

# Evaluating practical automated negotiation based on spatial evolutionary game theory

Siqi Chen<sup>1</sup>, Jianye Hao<sup>2</sup>, Gerhard Weiss<sup>1</sup>, Karl Tuyls<sup>3</sup>, and Ho-fung Leung<sup>4</sup>

<sup>1</sup> Dept. of Knowledge Engineering, Maastricht University, NL,  
siqi.chen@maastrichtuniversity.nl

<sup>2</sup> Massachusetts Institute of Technology, USA

<sup>3</sup> University of Liverpool, UK

<sup>4</sup> The Chinese University of Hong Kong, HK

**Abstract.** Over the past decade automated negotiation has developed into a subject of central interest in distributed artificial intelligence. For a great part this is because of its broad application potential in different areas such as economics, e-commerce, the political and social sciences. The complexity of practical automated negotiation – a multi-issue, incomplete-information and continuous-time environment – poses severe challenges, and in recent years many negotiation strategies have been proposed in response to this challenge. Traditionally, the performance of such strategies is evaluated in game-theoretic settings in which each agent “globally” interacts (negotiates) with all other participating agents. This traditional 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. This paper presents an approach to handle this type of local setting. Starting out from the traditional global perspective, the negotiations are also analyzed in a new fashion that negotiation locality (hence spatial information about the agents) is taken into consideration. It is shown how both empirical and spatial evolutionary game theory can be used to interpret bilateral negotiation results among state of the art negotiating agents in these different scenarios.

## 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 [11]. 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, economics, and social psychology.

Although automated negotiation has been a very active topic for decades, most of the research efforts in this area focus either on theoretical negotiation models or on simplified models for practical negotiation applications. Owing to the growing popularity

of the international agent-based negotiation competition ANAC [10], more recent research has concentrated on practical bilateral negotiation [3, 8]. They together advance the state of the art of negotiation theory to a more realistic and complex stage. This kind of negotiation normally shares the following five features that are still poorly understood. (1) Negotiations occur in continuous time. (2) The behavior model<sup>5</sup> of the opposing party is not available and can only be observed indirectly through the exchange of offers. (3) There are multiple items under negotiation. (4) The achievable profit through an agreement decreases over time. Finally, (5) participants have a private reservation value, which is set as the agent's minimal benefit when no mutually acceptable agreement can be found. Since the negotiation strategies discussed here are all implemented in the form of (software) agents, in this paper no explicit distinction is made between the terms negotiation strategy and negotiating agent.

There exist a number of good examples of research on complex negotiation such as [4, 5, 8, 16]. Williams et al. [16] employ Gaussian processes for optimizing an agent's own concession rate by predicting the maximal concession that the opponent is expected to make in the future. This strategy, known as IAMhaggler2011, made the third place in ANAC 2011. Another successful strategy based on Gaussian processes is described in [3], where Sparse Pseudo-input Gaussian processes are applied to alleviate the computational complexity of building an opponent model. Hao and Leung [8] propose a novel strategy, which was the winner of ANAC 2012. This method attempts at exploiting the opponent as much as possible by learning opponent behavior and also predicts the optimal offer for the opposing side to improve the acceptance probability of its own proposals, using a reinforcement-learning based approach. Chen et al. [5] adopt an approach called OMAC to complex negotiations that aims at learning an opponent's strategy by analyzing its behavior through discrete wavelet transformation and cubic smoothing spline. With the learnt opponent model, OMAC dynamically adjusts its concession rate in response to uncertainties in the environment. OMAC outperformed the five best agents of ANAC 2011 and was finally awarded the third place in ANAC 2012. To tackle the problem of limited experience available in a single negotiation, Chen et al. [4] then develop a strategy that is able to transfer knowledge efficiently from previous tasks on the basis of factored conditional restricted Boltzmann machines. In the latest edition of the negotiation competition, ANAC 2013 [1], agents can make use of their negotiation history to improve their performance in new encounters. According to the final results, the best-performing agent is Fawkes, which learns an opponent model by combining the two approaches proposed in [5, 8].

Many new and novel strategies for complex negotiations have been proposed, but they are primarily evaluated in terms of their scores in fixed negotiation tournaments [7, 14], where agents leave 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 works [4, 6] employ empirical game theory to investigate the fitness of the strategies (or so-called robustness in other research) 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 con-

---

<sup>5</sup> Because both an agent's utility function and bidding strategy is hidden, we will often use the term behavior model to refer to both as the "joint forces" that govern its negotiating behavior.

sidering the location of individuals. Against this background, the contributions of this paper are as follows. We first provide a standard evaluation of state-of-the-art negotiating agents, which is still missing from current literature. These agents are tested in a number of tournament competitions, where the domains are adopted from the most recent international agent-based negotiation competition (i.e., ANAC 2013). The results are also used as a basis of further analysis. Second, dependent on a strategy-pair payoff matrix (which comes from results of previous tournament competitions), the fitness of the strategies is studied using empirical game theory in a setting where only a few agents globally negotiate with all others. Lastly, we extend this setting to a more interesting but complicated one in which the number of players can 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. Spatial evolutionary game theory is applied to analyze the changes of each strategy share in the whole population. This allows to better understand the impact of different settings on negotiation strategies’ fitness.

The remainder of this paper is organized as follows. Section 2 briefly introduces the negotiating agents analyzed in this work as well as the test domains and the negotiation simulation environment used for the analysis. Section 3 shows results of the performance of a number of state of the art negotiating agents in tournament experiments. Section 4 provides a thorough game-theoretic analysis of the fitness of negotiation strategies in different cases. Finally, Section 5 concludes the paper and identifies some important research lines opened by the described work.

## 2 Agents and test domains

To provide an extensive coverage of advanced negotiation agents, this work considers those agents which are ranked the first three places in ANAC 2011 – 2013 and whose sources can also be publicly accessed. They together result in a highly competitive negotiation setting. An overview of these ANAC agents is given in Table 1. Due to space limitation, here we do not discuss the technicalities of these agents. The interested reader is suggested to refer to [1, 5, 10] for a thorough discussion.

**Table 1.** Overview of top three agents of ANAC 2011 – 2013.

<i>Agent</i>	<i>Affiliation</i>	<i>Achievement</i>
Fawkes	Delft University of Technology	1st in 2013
Meta Agent	Ben Gurion University of the Negev	2nd in 2013
TMF Agent	Ben Gurion University of the Negev	3rd in 2013
CUHKAgent	Chinese University of Hong Kong	1st in 2012
AgentLG	Bar-Ilan University	2nd in 2012
OMAC	Maastricht University	3rd in 2012
HardHeaded	Delft University of Technology	1st in 2011
Gahboninho	Bar Ilan University	2nd in 2011
IAMhaggler2011	University of Southampton	3rd in 2011

The test domain is another decision factor for the quality of evaluation results. On the contrary to previous work that only examines the efficiency of negotiating agents in a relatively small number of domains, for the sake of a high level of generality, we adopt the whole set of domains created for ANAC 2013 (18 domains in total). Moreover, to capture the influence of the discounting factor  $\delta$  and the reservation value  $\vartheta$  on the performance of agents, different values for these two parameters are considered. Thus, experiments are conducted with three discounting factors (i.e.,  $\delta = \{0.5, 0.75, 1.0\}$ ) and three reservation values (i.e.,  $\vartheta = \{0, 0.25, 0.5\}$ ), which produce nine ( $3 \times 3$ ) different scenarios for each domain. In doing so, possible bias on domain selection can be avoided. The agents cannot get chances to optimize their strategies in such a circumstance; on the other hand, a good spread of domain characteristics is also ensured.

The performance evaluation on negotiating agents is done with the simulation environment – Genius [9]. It is the official testbed of the ANAC competition, which allows to evaluate intelligent agents employing different negotiation strategies across a variety of application domains under real-time constraints. Incorporating many key features to support and analyze automated negotiation, Genius facilitates the research on the field and provides a standard platform for people from the community to easily compare newly developed agents with those existing ones. For each scenario of every single domain, we run a tournament ten times to guarantee results with statistical confidence. In each tournament agents repeat negotiation against the same opponent with different negotiation roles (i.e., buyer and seller role) as well as the order in which they start with bidding. If a negotiation fails (e.g., no agreement is made before/at the end of an encounter), then the disagreement solution applies, which means that each agent merely receives its own reservation value  $\vartheta$ .

### 3 Tournament competition results

Because of space limitation, the detailed negotiation settings are not presented here; the interested reader is suggested to refer to [3, 6] for more information. As specified in the previous section, nine variants of each of these 18 domains with different discounting factors and reservation values were used, totalling up to 162 scenarios. This resulted in a total number of 524,880 negotiations in the experiments (with each scenario repeated 10 times).

The overall performance of the agents is summarized in Table 2 with the mean utility and standard deviation. The best performance came from CUHKAgent and AgentLG (1st and 2nd in ANAC 2012), followed by TMF-Agent (3rd in 2013) with a very small difference. Meta-Agent (2nd in 2013), OMAC (3rd in 2012) and Hardheaded (1st in 2011) then took the fourth and sixth place in the competition, respectively. As the latest edition of the competition focuses more on learning and adaption in negotiation, it is surprising to see that ANAC 2013 agents that are given negotiation history to aid their performance in new tasks still (on average) lag behind those 2012 agents. These findings revealed an important fact that the scope of the tournament pool has a significant impact on the experimental results, and illustrated the necessity of a wide range of state-of-the-art benchmarking agents when assessing a negotiation strategy.

**Table 2.** Overall performance of all agents across all scenarios in descending order. The letter in bold of each strategy is taken as its identifier for the later game-theoretic analysis.

<i>Agent</i>	<i>Mean utility</i>	<i>Standard deviation</i>
<b>C</b> UHKAgent	0.656	0.0002
<b>A</b> gentLG	0.656	0.0003
<b>T</b> MF-Agent	0.648	0.0003
<b>M</b> eta-Agent	0.645	0.0004
<b>O</b> MAC	0.638	0.0002
<b>H</b> ardHeaded	0.635	0.0002
<b>G</b> ahboninho	0.626	0.0002
<b>F</b> awkes	0.619	0.0003
<b>I</b> AMhaggler2011	0.581	0.0001

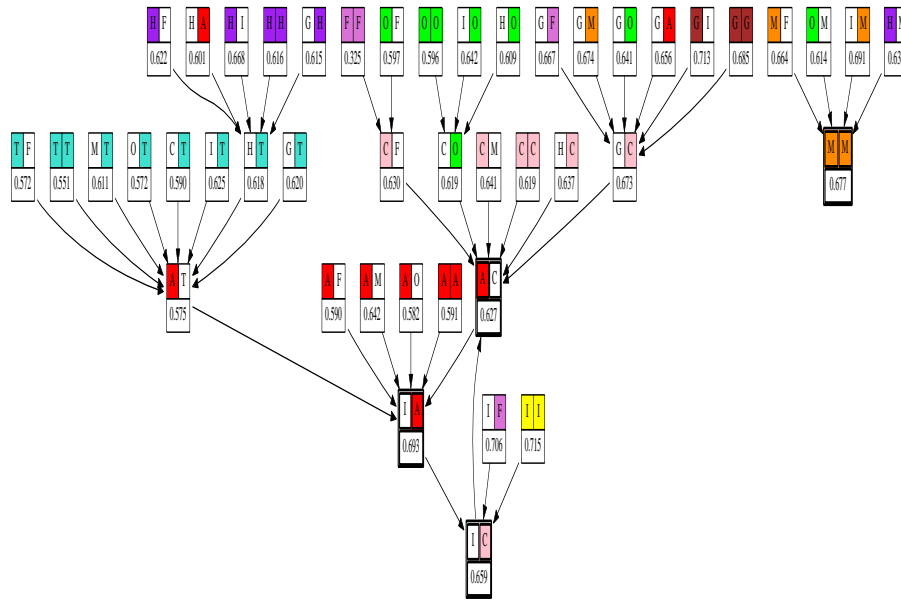
According to the experimental results, the strategy-pair payoff matrix is set up and shown in Table 3 (which are averaged over all negotiation encounters considered in our work). The first letter of each strategy is used as its identifier. As the matrix is symmetric, we only present the row strategy's payoff. On the basis of this payoff matrix, in the subsequent sections we perform the game-theoretic analysis of repeated negotiation scenarios that is more open than those in this section.

**Table 3.** Strategy-pair payoff matrix, where the score pair in each entry is averaged over all domains, with the score representing the column player's payoff (as the matrix is symmetric). The first letter (**bold**) of each agent is used as the identifier.

Payoff	<b>G</b>	<b>H</b>	<b>I</b>	<b>A</b>	<b>C</b>	<b>O</b>	<b>M</b>	<b>T</b>	<b>F</b>
<b>G</b>	0.686	0.672	0.639	0.707	0.731	0.698	0.694	0.706	0.674
<b>H</b>	0.559	0.616	0.592	0.608	0.647	0.640	0.630	0.658	0.609
<b>I</b>	0.788	0.743	0.715	0.766	0.796	0.750	0.773	0.772	0.758
<b>A</b>	0.607	0.593	0.621	0.591	0.582	0.581	0.609	0.536	0.562
<b>C</b>	0.616	0.626	0.522	0.672	0.619	0.628	0.621	0.630	0.592
<b>O</b>	0.585	0.579	0.534	0.584	0.610	0.596	0.592	0.599	0.573
<b>M</b>	0.654	0.648	0.609	0.676	0.661	0.636	0.677	0.653	0.662
<b>T</b>	0.534	0.577	0.478	0.614	0.550	0.545	0.570	0.551	0.515
<b>F</b>	0.660	0.635	0.653	0.617	0.668	0.621	0.666	0.629	0.325

## 4 Game-theoretic analysis of automated negotiations

We have deeply studied the strategy performance from the competition perspective. As discussed before, there however exists a significant limitation of this performance measure because it cannot give any indication about the fitness of these strategies in a open environment. It is unclear, for instance, which strategy will become the winner of a competition when the players are allowed to deviate, i.e., to switch to another strategy for the sake of better individual profits, or when the mixture of opponent strategies



**Fig. 1.** Deviation analysis for the two-player negotiation. Each node shows a strategy profile and the average score of the two involved strategies with the higher scoring one marked by a color background. The arrow indicates the statistically significant deviation between strategy profiles. The equilibria are the nodes marked with a thicker border.

changes. For the purpose of providing a broader view of the agents' performance, empirical game theory and spatial evolutionary game theory are both applied in this section to analyze the fitness of the strategies in two distinct cases. The following analysis is performed based on the payoff matrix shown in Table 3.

#### 4.1 Small number of players with global interaction

We start by studying strategy fitness in the case where several players negotiate with others. To appropriately address strategy fitness with global interaction, empirical game theory (EGT) analysis [12] is employed, which was initially developed to analyze the Trading Agent Competition (TAC). We consider the strategy deviations as discussed in [16], where there is an incentive for one agent to unilaterally change the strategy in order to statistically improve its own profit. The aim of using EGT is to search for pure Nash equilibria in which no agent has an incentive to deviate from its current strategy, or best reply cycle where there exists a set of profiles (i.e., the combination of strategies chosen by players) for which a path of deviations exists that connect them, with no deviation leading to a profile outside of the set. For convenience, these two types of states are both called empirical stable states.

We investigated strategy fitness with global interaction by means of EGT in the two different scenarios with increasing complexity below:

*Scenario 1*: a negotiation encounter between two players.

*Scenario 2*: the full tournament composed of nine players with nine strategies.

This is because the former represents the underlying bilateral negotiation, i.e., only two players participate in the game, and the other illustrates such kinds of negotiations in more complex tournaments. For brevity, we use the bold letters in Table 2 as the identifier for each strategy (e.g., **H** means Hardheaded, **C** means CUHKAgent). The set of strategies is given by  $\Sigma = \{\mathbf{G}, \mathbf{H}, \mathbf{I}, \mathbf{A}, \mathbf{C}, \mathbf{O}, \mathbf{M}, \mathbf{T}, \mathbf{F}\}$ .


In the first scenario, the resulting graph under EGT analysis contains  $\binom{|p|+|s|-1}{|p|} = \binom{10}{2} = 45$  distinct nodes, where  $|p|$  means the number of players and  $|s|$  the number of strategies. A *profile* is defined as the two strategies used by the players in the game (it is worth noting that the two players may use the same strategy). Furthermore, the score of a specific strategy in a particular profile is decided by the payoff matrix given in Table 3. The results are depicted in Figure 1. Each node represents a strategy profile being a mix of two strategies; an arrow indicates the statistically significant deviation to a different strategy profile. Please note that this figure only cares about the strategy mixture and therefore the player order is not taken into account. Under this EGT analysis, there exists one pure Nash equilibrium and a best reply cycle, highlighted by a thick border in Figure 1 as follows:

1. The players both use Meta-Agent, i.e.,  $[\mathbf{M}|\mathbf{M}]$ .
2. a best reply cycle consists of  $[\mathbf{A}|\mathbf{C}]$ ,  $[\mathbf{I}|\mathbf{C}]$  and  $[\mathbf{I}|\mathbf{A}]$ .

The sole equilibrium is the strategy profile  $[\mathbf{M}|\mathbf{M}]$ . This stable state only attracts few profiles. For the remaining states, there exists a path of statistically significant deviations that leads to one state of the best reply cycle. This cycle has a basin of attraction<sup>6</sup> of 89% of the profiles. The results of repeated single negotiations (i.e., between two players) show that there are four empirical stable states including four robust strategies – Meta-Agent, CUHKAgent, IAMhaggler2011 and AgentLG; moreover, Meta-Agent, CUHKAgent and AgentLG are the winning strategies in one or more states (i.e., they gain a higher score). In addition, the analysis also indicates that high-scoring strategies (e.g., TMF-Agent and OMAC) do not necessarily perform well in repeated single encounters, or in other words, they are not robust. It is, however, worth pointing out that the repeated single encounter analysis, while useful, cannot tell anything about the strategy robustness when the setup gets more complicated. Next, we turn our attention to a more complex setting composed of more players with more strategies.

For scenario 2, we consider tournaments consisting of nine players, where each can select one of the nine strategies introduced in Table 2. The results are given in Figure 2. Here a *profile* is defined as the mixture of strategies used by players in a tournament. The nodes in this figure consist of two rows. The top row explains the set of strategies selected by agents in the tournament; the second means the number of agents choosing each strategy. As Figure 2 tells, there is one pure Nash equilibrium that only includes Meta-Agent; in other words, all the players switch to this strategy in

<sup>6</sup> The basin of attraction [2] of a stable state is the number of profile states which converge to a stable state. The likelihood of reaching that stable state can be measured on the basis of the size of the basin of attraction.



**Fig. 2.** Deviation analysis for nine-player tournaments composed of nine strategies. The equilibrium is the nodes with no outgoing arrow and a thicker border.


the end. The basin of attraction of the equilibrium state includes 100% of the profiles. With that, the results of the second case suggest that in the setting of more complex tournaments, only Meta-Agent remains robust, among the four stable strategies that we have found in the repeated single negotiations.

The EGT analysis proves good fitness of four strategies – Meta-Agent, CUHK-Agent, AgentLG and IAMhaggler2011; especially Meta-Agent performs consistently well in both scenarios. It is very interesting to see that although not being the strongest agent in the competition results (refer to Table 2), Meta-Agent is more suitable for an open and competitive environment than others. A high performance in self-play and a fairly good relative advantage of this strategy over other competitors may account for its success.

#### 4.2 Large number of players with various interaction ranges

The EGT analysis is based on the assumption that each player interacts with all other involved players, that is, global interaction is assumed (e.g., [3, 4, 16]). This does not hold in many real-life cases; for instance, in diplomatic negotiations on a territorial dispute it is obvious that negotiation concerns only adjacent countries rather than all countries with which the disputing countries are in some relationship. As another example, the location of individuals is also of great importance for resources allocation in wireless





**Fig. 3.** A locality in a hexagon model.

sensor networks. 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 using the EGT approach; 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 two reasons, we investigate how a population of players or individuals 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 [13, 15], is applied to the tournament results. This allows 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 assumed to be located at a certain environmental position (also called cell) and its fitness is determined by the average payoff of its strategy playing against its neighbors (refer to Table 3). Take a simple case with three strategies as a toy example, where the center cell choosing strategy 1 meets its neighbors as shown in Fig. 3. The fitness of the center cell is the average payoff of playing against three opponents using strategy 1, one opponent using strategy 2, and two opponents using strategy 3 (i.e., with the neighbor distribution  $x=(\frac{3}{6}, \frac{1}{6}, \frac{2}{6})$ ), which is formally defined in Equation 1. The payoff matrix of the three strategies is given by matrix  $A$  below:


$$A = \begin{bmatrix} 4 & 10 & 0 \\ 1 & 4 & 9 \\ 3 & 7 & 4 \end{bmatrix}$$

where an entry  $A(i,j)$  is the payoff of strategy  $i$  against strategy  $j$ . Thus, the fitness ( $\rho$ ) of the center cell is  $\frac{22}{6}$ , following the equation below.

$$\rho = e_i A x^T \quad (1)$$

where  $e_i$  denotes the  $i$ -th row of a unit matrix  $e$  with the size of the number of strategies and  $A$  denotes the payoff matrix.

In our analysis we assume that there is a population of players using the strategy set ( $\Sigma$ ) consisting of the nine negotiation strategies, with a payoff matrix (see Table 3) suggesting utilities of any pair of strategies. Initially, every strategy has an equal population of 100 players randomly distributed over a  $30 \times 30$  two-dimensional hexagon lattice  $\Lambda$ . Each cell ( $I$ ) is occupied by a strategy and bordered with six other cells, that




**Fig. 4.** Strategy distributions over generations when players interact with their direct neighbors.

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 then 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.

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. 4 shows the strategy distributions over generations in the case of players interacting with others in their neighborhood. 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 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 (see Table 2), 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 natural 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. We show the result in Fig.5. In this case, only two strategies – Meta-Agent and Gahboninho exist, while Meta-Agent almost fully dominates the population share. Moreover, the time needed for players to converge to Meta-Agent also becomes shorter. As a matter of fact, when further extending agents’ interaction range to all other players (i.e., global interaction), all individuals switch to Meta-Agent in just few generations.

When comparing the results depending on empirical game theory and spatial evolutionary game theory, we found that Meta-Agent was a very successful strategy in various scenarios. To summarize, the more players and larger agent-interaction range in the game, the better performance it delivers.



**Fig. 5.** Strategy distributions over generations when players interact with their neighbors and neighbors' neighbors.

## 5 Conclusions

This paper presented a thorough review of the performance of practical negotiation strategies from the perspective of game theory. A wide range of high quality agents (using these strategies) from ANAC competitions were evaluated with respect to two different game-theoretic techniques. More importantly, this paper, as the very first work, studies the fitness of negotiation strategies in a repeated fashion 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. The detailed analysis conducted in the work provided a number of valuable new insights into the efficacy of a negotiation strategy. First of all, competition performance, while important, does not serve as a good indicator for an agent in an open environment in which agents, having freedom to change their strategies, repeatedly negotiate with others. Recall that some high-scoring agents like CUHKAgent and OMAC perform disappointingly in the new settings. Then, for the negotiations where a small number of players are involved with global interaction, fitness of strategies may be qualitatively different in settings with different complexity. For instance, AgentLG is not robust anymore in the second scenario where more players and strategies are available. Last but not least, the results obtained from spatial evolutionary game-theoretic analysis illustrate that agent-interaction range indeed has an effect on the evolution process of strategies in terms of when to reach a stable state and what share a strategy can occupy in the population. Moreover, the range of interaction seemingly boosts the performance of leading strategies. They together confirm the necessity of considerations of both local and global interaction in the performance analysis of the strategies.

Regarding future work, we believe it is worth investing research efforts in exploring several interesting questions. For example, how would the evolution process change if the impact of neighbors is weighted by their distance (e.g., weighted influence of neighbors), or if a player is allowed to choose any strategy from  $\Sigma$  rather than merely from its neighborhood (e.g., by modifying the natural selection process). Another important

avenue we see is to apply our method to the negotiation results between autonomous negotiating agents against human negotiators.

## References

1. The Fourth International Automated Negotiating Agent Competition (ANAC 2013). <http://www.itolab.nitech.ac.jp/ANAC2013/>.
2. T. Baarslag, K. Fujita, E. H. Gerding, K. Hindriks, T. Ito, N. R. Jennings, C. Jonker, S. Kraus, R. Lin, V. Robu, and C. R. Williams. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artificial Intelligence*, 198:73 – 103, 2013.
3. S. Chen, H. B. Ammar, K. Tuyls, and G. 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*, pages 707–714. ACM, 2013.
4. S. Chen, H. B. Ammar, K. Tuyls, and G. Weiss. Using conditional restricted boltzmann machine for highly competitive negotiation tasks. In *Proceedings of the 23th Int. Joint Conf. on Artificial Intelligence*, pages 69–75. AAAI Press, 2013.
5. S. Chen and G. Weiss. An efficient and adaptive approach to negotiation in complex environments. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 228–233. IOS Press, 2012.
6. S. Chen and G. Weiss. An efficient automated negotiation strategy for complex environments. *Engineering Applications of Artificial Intelligence*, 26(10):2613 – 2623, 2013.
7. K. Fujita, T. Ito, T. Baarslag, K. V. Hindriks, C. M. Jonker, S. Kraus, and R. Lin. *The Second Automated Negotiating Agents Competition (ANAC2011)*, volume 435 of *Studies in Computational Intelligence*, pages 183–197. Springer Berlin / Heidelberg, 2013.
8. J. Hao and H. Leung. ABiNeS: An adaptive bilateral negotiating strategy over multiple items. In *Proceedings of WI/IAT’2012*, pages 95–102. IEEE Computer Society, 2012.
9. K. Hindriks, C. Jonker, S. Kraus, R. Lin, and D. Tykhonov. Genius: negotiation environment for heterogeneous agents. In *Proceedings of the 8th international joint conference on autonomous agents and multiagent systems*, pages 1397–1398. ACM, 2009.
10. T. Ito, M. Zhang, V. Robu, S. Fatima, and T. Matsuo, editors. *New Trends in Agent-based Complex Automated Negotiations*. Springer-Verlag, 2012.
11. N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.
12. P. R. Jordan, C. Kiekintveld, and M. P. Wellman. Empirical game-theoretic analysis of the tac supply chain game. In *Proceedings of AAMAS’2007*, pages 1188–1195. ACM, 2007.
13. T. Killingback and M. 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.
14. I. Marsa-Maestre, M. A. Lopez-Carmona, T. Ito, M. Zhang, Q. Bai, and K. Fujita, editors. *Novel Insights in Agent-based Complex Automated Negotiation*. Springer, 2014.
15. G. Szabó and G. Fáth. Evolutionary games on graphs. *Physics Reports*, 446(4):97–216, 2007.
16. C. Williams, V. Robu, E. Gerding, and N. Jennings. Using gaussian processes to optimise concession in complex negotiations against unknown opponents. In *Proceedings of the 22th Int. Joint Conf. on Artificial Intelligence*, pages 432–438. AAAI Press, 2011.