in its local scale. Calculating the fitness of each cell in the field is simultaneously performed. After this, each cell imitates which one has the highest fitness of its neighborhood (including itself). In this way the natural selection process (i.e., how to choose the new strategy of the cell for the next generation) is well defined.

For clarity, Figure 1 gives a running example about how these strategies evolve, specifically, the dynamics of strategy change over generations. In the beginning (Figure 1(a)), every cell is randomly assigned with a strategy with each marked by a distinct color, while the shares of the strategies in the population are kept being equivalent. Later, in the fifth generation (Figure 1(b)), one strategy (red)

¹ Dept. of Knowledge Engineering (DKE), Maastricht University, NL, email: siqi.chen@maastrichtuniversity.nl

² Massachusetts Institute of Technology, USA, email: jianye@mit.edu

³ Maastricht University, NL, email: gerhard.weiss@maastrichtuniversity.nl

⁴ University of Liverpool, UK, email: k.tuyls@liverpool.ac.uk

⁵ The Chinese University of Hong Kong, HK, email: lhf@cuhk.edu.hk

⁶ The average payoff is calculated through a wide range of ANAC tournaments [4].

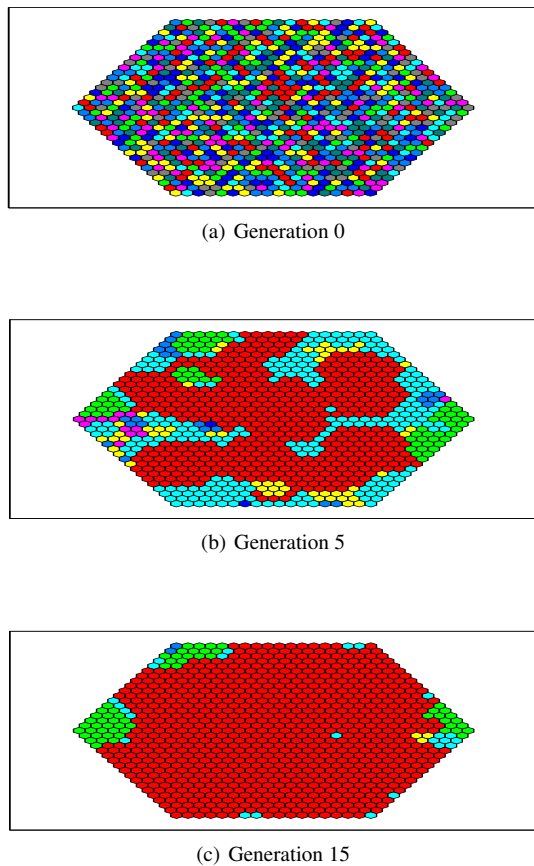


Figure 1. Evolution of strategy distribution.




Figure 2. Strategy distributions over generations when players interact with its direct neighbors.

becomes one of the dominant strategies, whereas a number of strategies are fading away and some even has been exterminated. As evolution continues, there are more strategies vanishing, and the proportion of the red strategy in population gradually grows (as shown in Figure 1(c)).

To obtain results with high statistical significance, we ran the simulation 10,000 times with random initialization of the location arrangement of the nine strategies. Fig. 2 shows the strategy distributions over generations in the case of interaction between nearest neighbors. As can be seen, this spatial evolutionary game, after around 25 generations, ends up with a co-existence of three strategies – Meta-Agent, AgentLG and Gahboninho. Further, the strategy




Figure 3. Strategy distributions over generations when players interact with its neighbors and neighbors' neighbors.

Meta-Agent plays a dominant role in population shares, attracting more than 96% of the individuals. In spite of being the best one in competitions, CUHKAgent is exterminated like other weak strategies. With a poor performance in competitions, the survival of Gahboninho as the second largest proportion (yet quite small) in the population is surprising.

However, if the nature selection process is modified such that a player's interaction range is extended to its neighbors' neighbors, then the difference between Meta-Agent and others would be enlarged. In this case as shown in Fig.3, only Meta-Agent and Gahboninho exist, with Meta-Agent almost fully dominating the population. Moreover, the generations needed for players to converge to Meta-Agent also become shorter. In fact, when further extending agent's interaction range to all other players (i.e., global interaction), all individuals switch to Meta-Agent just in few generations. To summarize, the more players and larger agent-interaction range in the game, the better performance Meta-Agent delivers.

3 Conclusions

This paper presented a new method to evaluate the performance of state-of-the-art agents in complex automated negotiations from the perspective of game theory. More specifically, our work, as the very first work, studies the fitness of negotiation strategies in repeated negotiation games where the number of participating players is large and the location of players serves as an important factor of how to decide their new strategies.

REFERENCES

- [1] Siqi Chen, Haitham Bou Ammar, Karl Tuyls, and Gerhard Weiss, 'Optimizing complex automated negotiation using sparse pseudo-input Gaussian processes', in *Proceedings of the 12th Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems*, pp. 707–714. IFAAMAS, (2013).
- [2] Siqi Chen, Haitham Bou Ammar, Karl Tuyls, and Gerhard Weiss, 'Using conditional restricted boltzmann machine for highly competitive negotiation tasks', in *Proceedings of the 23th Int. Joint Conf. on Artificial Intelligence*, pp. 69–75. AAAI Press, (2013).
- [3] Siqi Chen and Gerhard Weiss, 'An efficient automated negotiation strategy for complex environments', *Engineering Applications of Artificial Intelligence*, **26**(10), 2613 – 2623, (2013).
- [4] Katsuhide Fujita, Takayuki Ito, Tim Baarslag, Koen V. Hindriks, Catholijn M. Jonker, Sarit Kraus, and Raz Lin, *The Second Automated Negotiating Agents Competition (ANAC2011)*, volume 435 of *Studies in Computational Intelligence*, 183–197, Springer, 2013.
- [5] Timothy Killingback and Michael Doebeli, 'Spatial evolutionary game theory: Hawks and doves revisited', *Proceedings of the Royal Society of London. Series B: Biological Sciences*, **263**(1374), 1135–1144, (1996).
- [6] György Szabó and Gábor Fáth, 'Evolutionary games on graphs', *Physics Reports*, **446**(4), 97–216, (2007).