Decision Support Systems 60 (2014) 1–9

Contents lists available at ScienceDirect

Decision Support Systems

journal homepage: www.elsevier.com/locate/dss

Training with automated agents improves people's behavior in negotiation and coordination tasks☆

Raz Lin [a,*], Ya'akov (Kobi) Gal [b,d], Sarit Kraus [a,c], Yaniv Mazliah [b]

[a] Department of Computer Science, Bar-Ilan University, Ramat-Gan 52900, Israel
[b] Department of Information Systems Engineering, Ben-Gurion University of the Negev, Be'er-Sheva 85104, Israel
[c] Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, USA
[d] School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA

ARTICLE INFO

Available online 5 June 2013

Keywords:
Training
Automated agents
Automated negotiation
Coordination

ABSTRACT

There is inconclusive evidence whether practicing tasks with computer agents improves people's performance on these tasks. This paper studies this question empirically using extensive experiments involving bilateral negotiation and three-player coordination tasks played by hundreds of human subjects. We used different training methods for subjects, including practice interactions with other human participants, interacting with agents from the literature, and asking participants to design an automated agent to serve as their proxy in the task. Following training, we compared the performance of subjects when playing state-of-the-art agents from the literature. The results revealed that in the negotiation settings, in most cases, training with computer agents increased people's performance as compared to interacting with people. In the three player coordination game, training with computer agents increased people's performance when matched with 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 interacting with human counterparts.

1. Introduction

Settings in which people and computers make decisions together arise in a wide variety of application domains (e.g., hospital care-delivery systems, system administration applications) as well as in virtual reality and simulation systems (e.g., disaster relief, military training). 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. However, there is scant work on the influence of autonomous agents on people's behavior.

The evidence on the use of computer agents to change people's behavior in strategic settings is inconclusive. On the one hand, autonomous agents designed by researchers and students commonly use opponent modeling, game theoretic reasoning and machine learning, approaches that allow them to perform successfully in their respective setting [13]. On the other hand, when deciding whether to cooperate, people prefer to cooperate with other people rather than with computer agents. In particular, people have been shown to offer less to computer agents when making agreements than to people [24].

To address this gap, we study the question of whether using automated agents to train people can improve people's performance in two representative settings involving negotiation and coordination among multiple participants. We propose two methods for training people in these settings that are evaluated empirically in extensive experiments. The first training method involves people practicing a given task with other participants (whether other people, or computer agents that are designed by researchers and students). The second method involves people designing an automated agent to serve as their proxy in the given task. We compared the efficacy of these approaches by measuring people's behavior 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 tasks is that their behavior depends in part on the strategies of the other participants. We therefore used a standardized agent to interact with people when comparing between their performances in the testing phase. This agent was chosen from the state-of-the-art in each of the respective settings, 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 consists of three settings. The first two consisted of different types of strategic multi-attribute bilateral negotiation tasks of imperfect information. The first simulated a job

☆ Preliminary results of this research were published in the proceedings of AAMAS 2009.

* Corresponding author.
E-mail address: linraz@macs.biu.ac.il (R. Lin).

0167-9236/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.dss.2013.05.015

interview between an employer and candidate, while the second simulated diplomatic negotiations (preliminary results on the first setting were published by Lin et al. [16]). In both cases, an agreement consisted of an assignment of possible values for each of the attributes, and the negotiation was conducted using an alternating offer protocol. The third setting was purely competitive and consisted of a three-player multi-round coordination game commonly used in the literature to evaluate computer agents [25]. We compared people's performance in these settings under some or all of the following training conditions:

- classical role playing (training) with another human counterpart;
- training with an automated agent;
- designing and coding an automated agent to act as a proxy.

During the testing and training phase, subjects were not told that they were interacting with an agent. Thus, any difference in their behavior can be attributed to the history of their prior interaction in the training phase.

Results showed that training with state-of-the-art agents helped people improve their performance for all role contingencies in the job candidate and the coordination setting and all role contingencies but one in the diplomatic negotiation setting. Training with agents designed by the subjects themselves improved their performance for all role contingencies in the job candidate and the coordination setting, but had a negligible effect in the diplomatic negotiation setting. Further analysis revealed that in the coordination game, training with people improved the performance of those people that coordinated more often with the standardized agent.

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 skillful 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 the paper is organized as follows. In Section 2 we review related work focused on the evaluation of training methods and the use of simulation and role-playing for training. Sections 3 and 4 present experiments and results for the negotiation and coordination settings in our study. Finally, we conclude the paper with open questions and future directions for research.

2. Related work

We first discuss related work relating to training people to perform negotiation tasks. The use of simulations and role-playing is common for training people in negotiations (e.g., the Interactive Computer-Assisted Negotiation Support system (ICANS) [22], the InterNeg Support Program for Intercultural REsearch (INSPIRE) [9] and virtual humans for training [8]). Surprisingly, little research has been conducted that measures the effect of simulations and role-playing directly on people's negotiation skills, despite underlying assumptions that role-playing improves people's negotiation skills [5,21]. Specifically, several works have evaluated the role of simulation in training students' skills as diplomatic negotiators using questionnaires and subjective reporting [4,20]. Susskind and Corburn [21] study the usefulness of negotiation simulations by questioning leading practitioners in the field about why and how they use simulations to teach negotiation. Kenny et al. [8] and Traum et al. [23] have used virtual humans to facilitate people's negotiation, leadership and interviewing skills. These virtual humans were tested in several negotiation scenarios in social and military contexts in which culture plays a crucial role. Lennon et al. [11] have studied the extent to which training improves people's negotiation skills across cultures, as measured by their performance in a post-training negotiation task. There is no prior work that uses automated agents for the purpose of improving human performance in negotiation.

Another strand of research has studied the role of media, GUIs and decision support tools on people's negotiation behavior. Ross et al. [19] and Butler [2] studied whether watching negotiation simulations on video helped students increase their learning of negotiation concepts, as measured by students' reaction to the video and their ability to recognize pivotal points in the negotiation process. Other works have studied the role of web-based GUIs for facilitating negotiation [9,12]. None of these methods have measured the effect of these support tools on people's performance in real time.

The use of automated agents in human–computer negotiations is a burgeoning field in Artificial Intelligence. For a comprehensive summary, see the survey by Lin and Kraus [13]. Most work in this field has focused on the design of agents that can reach more beneficial agreements than do people [3,6,10,14]. Notable exceptions include Kamar et al. [7] who designed a computer agent that used collaborative decision-making strategies to interact with people in a cooperative game, and Bachrach et al. [1] who showed that agents playing strategies who implement solution concepts from cooperative game theory can play well with people in a weighted voting game. None of these works have studied the effect of prior play in coordination games on people's performance.

Table 1
Number of subjects in each evaluation method in the Job-Candidate and Britain–Zimbabwe domains.

Approach/role	Employer	Job candidate	Britain	Zimbabwe
Control group	18	16	15	15
Training via Human Negotiation	18	18	20	20
Training via Automated Negotiator	20	20	18	18
Training via Agent Design	19	19	15	N/A

Table 2
Comparison of the average scores and standard deviation of human negotiators using different training methods and the control group.

Method	Role	Average	Std.	*p*-Value
Control group	Employer	431.78	80.83	
	Job Can.	320.5	112.71	
	Britain	335.33	194.62	
	Zimbabwe	−320.07	274.42	
Training via human negotiation	Employer	448.56	66.08	0.25
	Job Can.	**383.83**	112.73	0.05
	Britain	366.45	198.65	0.32
	Zimbabwe	−268.7	301.93	0.3
Training via agent design	Employer	466.84	46.26	0.06
	Job Can.	**391.53**	76.75	0.02
	Britain	422.93	162.77	0.09
	Zimbabwe	N/A		
Training via automated negotiator	Employer	**468.6**	38.94	0.04
	Job Can.	**433**	102.84	0.002
	Britain	301.22	182.14	0.3
	Zimbabwe	**−44.6**	196.19	<0.002

3. Training methods in bilateral negotiation

In this section we study whether role-playing with people or training with automated agents can enhance the negotiation experience by improving the negotiation skills of human negotiators.

3.1. The bilateral negotiation settings

Following Lin et al. [14] we consider a bilateral negotiation settings in which two agents, either automated negotiators or people,