Coordination in Multi-Player Human-Computer Groups

Yaniv Mazliah and Ya'akov (Kobi) Gal

Department of Information Systems Engineering, Ben-Gurion University of the Negev, Be'er-Sheva 85104, Israel

Abstract. There is inconclusive evidence whether practising tasks with computer agents improves people's performance on these tasks. This paper studies this question empirically using extensive experiments involving human subjects which were trained by computer agents to play a three-player coordination task that is a common test bed in AI to evaluate computational cooperation strategies. Following training, we compared the performance of subjects when playing state-of-the-art agents from the literature. The results revealed that training with computer agents increased people's performance when compared to the state-of-the-art agent. These results demonstrate the efficacy of using computer agents as tools for improving people's skills when interacting in strategic settings, saving considerable effort and providing better performance than when training with human counterparts.

1 Introduction

Settings in which people and computers cooperate arise in a wide variety of application domains (e.g, hospital care-delivery systems, systems administration applications) as well as in virtual reality and simulation systems (e.g, disaster relief, military training) [1–4]. The automated computer agents in these settings are designed for the purpose of supporting people, acting as proxies for individuals or organizations, or working autonomously to carry out the actions for which they are responsible.

There is a breadth of prior work on the design and evaluation of computer agents for cooperating with people. However the effects of using computer agents to change people's behaviour in strategic settings is inconclusive. On the one hand, autonomous agents designed by researchers and students commonly use opponent modelling, game theoretic reasoning and machine learning, approaches that allow them to perform successfully in their respective setting [5]. On the other hand, when deciding whether to cooperate, people prefer to cooperate and coordinate with other people rather than with computer agents [6].

To address this gap, we study the question of whether using automated agents to train people can improve people's performance in a representative setting involving coordination among groups of three participants. Our method involves people practising to coordinate in a given task with other participants—whether other people, or computer agents. We compared people's behaviour during training with that of their performance during a separate testing phase conducted on the same task. A challenge to evaluating people's performance in these multi-participant task settings is that their behaviour depends in part on the strategies of the other participants. We therefore used a standardized agent to interact with people when comparing between their performance in the testing phase. This agent was chosen from the state-of-the-art, meaning that its proficiency was already demonstrated when interacting with other computer agents (or people) in separate studies. The use of the standardized agent provided an objective metric with which to evaluate people's performance.

Our empirical methodology consisted of a three-player multi-round coordination game of imperfect information that is used as a research test-bed for evaluating novel agent designs in Artificial Intelligence [7]. We compared people's performance in these settings when interacting with other people or with another computer agent. All situational contexts were kept identical for subjects in both conditions. Thus, any difference in their behavior can be attributed to the history of their prior interaction in the training phase.

Results show that using automated agents to train people improved their performance when compared to the standardized agent's performance. Also, training with people improved the performance of those people that coordinated more often with the standardized agent. Further analysis revealed that the top-scoring people preferred to cooperate with the computer agent rather than with other people, causing the average human score to drop. In both settings, there was no significant improvement in people's performance after training with people.

These results have insight for agent designers for human-computer decision-making as well as social scientists. They suggest that in settings requiring coordination and agreements, people can learn to be more skilful by learning to play from computer agents. These agents can be used as tools for training people in such tasks. This can result in considerable savings in cost and effort as compared to using people for training purposes.

The remainder of this section briefly discusses work in behavioural economics related to human behaviour in coordination games, and work in computer science related to the design of computer agents for coordination with humans. The rest of the paper is organized as follows. Section 2 describes the Lemonade Stand Game and presents tactics for playing the game. Section 3 presents our empirical methodology and results. Finally, we conclude the paper with future directions for research.

1.1 Related Work

Past work in behavioral economics has studied the extent to which people learn to play equilibrium strategies in coordination games [8–10]. However, there is scant work on predicting people's coordination strategies in coordination games. Classical game theoretic techniques are not successful at predicting people's play in coordination games [11]. Recently, Zuckerman et al. [12] used machine learning to predict people's coordination strategies in different types of coordination strategies. None of these works have studied the effect of prior play in coordination games on people's performance. Predicting cooperation is especially hard when there several Nash equilibrium to the game, It is difficult to predict which focal point will the game progress to. When observing people's behavior we also have to take into to account that people not always converge into such equilibrium.

Our work is also novel in studying human coordination in groups of more than two players. Prior work has shown that coordination is challenging for people even in two player game settings. For example, in the Battle of the Sexes game, people mis-coordinate significantly more often than that predicted by the mixed Nash equilibrium of the game, while introducing communication into the game minimized the mis-coordination [13]. Our study demonstrates the efficacy of using computer agents to improve human coordination in more challenging three-player settings.

2 The Lemonade Stand Game

The Lemonade Stand Game (LSG) was originally proposed as a test-bed for the evaluation of opponent modelling and machine learning techniques [7]. In this game there are twelve possible actions for each player, representing possible locations to setup a lemonade stand on the beach of an island. The twelve locations are uniformly spread around the perimeter, in the same way that hours are displayed on the face of a clock. Players choose their locations on the board simultaneously. The players in the game represent lemonade vendors who compete to serve customers in their vicinity. The utility to each player is the sum of its distance from the nearest player in a clockwise direction and the nearest player in a counter-clockwise direction. Distances are measured by counting the number of positions between players. Actions are taken simultaneously by all players. This is analogous to the profit each player would make if customers buying lemonade are uniformly distributed around the island. If more than one player are positioned in the same location, they share the profit they incur in that location. In this case, both of the two players in the "collision" receive a score of 6, whereas the third player receives a score of 12. If all three players are positioned in the same position, then each receive a score of 8. The rules of the game are commonly changed every year to promote active research, such as varying the utility function or the number of games played in succession with the same agent players. In this study we confined ourselves to the original game description used in the 2009-10 competitions.

A snapshot of the GUI for playing the LSG is shown in Figure 1. The snapshot is shown for a particular round from the perspective of the player whose position is represented by the red disc at 6 o'clock. This player scored 10 points in this round. This is because the distance of the player from the "green player" at 2 o'clock is four positions, and the distance of the player from the "blue player" at 8 o'clock was 6 positions. The score shown in the upper right-hand-side of the figure in both table and graph formats. In addition, the cumulative score for all players is displayed on the upper left-hand-side of the figure. The bottom part of the Figure is a history panel in which participants can observe the results of all of the prior rounds of the game.

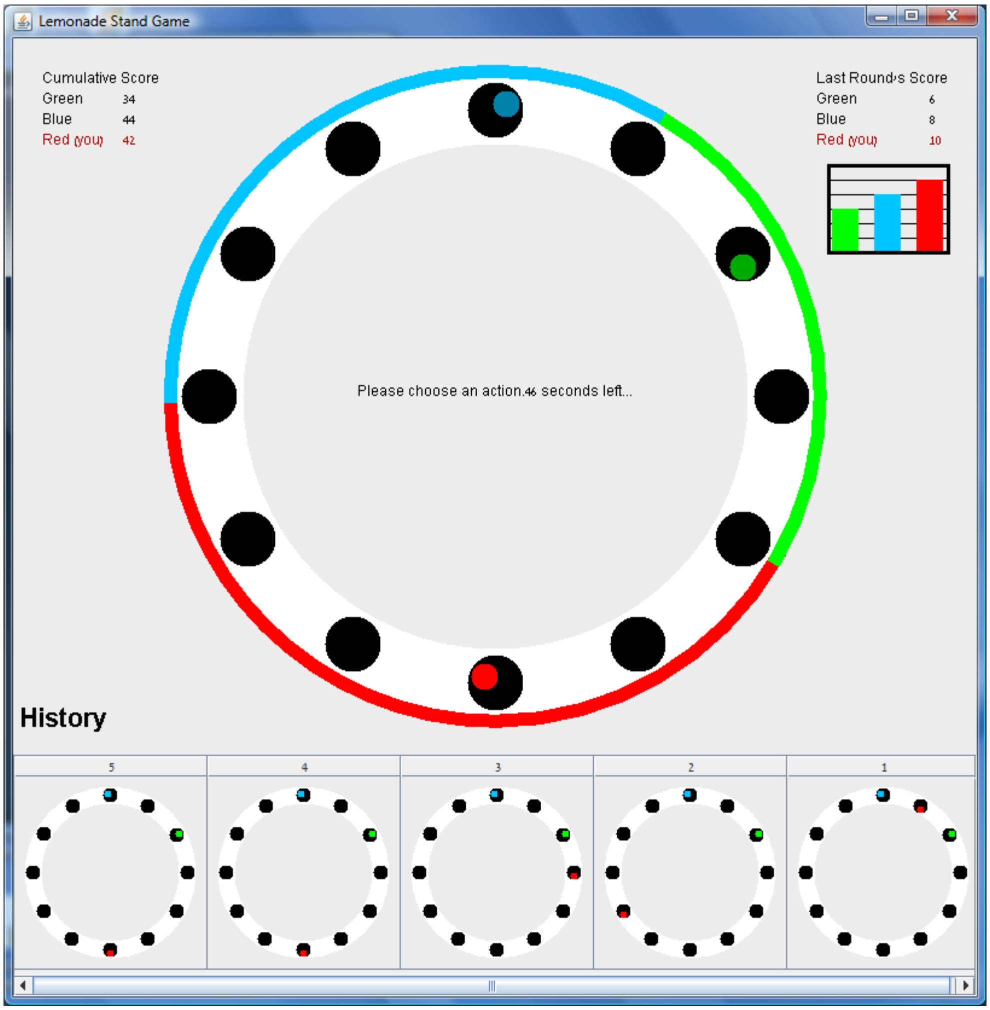

Fig. 1. A snapshot of the Lemonade Stand Game showing the main board panel from the perspective of one of the players, as well as the history of past rounds.

The advantage of using the LSG as a setting for training people's play is twofold. First, its rules are simple and intuitive. However, it is challenging for people to play. The game is inherently competitive, meaning that when a player gains from playing a particular strategy, other players necessarily lose. However, It also allows for cooperation between players to coordinate an "attack" on the third player. To succeed, players not only need to reason about what strategy the other players are using but also to try to influence their strategy over time. In addition, player's outcomes are affected by the strategies of two other players. This makes it more difficult for people to identify and learn beneficial strategies to play during training. Second, there is a publicly available library of agents designed by experts to compete in an annual tournament.

2.1 Strategies in the Game

In this section we describe key strategies towards playing the game [14]. A visualization of these strategies on a Lemonade stand board game is shown in Figure 2. In one of these tactics, called "stick", a player chooses to remain in the same location in two consecutive rounds.

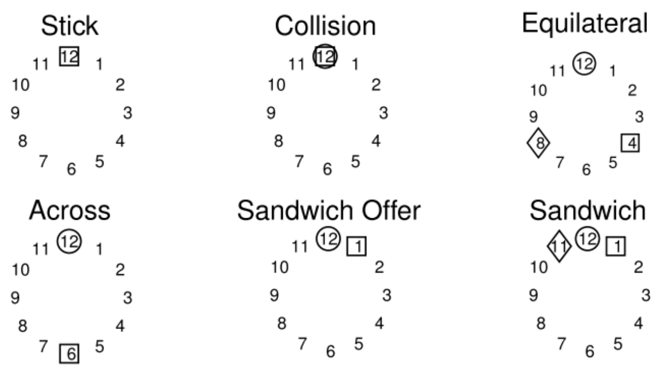

Fig. 2. Key strategies in the Lemonade Stand Game.

The situation in which both players are positioned directly across from each other results is called "across". This strategy guarantees a payoff of between 6 and 12 to each of the players that are situated across from each other, while the third player that is "locked in" receives a payoff of 6. The "Equilateral" positioning in which the players are situated at an equal distance from each other incurs a score of 8 points to each player. Another more aggressive form of cooperation is the "Sandwich" strategy, in which one of the players is positioned directly between the two other players, provides a higher outcome of 11 for both of the agents positioned at the edges and a low score of two for the agent in the middle. To achieve a sandwich outcome to the game a player might try to use a "Sandwich offer".If a player identifies a sticking player he can move to the location next to him, effectively offering the third player the option to perform a "Sandwich" collaboration move. "Collision" is a scenario where two players choose the same position. While this is forcing both player to a low score of 6, it could be used to encourage a player to leave a certain position.

There are several classes of Nash equilibria strategies in the game. One such class includes the situations in which the three players are located in different locations on the board and each player earns at least 6 points is a Nash equilibrium in the game. This is because no player can increase its score by deviating in this situation. For example, the equilateral positioning is a Nash equilibrium. This might seem like a good strategy profile for the game. However the across strategy profile is also a Nash equilibrium, and dominates the equilateral positioning from the point of view of the two players situated across from each other. Another class of Nash equilibrium strategies includes the situations in which two players are situated in the same position, while the third is positioned directly across from them. The "Sandwich" strategy profile is not a Nash equilibrium, because the low scoring player in the middle has an incentive to deviate from its position.

3 Empirical Methodology

We recruited 56 undergraduate students from Ben-Gurion university to play the game. Ages ranged from 24 to 30, 58% males and 42% females. All subjects played 90 rounds of the LSG, divided into three epochs of thirty games each. The first two epochs (called "training epochs") were used to train the people to play the game, and their performance were measured when playing the final "testing epoch". The player configuration in the training epochs varied as follows. In the "all-human" training condition, the player configuration included only human players. In the "single-agent" training condition, the player configuration included two human players and a single agent player. In the "two-agent" training condition, the configuration included a single human player and two agent players. The player configuration for the testing epoch included two human players and the standardized agent player.

Subjects were randomly divided into the single- and two-agent training conditions, as well as a baseline condition in which subjects played a single testing epoch with the standardized agent. Prior to playing the game all subjects were introduced to the game instructions, as well as the basic strategies in the game that were described in the previous section. Altogether, there were 16 games played in the no-training condition, 19 games played in the all-human and single-agent training conditions, and 12 games played in the double-agent training condition. Subjects were paid according to their total score in all three epochs. Payments were score based and not were not related to the score other subjects achieved.

To choose the standardized and training agent we ran an independent tournament which evaluated the agent entries submitted to the 2009 and 2010 agent competitions [15]. The tournament consisted of 30 rounds, the same number of rounds in the testing epoch.[1] The winner of the tournament, called EA^2 (which also won the 2009 competition) was chosen to be the standardized agent for testing performance after training. The training agents were chosen to be the second

[1] The actual 2009 and 2010 tournaments ran 1000 rounds.

and third runner ups in the competition. All of the automated agents we used in the study focus on forming across manoeuvres with one of the other players. This is communicated by playing stick in the same position for several rounds, thus allowing a second player to take the position across, or by moving to the position directly across from another player that is playing stick. However the agents also had distinguishing factors. The standard agent EA^2,which we used in the testing, had a model it built for each player. Using this model the agent was able to decide who is a better partner for cooperation. The agents used for training used less sophisticated approach which was to try and cooperate with one player, if that player did not cooperate than the agent tried to cooperate with the other player.

3.1 Results

All results reported as significant in the following section were confirmed in the $p < 0.05$ range using single-factor ANOVA tests. We hypothesized that subjects that trained with two agents would increase their performance in the testing epoch when compared to (1) subjects that trained with a single agent; and (2) the standardized agent. We first compare people's play in the various conditions to that of the standardized agent. Figure 3 shows the average aggregate performance of people and of the standardized agent in the testing epoch. As shown by the Figure, the EA^2 agent significantly outperformed people in the all-human and all-human training conditions. However, the difference in performance between the EA^2 agent and people was not significant in the single-agent and double-agent training conditions. (This can also be seen in the Figure due to overlapping standard error bars).

We conjectured that the reason for the improvement in people's performance was that people learned to play beneficial strategies from interacting with agents in the training epochs. To examine this, we analysed people's behaviour in the game. Table 1 shows the frequency of follow and stick strategies in the testing epoch for the all-human and double-agent training condition. As shown by the table, the standardized agent engaged in significantly more follow and stick strategies when compared to people in the all-human condition. Thus people were less likely to cooperate than the standardized agent during the test epoch, and played more erratically. However, there was no significant difference in the number of follow and stick actions of the standardized agent and people in the double-agent condition. Consequently, the number of possible across outcomes for the standardized EA^2 agent (also shown in the table) was significantly higher than people in the no-training condition but not in the two-agent training condition. As shown by the table, people engaged in significantly more follow, stick and across strategies in the two-agent training condition (shown in boldface) than in the single-agent training condition. The difference between the number of follow and stick strategies played by people in the single-agent and no-training condition exhibited a similar pattern. Therefore, we attribute the improvement in people's play in the single- and double-agent condition to their increased use of cooperative strategies. Lastly, as shown by Figure 3, people were not able to outperform the agent after training. We attribute this to the inherent difficulty of playing the state-of-the-art agent for this game.

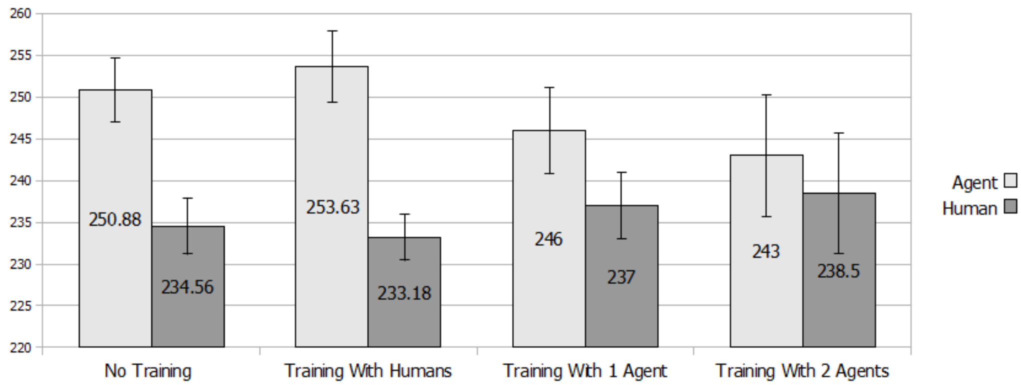

Fig. 3. Performance Comparison: People versus the standardized EA^2 agent

	All-Human			Two-Agent		
	Follow	Stick	Across	Follow	Stick	Across
People	8.82	7.13	9	**17**	**21**	**25.33**
EA^2	**13.4**	**20.8**	11.58	23.08	21.33	21.92

Table 1. Number of "Follow" and "Stick" strategies used by human players in the testing epoch

3.2 Comparing People's Performance Across Conditions

We now compare people's performance scores across the various conditions. The number of follow, stick and across outcomes for people in the two-agent training condition was significantly higher than in the no-training and all-human training condition. Thus people learned to be more cooperative when training with two agents. As shown by Figure 3, people's average score in the two-agent training condition (238 points) was higher than in the no-training (234 points) and single-agent training condition (237), but this difference was not significant.

To explain this discrepancy we distinguish between the top- and low-scoring human players in each game (the people who scored the highest and lowest scores in each game). Because the LSG is constant game, one player's win is another player's loss. In the context of the lemonade stand game, this means that when two players coordinate and play across, they necessarily earn more points than the player that is left out. We hypothesized that top-scoring human players coordinated more often with the standardized agent than low-scoring players, allowing them to outperform the low-scoring players.

Figure 4 shows the average performance of low- and top-scoring human players in all conditions. As shown by the Figure, the top-scoring players in the two-agent training condition outperform top-scoring players in the one-agent and all-human training conditions. Also shown in the figure is that there is no difference in the performance of low-scoring agent across conditions. Table 2 shows the frequency of across outcomes in the two-agent training condition for low- and top-scoring human players. As shown by the table, the top-scoring players played significantly more across outcomes than low-scoring players in each of the conditions. In addition, top-scoring players achieved significantly more across outcomes in the two-agent training condition than in the all-human and single-agent training condition. This confirms our hypothesis, in that the success of top-scorers in the game is attributed to their increased coordination with the standardized agent (rather than the other human player).

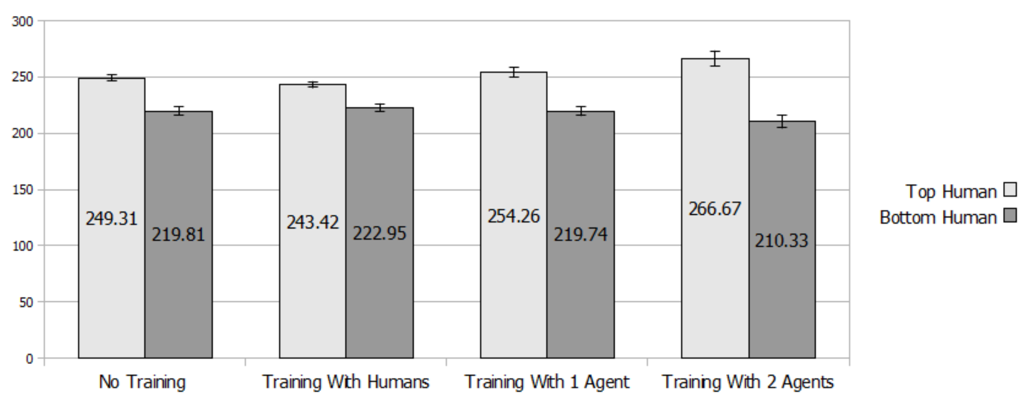

Fig. 4. Performance Comparison: People versus people

	Follow	Stick	Across
low-scoring	16	15	14.92
top-scoring	**26.6**	**20**	**25.3**

Table 2. Number of “follow” and “stick” strategies used by top- and low-scoring human players in the two-agent training condition

4 Conclusions

In this paper we present some interesting results in the Lemonade Stand Domain, we try to answer the question can people learn and improve using simulations with automated agents. Our results show that such training is possible, and is superior to training with inexperienced people, this is true even though the domain requires players to adapt to other players and try to influence their behaviour. It is interesting to note that people were able to learn from agents even when training with one human and one automated agent, thus people are able to learn even though they don't observe two players who has some plan of how to coordinate. In future work we are planning to study the effect of computer agents on people playing other people as well as agents written by human subjects that are not experts in the game.

5 Acknowledgements

Thanks to Martin Zinkevich for developing the GUI interface to the Lemonade Stand Game.

References

1. Schurr, N., Patil, P., Pighin, F., Tambe, M.: Using multiagent teams to improve the training of incident commanders. In: Proc. 5th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS). (2006)
2. Murphy, R.R.: Humanrobot interaction in rescue robotics. IEEE Transactions On Systems, Man, and CyberneticsPart C: Applications and Reviews **34**(2) (May 2004)
3. Kamar, E., Horvitz, E., Meek, C.: Mobile Opportunistic Commerce: Mechanisms, Architecture, and Application. In: Proc. 7th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS). (2008)
4. Rajarshi, D., Hanson, J.E., Kephart, J.O., Tesauro, G.: Agent-human interactions in the continuous double auction. In: Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI). (2001)
5. Lin, R., Kraus, S.: Can automated agents proficiently negotiate with humans? Communications of the ACM **53**(1) (2010)
6. Nass, C., Reeves, B.: The media equation: How people treat computers, televisions, and new media as real people and places (1996)
7. Zinkevich, M., Bowling, M., Wunder, M.: The lemonade stand game competition: solving unsolvable games. ACM SIGecom Exchanges **10**(1) (2011) 35–38
8. Mehta, J., Starmer, C., Sugden, R.: The nature of salience: An experimental investigation of pure coordination games. The American Economic Review **84**(3) (1994) 658–673
9. Crawford, V.: Adaptive dynamics in coordination games. Econometrica: Journal of the Econometric Society (1995) 103–143
10. Cachon, G., Camerer, C.: Loss-avoidance and forward induction in experimental coordination games. The Quarterly Journal of Economics **111**(1) (1996) 165

11. Van Huyck, J., Battalio, R., Beil, R.: Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. The Quarterly Journal of Economics **106**(3) (1991) 885
12. Zuckerman, I., Kraus, S., Rosenschein, J.: Using focal point learning to improve human–machine tacit coordination. Autonomous Agents and Multi-Agent Systems **22**(2) (2011) 289–316
13. Straub, P.: Risk dominance and coordination failures in static games. The Quarterly Review of Economics and Finance **35**(4) (1996) 339–363
14. Wunder, M., Kaisers, M., Littman, M., Yaros, J.: A cognitive hierarchy model applied to the lemonade game. In: AAAI Workshop on Interactive Decision Theory and Game Theory (IDTGT). (2010)
15. Zinkevich, M., Bowling, M., Wunder, M.: The lemonade stand game competition: solving unsolvable games. ACM SIGecom Exchanges **10**(1) (2011) 35–38