Revised version appeared in:

Revised Selected Papers from Web 2011 Proceedings

2012 (356-366)

An Experiment in Human – to - Software Agent Negotiations

Rustam Vahidov[1], Gregory Kersten[1] & Raafat Saade[1]

[1] Concordia University, Montreal, Canada

Abstract

Negotiation is a powerful mechanism for facilitating effective economic exchanges. Electronic negotiations allow participants to negotiate online and use analytical support tools in making their decisions. Software agents offer the possibility of automating negotiation process using these tools. This paper aims at investigating the prospects of agent-to-human negotiations in B2C contexts using experiments with human subjects. Various types of agents have been configured and paired up with human counterparts for negotiating product sale. The paper discusses the results obtained both in terms of objective, as well as subjective measures.

http://interneg.org/

Acknowledgments: The publication of the InterNeg Research Papers has been supported by the Natural Sciences and Engineering Research Council, the Social Sciences and Humanities Research Council, and the J. Molson School of Business, Concordia University.

1. Introduction

Negotiation is an important mechanism for facilitating economic transactions. In the course of negotiations parties exchange offers in order to jointly explore the possibilities of finding acceptable solutions. Negotiations involving more than issue allow for more degrees of freedom in search for agreements that would be beneficial to the negotiators due to the asymmetry of their preference structures.

Online negotiations supported by electronic negotiation systems allow the parties exchange offers over the internet [1]. In addition to enabling anytime/anywhere mode of interactions, they may also incorporate analytical facilities for supporting negotiators in their preparation and conduct of negotiations. This support can range from such tools as those for capturing and modeling negotiator's preferences, to providing active advice and critique, and all the way to complete automation of the negotiation conduct.

Despite early optimistic expectations of the growth of negotiations as one of the primary mechanisms of conducting online transactions, in the reality only few commercial sites offer such capabilities to their customers. One such website that allows customers to make (a limited number of) offers is Priceline.com. One possible explanation to the scarcity of negotiating websites is that negotiations imply a relatively high cognitive load, especially if multiple issues are involved (e.g. price, warranty, product attributes, shipment, etc.). This load may translate into a prohibitive cost when day-to-day transactions are concerned.

Software agents may circumvent this problem by automating negotiation process while working with customers towards an acceptable deal. Moreover, they can also ensure consistency in reaching negotiation outcomes according to the set policies. They can be configured to behave in a competitive or collaborative fashion, depending on the context and the needs of a business. However, up to date little experimental work has been done in assessing the potential of human customer vs. software agent negotiations in terms of objective and subjective variables.

The purpose of this work is to investigate the prospects of human – software agent negotiations in experimental settings. To this end an electronic negotiation system incorporating software agents have been built. The system was used in experiments with human subjects to measure such outcomes as utility of agreements and number of agreements. Additionally, such subjective variables as satisfaction and perceived usefulness were also measured.

2. Related Work

Research on automated negotiations involving software agents has been extensive [2, 3]. While thorough coverage of the past work in the area is well beyond the scope of this paper, we will review the representative publications in the context of business exchanges. One could categorize these in accordance with the context of interactions (i.e. C2C, B2B, B2C), and the extent of automation.

One well-known early work in this direction was the construction of the Kasbah electronic marketplace [4, 5]. Targeting primarily the C2C domain the marketplace allowed human users to configure agents, which would then be sent to the marketplace to negotiate with each other. Three types of agents ranging from competitive to the conceding ones were provided.

Negotiations included a single issue, i.e. price.

In B2B applications software agents have been proposed for automating various aspects of supply chain management. For example, in [6] an agent- supply chain formation. The agents acting as brokers representing various entities within supply chain negotiated agreements with each other in building up the chain.

There has also been work targeting the B2C transactions. In [7] the authors proposed an agent-based architecture for automated negotiations between businesses and consumers. The buyer agents incorporated such components as searcher and negotiator, while seller agents featured negotiator module whose strategy was set by the sales department. In [8] the authors have proposed an intelligent sales agent with the capabilities for negotiation and persuasion. The agent employed reinforcement learning in the process. In their experiments with human subjects they found that the agent using persuasion capability has increased buyer's product valuation and willingness to pay.

It has been argued by many that complete automation of real-life negotiations, in particular in business contexts does not seem to be a viable solution (e.g. [9]). Automation in general is applicable only when tasks concerned are well-structured, which is rarely the case in many business situations. However, since efficient policies can be set for multiple daily interactions with the customers regarding the sales of products and services, it seems that a relatively high level of automation may be feasible.

While the work reviewed above concerns fully automated negotiations, there has been some research into sharing responsibilities between human negotiators and negotiation agents. In [10] a system has been proposed where agents actively supported human decision making in the negotiation process. An agent advised the human user on the acceptability of the received offer, helped with the preparation of the counter-offer, and critiqued offers composed by the users when it deemed necessary to intervene. In [11] an agent-based architecture was proposed for managing multiple negotiations. In this architecture a fleet of agents negotiated deals with customers. These negotiations were monitored by a coordinating agent, which, based on the analysis of situation instructed the negotiating agents to adjust their strategies and reservation levels within the limits of its authority. The overall process was monitored by a human user who could intervene to make changes if necessary.

The current work is aimed at investigating how software agents perform in agent-to-human dyads as compared to human-human dyads. Various types of agents following different strategies have been configured for the comparison of their performance.

Subjective measures have also been employed to measure the perceptions on the human side.

3. Negotiation Case and Configurations of Agents

The negotiation case developed for the experimental study concerned the sale of a desktop computer. There were five issues including the price, type of monitor, hard drive, service plan, and software loaded. Each option for each issue had a corresponding level of utility (attractiveness), these levels being different for the buyers vs. sellers. In order to calculate the total utility of the offer the issues were assigned different weights. These were then used in an additive utility function to estimate the level of attractiveness of an offer. Agents used this information in order to decide on the acceptability of the received offers and generate offers.

All agents acted on the seller side, and they were not aware of the buyers' preference structures. The weights were slightly different for sellers vs. buyers to facilitate integrative negotiations. Thus, agents would decide on the utility of the next offer first, according to their concession schedules, and then generate the corresponding offer.

We have chosen to use five different concession schedules, three of which were similar to those used in Kasbah experiments. These included: competitive, neutral, collaborative, competitive-then-collaborative, and tit-for-tat strategies. The competitive agents (CM) tend to make smaller concessions in terms of utility of generated offers in the beginning of the negotiation period. However, as they approach the end of the period, they would start making larger concessions in search of an agreement (figure 1).

Neutral strategy (NT) dictates that an agent concedes the constant amount of utility regardless of the time period, i.e. the concession schedule is linear (figure 2). Collaborative schedule (CL) implies making large concessions in the very beginning of the negotiation period in search of a quick agreement. This represents the case where an agent is anxious to sell the product. However, as the agent quickly drops the utility close to the reservation levels, it cannot make large concessions later in the process (figure 3).

Competitive-then-collaborative schedule (CC) models more complex behavior of the agents. In the beginning of the process an agent behaves competitively, however, in the middle of the negotiation period it changes its profile to a collaborative one. Thus, there is an inflexion point in an agent's schedule (figure 4). The reason for introducing this strategy is to imitate the situation when an agent's behavior adjusts due to the overall situation in the market (e.g. the product is not selling well). Moreover, the CC schedule allows introducing less predictable non-obvious behavior, which may be characteristic of human negotiators. (Little circles appearing on the screenshots are used to graphically define the shapes of the curves.)

Figure 1. Competitive schedule

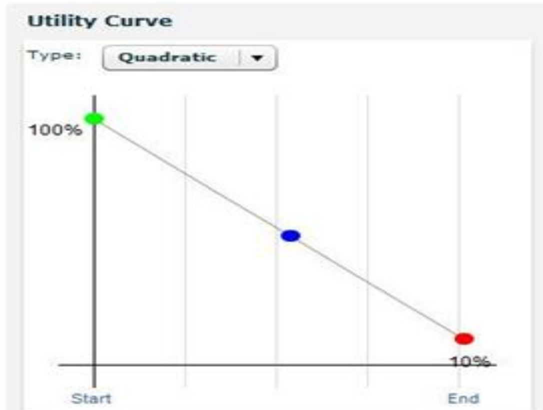

Figure 2. Neutral schedule

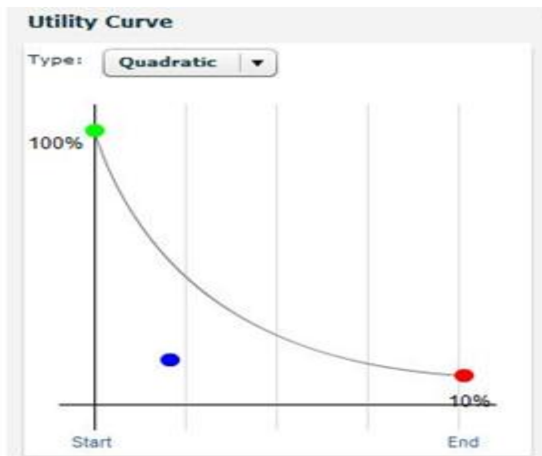

Figure 3. Collaborative schedule

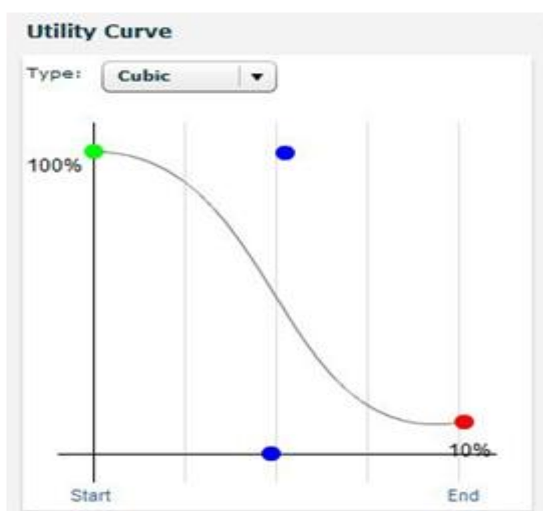

Figure 4. Competitive-then-collaborative schedule

The final strategy used is tit-for-tat. These agents do not rely on utility calculations. Rather, they watch the opponent moves and simply mirror them in composing counter-offers. In other words, when an opponent makes a new offer an agent determines the difference between this offer and the previous one made by the opponent, and applies the same difference to its own offer. If, say an opponent made a large change to a price, the agent would do the same.

The agent follows the following algorithm. In the beginning of the process it makes an offer that has highest utility to an agent. It then waits for the opponent to respond. If an opponent agrees, the process terminates. If an opponent makes a counter-offer the agent calculates its acceptable utility level according to the concession schedule employed. If the opponent's offer is equal or higher than the acceptable utility, the agent accepts the offer. Otherwise, the agent generates a new offer according to the acceptable utility level. It takes the opponent's offer as a starting point, and employing hillclimbing algorithm changes it to get close to the set utility level. This heuristic method is used instead of analytical one, since most of the issues are not continuous variables. It then sends this offer to the opponent.

4. Variables and Experimental Setup

In the current work we were interested in the objective outcomes of agent – human negotiations, as well as subjective variables capturing human perceptions of the process, outcomes and system. The objective variables included the utility of the agreements, and the proportion of agreements achieved. These relate to the economic benefits of agent-human negotiations. The subjective variables included satisfaction with the outcomes, satisfaction with the process, ease of use, and perceived usefulness of the system. These are important indicators from the information systems literature, especially relating to the acceptance and use of the system by human users.

The subjects in the study were university students enrolled in the introductory course on information technology. Thus, the negotiation case was well in line with the learning objectives of the course. The treatments included pairing up the subjects with various types of agents described in an earlier section. We also paired up humans with humans in a control group.

The experiment was conducted via the web, whereby subjects could perform their tasks from any location in an asynchronous mode during a two-day period. The subjects were invited to join the negotiations via email containing the link to the system. Negotiations began by sellers making the first offer. The agent sellers then checked for the status of negotiations at fixed intervals of time (every 3 hours). At those points of time, if they have not received new offers, they would wait until the next period of time elapsed. If an offer was received they would evaluate it and would either accept it, or would make a counter-offer. Human subjects were free to terminate the negotiation at any time without reaching an agreement with their counter-parts. After either reaching an agreement, or terminating the negotiations the human subjects were asked to complete a questionnaire measuring their perceptions of the outcome, process, and the system. One final question read: "I was negotiating with: 1) a human; 2) a computer: 3) not sure."

5. Results

For the analysis of the results we have selected only those negotiation instances, which featured at least four offers in total. The rationale for this decision was to include only those cases where the subjects took the task seriously. Thus, we ended up having 436 usable negotiation instances. Of these, 65% ended up in an agreement, while in 35% of cases the agreement was not reached.

Figure 5 shows the results of the question related to whether the participants guessed correctly if they were negotiating with humans or computers. The left side shows the results from human-agent dyads, and the right side shows human-human ones. The leftmost bar in each group indicates the number of responses that read “human”, the middle one relates to “computer” responses, and the last one shows “not sure” responses. As one can see, the majority of subjects in the agent-human dyads were not sure if they were interacting with the humans or agents (183 responses). This was followed by the group of subjects who had thought they were negotiating with other humans (114). The smallest group consisted of those who guessed correctly that they were interacting with agents (65). It is interesting to note that some subjects in the human-to-human dyads thought they were interacting with a computer (2 out of 30).

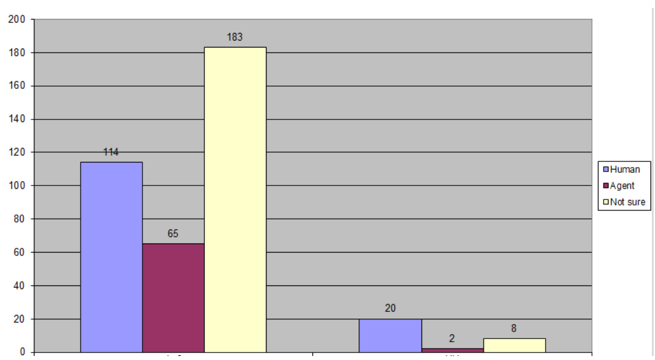

Figure 5. “I was negotiating with…” agent - human dyads vs. human – hyman dyads

The distribution of answers depended on the type of the agent strategy employed. For example, in competitive-then-collaborative category much larger proportion of subjects thought they were negotiating with a human counter-part as compared to those who had an impression they were dealing with a machine (25 vs. 8). This can be explained by the fact that CC concession schedule results in more complex behavior, less obvious behavior that could be more readily ascribed to humans, rather than machines. Similar, though less prominent results were obtained in competitive agent category (33 vs. 15). On the other hand, the collaborative category was the only one where the number of “human” vs. “machine” responses was equal (21 each). Perhaps, the subjects expected their human counterparts to be more competitive, rather than conceding.

Table 1 shows the proportions of agreements for different compositions of dyads. The largest proportion of agreements was reached in the collaborative agent category. This an intuitive result, since collaborative agents make large concessions early in the negotiations process, and thus they have a higher chance of making a deal with the human counterparts. It is interesting to see that human-to-human dyads have a second-lowest record in terms of proportion of agreements made. Thus, the majority of agent-involved dyads have reached more agreements than purely human dyads.

Competitive agents were able to reach an agreement in 53% of cases. Competitive-then-collaborative agents have made agreements in 75% of cases, falling between the CL and CM categories, but higher than neutral category. The lowest number of agreements was achieved in tit-for-tat category. This is the only agent strategy that does not employ utility function, and, thus it does not necessarily drop its utility level to the minimum towards the end of the period. Overall, agent-human pairs achieved agreements in 66% of cases vs. 50% exhibited by HH dyads.

Table 1. Proportions of agreements

Category	Agreements, %
All agent categories	66
Competitive	53
Neutral	70
Collaborative	82
Competitive-collaborative	75
Tit-for-tat	43
Human-human	50

Table 2 compares the utilities of reached agreements for sellers and buyers across different categories. In human-human dyads the sellers achieved much lower utility levels than buyers. This could be explained by the adopted reference frames. Since both sellers and buyers in this category were undergraduate student subjects, they tended to shift the price levels downwards to what they consider to be acceptable regions. Nonetheless, as it can be seen from the table, the human sellers had reached the lowest levels of utility. The highest average utility was achieved by tit-for-tat agents (72.4%). However, as already mentioned, they performed worst in terms of proportion of agreements reached. In terms of proportion of agreements the competitive agents have performed slightly better than human sellers. However, utility-wise these agents have considerably outperformed their human "colleagues" (63.2% vs. 35.9%). Collaborative agents did only slightly better than humans, reaching 36.5% utility. However, they had much higher proportion of agreements. Competitive-then-collaborative agents have reached the average utility level of 40.4%, and the neutral ones had a slightly higher value of 43.8%. Overall, agents did better than human negotiators (46.8% vs. 35.9 %).

Table 2. Utilities of agreements

Category	Seller utility, %	Buyer utility, %
All agent categories	46.8	65.6
Competitive	63.2	44.9
Neutral	43.8	69.7
Collaborative	36.5	79.0
Competitive-collaborative	40.4	71.9
Tit-for-tat	72.4	36
Human-human	35.9	73.0

In order to compare the subjects' perceptions a questionnaire was used with three items per construct measuring perceived usefulness, perceived ease of use and satisfaction with the outcome, and four items measuring satisfaction with the process. Factor analysis resulted in an acceptable pattern of loadings (Table 3). We have then used item averages for factors to compare across different categories. Results are shown in table 4.

Table 3. Factor analysis results

	Factor			
	PU	PEU	SP	SO
SO1	.146	.088	.402	**.532**
SO2	.207	.078	.258	**.789**
SO3	.235	.063	.345	**.721**
SP1	-.127	-.010	**-.404**	-.218
SP2	.112	.114	**.604**	.102
SP3	-.131	.002	**-.426**	-.175
SP4	.207	.067	**.712**	.263
PU1	**.684**	.194	.192	.172
PU2	**.816**	.130	.190	.177
PU3	**.809**	.164	.268	.204
PEU1	-.022	**-.635**	-.102	-.042
PEU2	.207	**.845**	.078	.062
PEU3	.191	**.804**	-.026	.059

There were no significant differences among the categories in terms of satisfaction with the process, perceived usefulness, and ease of use. There were some significant differences regarding satisfaction with the outcome, which is understandable. In particular, tit-for-tat and competitive strategies yielded lower satisfaction levels than some other strategies, such as collaborative. As human subjects had lower utility values of their agreements they also felt less satisfies with the outcomes. None of the categories yielded significantly different results as compared with human-human interactions.

Table 4. Comparison of item averages

	Factor			
	SO	SP	PU	PEU
All agents	4.35	4.06	3.48	3.34
CM	3.89	4.19	3.38	3.38
NT	4.60	4.12	3.48	3.31
CL	4.83	4.02	3.44	3.31
CC	4.73	4.04	3.79	3.35
TT	3.35	3.79	3.24	3.33
HH	4.03	3.97	3.36	3.34

6. Conclusions

The purpose of this study was to experimentally investigate the promises of agent-human negotiations in B2C context. To this end various types of agents were configured to conduct negotiations with human subjects. The question of whether humans were able to tell if they were negotiating with machine has important implications, since if they did they would be, in principle, able to predict the opponents moves. Findings indicate that, in most cases, the subjects were not able to make a correct guess. This is especially true when agents employed a complex concession pattern, i.e. compete-then-collaborate.

In regards with the objective outcomes the results show that human negotiators performed worst as compared to agents in terms of utility of agreements. They were also second worst in terms of number of agreements.

When it comes to selling products and services or retaining customers, human representatives of companies do sometimes negotiate with their customers. Some of these negotiations nowadays occur through electronic media, using such facilities as e-mail and chat. Thus, in this study we also looked at perceptive measures related to system acceptance and usage. We found no significant differences between agent-human vs. human-human dyads.

One possibility for future work could be conducting experimental studies where agent and human negotiators could add issues in the course of negotiations.

References

G. Kersten and S. J. Noronha, "WWW-based Negotiation Support: Design, Implementation, and Use," *Decision Support Systems,* vol. 25, pp. 135-154, 1999.

C. Beam and A. Segev, "Automated Negotiations: A Survey of the State of the Art," *Wirtschaftsinformatik,* vol. 39, pp. 263-268, 1997.

N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. J. Wooldridge, and C. Sierra, "Automated Negotiation: Prospects, Methods and Challenges," *Group Decision and Negotiation,* vol. 10, pp. 199-215, 2001.

A. Chavez, D. Dreilinger, R. Guttman, and P. Maes, "A Real-life Experiment in Creating an Agent Marketplace," in *Software Agents and Soft Computing*, H. S. Nwana and N. Azarmi, Eds.: Springer-Verlag, 1997, pp. 160-179.

P. Maes, R. H. Guttman, and A. G. Moukas, "Agents that Buy and Sell," *Communications of the ACM,* vol. 42, pp. 81-87, March 1999.

M. Wang, H. Wang, D. Vogel, K. Kumar, and D. K. W. Chiu, "Agent-based negotiation and decision making for dynamic supply chain formation," *Engineering Applications of Artificial Intelligence,* vol. 22, pp. 1046-1055, 2009.

C.-C. Huang, W.-Y. Liang, Y.-H. Lai, and Y.-C. Lin, "The agent-based negotiation process for B2C e-commerce," *Expert Systems with Applications,* vol. 37, pp. 348-359, 2010.

S.-l. Huang and F.-r. Lin, "The design and evaluation of an intelligent sales agent for online persuasion and negotiation," *Electronic Commerce Research and Applications,* vol. 6, pp. 285-296, 2007.

R. Lin and S. Kraus, "Can automated agents proficiently negotiate with humans?," *Communications of the ACM,* vol. 53, pp. 78-88, 2010.

E. Chen, R. Vahidov, and G. E. Kersten, "Agent-supported negotiations in the e-marketplace," *International Journal of Electronic Business* vol. 3, pp. 28-49, 2005.

R. Vahidov, "Situated Decision Support Approach for Managing Multiple Negotiations," in *Negotiation, Auctions, and Market Engineering*. vol. 2, H. Gimpel, N. R. Jennings, G. E. Kersten, A. Ockenfels, and C. Weinhardt, Eds.: Springer Berlin Heidelberg, 2008, pp. 179-189.