

Bridging the Gap: Face-to-Face Negotiations with Automated Mediator

Raz Lin and Yehoshua Gev and Sarit Kraus

Abstract—Mediation is an important paradigm for dispute resolution. It can lead to “win-win” situations and benefit all parties. Yet, when people negotiate they demonstrate bounded rationality in their actions and diversity in their behaviors. This increases the difficulty to design automated mediators. To be successful, the mediator must take this into account and propose solutions deemed relevant, otherwise it will lose the focus and trust of the negotiators. To tackle these difficulties we present *AniMed** – an automated animated mediator, incorporated with a novel proposal generation strategy, aimed to increase the social benefit of the negotiating parties. To validate the benefits of using *AniMed** in negotiations, experiments were conducted with 130 people negotiating with each other. The results demonstrate the significant increase both in the social welfare and the individual utilities of both parties, compared to negotiations in which another version of reduced functionalities of *AniMed** was used, as well as negotiations involving another state-of-the-art automated mediator or no mediator.

Index Terms—bilateral negotiations, automated mediation, incomplete information.



1 INTRODUCTION

NEGOTIATIONS are procedures for resolving opposing preferences between two or more parties by means of discussion. The goal is to reach an agreement, i.e. a mutually accepted solution, without resorting to a struggle. Mediation, which is the involvement of a third party in the negotiation process, dates back to Ancient Greece [1] and has evolved to assist the negotiating parties in reaching beneficial solutions and increasing their social welfare. In many situations, the mediator does not have the authority to impose a solution on the parties or the power to compel them to uphold the agreement reached (unlike arbitration). Usually the mediator is neutral (unbiased) and objective.

Automated mediators, intelligent agents that take the role of an active mediator in the negotiation process, can play an important role in bridging the gap between people as they negotiate. They offer a discrete, impartial facilitator that might be

more trusted than a human one. The computational resources of automated mediators may also be more useful when incomplete information exists and there is uncertainty regarding the preferences of the parties, as the difficulty for a human mediator only increases. Yet the use of automated mediators is far from widespread, perhaps due to the difficulties in bridging between people, or due to their computerized nature.

In this paper we introduce *AniMed** – a domain-independent automated vivid and animated mediator designed to improve the social welfare of people in bilateral negotiations. *AniMed**, an English speaking avatar, interacts with people who negotiate face-to-face by means of a video-conference. *AniMed**'s novel design allows it to propose solutions that are in the context of the current negotiation state. This strategy differentiates it from other automated mediators found in the literature. *AniMed** also has the capability of proposing partial solutions, and by doing so it provides the negotiators with the option to incrementally strive for a beneficial solution. Moreover, the strategy incorporated in *AniMed** does not rely on the structure of the utility function of both negotiators, but rather constructs a preference relation between the possible solutions. Thus, *AniMed** has a generic strategy mechanism, allowing it to be matched and mediate proficiently

-
- S. Kraus is with the Department of Computer Science, Bar-Ilan University, Ramat-Gan 52900, Israel, and also with the Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 USA.
e-mail: sarit@cs.biu.ac.il
 - R. Lin and Y. Gev are with the Department of Computer Science, Bar-Ilan University, Ramat-Gan 52900, Israel.
e-mail: linraz@cs.biu.ac.il

with many possible types of negotiators without any restriction to any specific domain. In addition, *AniMed** was supplied with a facilitation mechanism that provided the negotiators a “negotiation calculator” which enabled them to calculate the utilities of each solution at any given time. Lastly, as *AniMed** was built on top of GENIUS, a generic negotiation framework [2], it will be available for the public and can be modified and used in numerous domains. To validate the benefits of *AniMed* we also designed two reduced functionalites versions of it – *AniMed* which did not include a facilitation mechanism and *Animed⁻* which did not include an animated avatar (the mediator’s proposals were shown in plain text).

*AniMed**, along with two reduced functionalities version of it, were evaluated in experiments with 130 people who negotiated face-to-face on a neighbor dispute domain [3] by means of video-conferences. The negotiation involved uncertainty with respect to the utility values of opposing parties. This uncertainty was also shared by the mediator, that had information solely on the preference relation between the issues and values under negotiation. *AniMed** significantly increased the individual utility score and the social welfare, measured by the sum of utilities, of both negotiators, compared to experiments in which *AniMed⁻*, another state-of-the-art mediator [4] or no mediator were involved. The results also indicate that while people are content with the solutions they achieve without any mediator involved, better solutions can be achieved when *AniMed** is present.

2 RELATED WORK

Mediator agents can be used in many settings. For example, they can be used to solve distributed constraint satisfaction problems (e.g., [5]) or as self-interested shopping agents [6]. However, not many have focused on automated mediators in human-human negotiations. Some are discussed in the context of online dispute resolutions, which are mostly alternative services to litigation. For example, eBay’s resolution center (<http://resolutioncenter.ebay.com/>) tries to facilitate the resolution of conflicts between buyers and sellers.

e-Alliance [7] is an automated mediator that offers support for multi-issue, multi-participant (different partners can be involved) and multiple-cycle (cycles of proposals and counter-proposals over the

same set of attributes) negotiations. These characteristics make the facility flexible enough for use in different domains. While e-Alliance was developed for agent-agent interactions, we are interested in the problem of human-human interactions.

AutoMed [4] is an automated mediator that most resembles our proposed mediator. *AutoMed* monitors the exchange of offers and actively suggests possible solutions. It uses a qualitative model for the negotiator’s preferences, without past knowledge. The suggestions it makes are Pareto-optimal according to its estimations. *AutoMed* was evaluated with human negotiators who negotiated using a computer system by exchanging offers selected from drop-down lists. *AutoMed* participates as a third-party that sends suggestions via the system. However, *AutoMed* has its limitations. Mainly it does not suggest incremental (partial) solutions nor does it provide explanations for its suggestions. Moreover, *AutoMed* constrains the negotiators to negotiate through the system, while a more natural approach would be to negotiate face-to-face. These drawbacks are nonexistent in *AniMed**, allowing it to generate more satisfactory solutions that are deemed more relevant by the negotiating parties.

The incorporation of an animated avatar is also supported by Yee *et al.* [8] who have investigated the impact of avatars on user experiences. They showed that using visualization leads to better interactions with people. Our experiments have indeed shown the different outcomes when an animated mediator was used compared to others.

3 NEGOTIATION PROBLEM DESCRIPTION

We consider the problem of a proficient automated mediator to be a key to improving the performance of two human negotiators who strive to reach a solution on conflicting issues. The mediator is situated in finite horizon bilateral negotiations with incomplete information between two people. The negotiation consists of a finite set of multi-attribute issues and time constraints.

Following Chalamish and Kraus [4], let $\mathcal{I} = \{I_1, I_2, \dots, I_m\}$ be a finite set of *issues* over which the negotiators disagree. Each I_j can be assigned a *value* from a finite domain of I_j , $dom(I_j) = \{v_1^j, v_2^j, \dots, v_b^j\}$ of possible values, where b is the maximum number of values in all issue domains and

$b_j \leq b$. Let $\vec{a} \in \text{dom}(I_1) \times \text{dom}(I_2) \times \dots \times \text{dom}(I_m)$ denote a possible *solution* reached in the negotiation and \mathcal{A} denote the list of all possible solutions. We will denote the total possible solutions in \mathcal{A} by k , that is $|\mathcal{A}| = k$. The status quo is also a solution and thus it is a member of \mathcal{A} . A *partial solution* is an assignment of values to issues, where not all issues are assigned a value. The unassigned values are annotated with the empty value ε .

The negotiation can end when (a) the negotiators reach a full solution, (b) one of the negotiators opts out, thus forcing the termination of the negotiation with an opt-out outcome, or (c) a predefined deadline is reached, whereby a status quo outcome is implemented. The negotiators can propose any number of possible solutions, while the other negotiator can either accept or reject. Each negotiator can either propose a solution which consists of all the issues in the negotiation or a partial solution.

Each negotiator is given the exact valuation of the values of $\text{dom}(I_j)$, and thus can evaluate the utility of a solution as the sum of these valuations. However, this is private information and is not shared between the negotiators. On the other hand, each negotiator shares with the mediator some information. For each issue in the set of issues, $\mathcal{I} = \{I_1, I_2, \dots, I_m\}$ and a set of values $\text{dom}(I_j) = \{v_1^j, v_2^j, \dots, v_{b_j}^j\}$ for each $I_j \in \mathcal{I}$, an importance level, $\lambda : \mathcal{I} \rightarrow \{1, \dots, m\}$, is assigned. Also a full order, \prec , over the values of $\text{dom}(I_j)$ exists. Thus, the importance level per issue, as well as the full order of values per issue, are shared with the mediator.

However, based on this information, the mediator cannot still decide which solution is preferred by a negotiator. Suppose $v_1 \prec v_2 \prec \dots \prec v_{b_j}$ are the ordered values of issue I_j . The rank value of v_i , $r(v_i) = \frac{i-1}{b_j-1}$, $1 \leq i \leq b_j$, assigns each value a rank between 0 and 1. Using these ranks, the mediator determines the scoring value of a solution, $\vec{a} = (v^1, \dots, v^m)$ as $s(\vec{a}) = \sum_{i=1}^m \lambda(i)r(v^i)$.

Now *AniMed** can *estimate* whether one solution \vec{a}_1 is at least as good as \vec{a}_2 according to negotiator n , denoted $\vec{a}_1 \succeq_n \vec{a}_2$ if $s^n(\vec{a}_1) \geq s^n(\vec{a}_2)$. If $s^n(\vec{a}_1) > s^n(\vec{a}_2)$ *AniMed** will assume that \vec{a}_1 is better than \vec{a}_2 according to negotiator n , denoted $\vec{a}_1 \succ_n \vec{a}_2$. We will also denote by $\text{order}_n(\cdot)$ the ranking of each solution compared to all other solutions, based on this estimation. Thus, $\text{order}_n(\vec{a})$ is a number between 1 and k , where k implies the estimated

best solution for negotiator n .

4 NEIGHBOR DISPUTE DOMAIN

For the implementation and experiments we chose a neighbor dispute domain [3]. In this domain, a negotiation takes place between two tenants, Alex and Tyler, due to a neighbors' dispute. Both negotiators need to negotiate in order to resolve the dispute, or otherwise be forced to undergo a lengthy and costly dispute resolution process. The issues under negotiation are with respect to: (a) throwing out the trash, (b) using a shared basketball court, (c) making noise, (d) using the patio, and (e) the parking of friends in the apartment's parking lot.

The scenario was symmetric for both negotiators, in the sense that the negotiators could compromise and make tradeoffs between the issues and the gains and losses were equivalent. On the parking lot issue both negotiators received the same utility. On two other issues – the basketball court and the patio – the more one gained, the less the other gained. These two issues had the same scale in the utility scores. For the last two issues – trash and noise – the negotiators could use tradeoffs between the values of both issues to gain the same utilities (detailed score functions for the domain can be found at <http://u.cs.biu.ac.il/~linraz/Papers/neighbors-utilities.pdf>). The utility values ranged from 0 to 1,000 for both negotiators, where the Pareto-optimal solution generated a utility of 720 for both.

The negotiation was limited to 28 minutes. If the negotiators did not reach a solution by the end of the allocated time, the negotiation ended and both tenants would be required to undergo a costly dispute resolution session. This outcome was modeled for both negotiators as the status quo outcome. Each negotiator could also opt-out of the negotiation if it felt that the prospects of reaching a solution with the opponent were slim and that it was no longer possible to negotiate.

5 MEDIATOR DESIGN

The design of *AniMed** is built on top of the GENIUS infrastructure, which is an integrated environment for supporting the design of generic automated negotiators [2]. This environment is rich and supports bilateral multi-issue and multi-attribute

negotiations, both with human counterparts and automated agents. An example snapshot of a negotiation interface is given in Figure 1. The focus of the research was to design an automated formulating mediator (as opposed to a manipulative one). That is, the agent tries to propose solutions and help the negotiators reach a mutually agreed outcome. *AniMed** is not topic embedded, allowing it to be used in many scenarios, and it is aimed to be proficient in bilateral negotiations involving people.

*AniMed** was implemented using two main considerations. First, a proficient strategy was used to enable it to generate solutions deemed relevant by the negotiating parties. To achieve this, *AniMed** utilizes recent solutions proposed by the negotiators when generating its own solutions, thus centering its solutions on the current context of the negotiation. The second consideration involves its user interface design. *AniMed** was implemented with a rich animated interface to make it appealing and user-friendly for people (see Figure 2). The animated design of *AniMed** and the structure of the experiments were also motivated by findings of Nass and Moon [9], with respect to human-human versus human-computer interaction. Thus, when the negotiation begins, the automated mediator introduces itself in a way that will allow both sides to trust it.

The motivation behind the strategy design of *AniMed** was to generate solutions that would maximize the social welfare of both negotiators. However, this is a difficult task due to conflicting interests between the negotiators, as well as due to the fact that *AniMed** is unaware of the exact utilities of the negotiators. To overcome this, *AniMed** starts proposing only after both negotiators have proposed or accepted a solution in the past. It uses this information to try to find a set of solutions that can still increase both negotiators' utilities. One of its strengths is its ability to provide a solution only on a subset of the issues under negotiation, allowing the negotiators to incrementally improve the final solution. In addition, to prevent the negotiators from labeling its solutions as irrelevant, *AniMed** does not propose any solutions that are identical to the last solutions proposed by the negotiators.

The strategy used by *AniMed** consists of five steps, which we describe hereafter using an example to better illustrate it. We will label the negotiators as *A* and *B*. Assume there exist 7 possible solutions.

Also assume negotiator *A* and *B* just proposed solutions labeled 3 and 5, respectively. Table 1 lists the information about the domain and the steps taken by *AniMed** to decide on a solution to propose.

The first step in the algorithm is taken before the negotiation starts. *AniMed** determines the scoring value of all solutions and orders them, as described in Section 3. Then, during the game play, *AniMed** chooses its suggestions based on these orderings and on the last solutions proposed by the negotiating parties, which are highlighted in Table 1.

The next two steps are motivated by the strategy of *AutoMed*. The second step in *AniMed**'s strategy is to discard all solutions that, for each party, have a lower ordering than her last proposal. Thus, *AniMed** removes solutions with lower ordering for negotiator *A* (#1, #2) and *B* (#7) while it keeps the four other solutions. Then, in its third step, *AniMed** searches for any non-Pareto optimal solutions and removes them as well. In our example, one of the solutions was non-Pareto optimal (#5).

In the fourth step *AniMed** orders all remaining solutions based on the following criteria. First, it orders them based on the solutions' joint ordering (that is, the sum of both orderings, marked as "joint order" in Table 1). If the solutions have the same joint ordering, it compares them to previous solutions proposed by each negotiator. However, in order to allow *AniMed** to propose solutions which are in context with the solutions previously suggested by the negotiating parties, solutions that are more similar to previously suggested solutions and measured by the number of similar values between the solutions, are ordered higher. If there are still solutions with the same rank, it orders them based on the absolute difference in their ordering (marked as "diff order" in Table 1). Algorithm 1 presents the pseudo-code of the algorithm for generating a full proposed solution.

*AniMed** now has a full solution that it can propose. However, from preliminary experiments, we observed that the dynamics of the negotiation mainly involve partial solutions. Thus, the fifth step in *AniMed**'s strategy is to generate partial solutions that could still benefit the negotiators. *AniMed** incorporates one of the two mechanisms below for generating partial solutions:

- First, *Animed** generates a set of partial solutions based on joint-value issues (that is, issues with values that are estimated as gener-

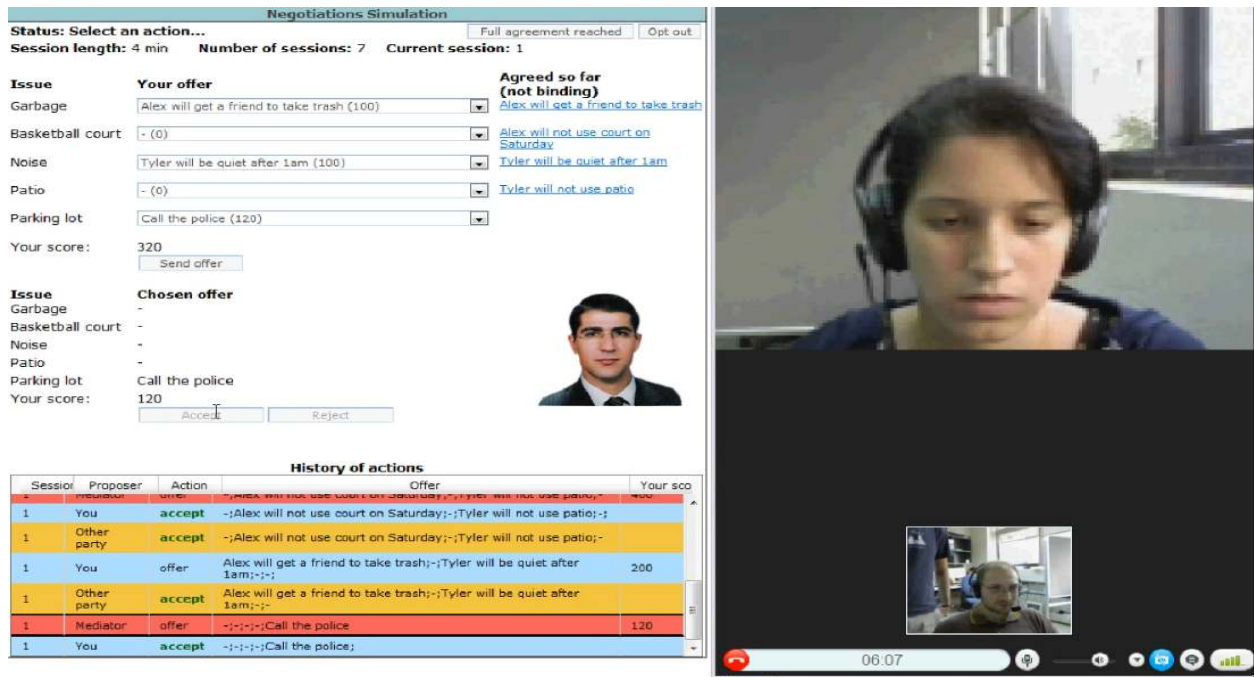


Fig. 1. An example of a negotiation snapshot using GENIUS.

Solution Idx	$order_A(\vec{a})$	$order_B(\vec{a})$	joint order	diff order
1	1	7	8	6
2	2	6	8	4
3	3	5	8	2
4	4	4	8	0
5	5	2	7	3
6	6	3	9	3
7	7	1	8	6

TABLE 1

A sample domain for choosing a solution to propose. For each negotiator, A and B , the possible solutions are ordered by her own preferences. The last solutions proposed by the negotiators are emphasized.

ating higher utilities for both parties, based on their importance). For example, in the negotiation demonstrated in Figure 1, several offers were exchanged by the parties until the parties agreed on four out of the five issues under negotiation (the trash issue, court, patio and quiet). The mediator then proposes a partial solution "Call the police" for the parking issue which is Pareto optimal for both.

- Second, *AniMed** applies a trade-off scheme between the issues to choose the partial solution to be proposed. This is done by calculating

Algorithm 1 Generating A Full Proposed Solution

```

1: for all possible solutions  $\vec{a}$  do
2:   Insert  $\vec{a}$  to OrderedList $_A$ , OrderedList $_B$ 
3:   Using  $order_A(\vec{a})$ ,  $order_B(\vec{a})$ 
4: end for
5: if Both sides interacted then
6:   for all possible solutions  $\vec{a}$  do
7:     if  $order_A(\vec{a}) < order_A(\text{last solution } A \text{ proposed})$  then
8:       remove  $\vec{a}$  from OrderedList $_A$ 
9:     end if
10:    if  $order_B(\vec{a}) < order_B(\text{last solution } B \text{ proposed})$  then
11:      remove  $\vec{a}$  from OrderedList $_B$ 
12:    end if
13:  end for
14:  SolutionsList =
  Intersect(OrderedList $_A$ , OrderedList $_B$ )
15:  SolutionsList = ParetoSolutions(SolutionsList)
16:  Define  $jointOrder(\vec{a}) = order_A(\vec{a}) + order_B(\vec{a})$ 
17:  Define  $diffOrder(\vec{a}) = abs(order_A(\vec{a}) - order_B(\vec{a}))$ 
18:  Sort SolutionsList
  Using  $jointOrder(\vec{a})$ 
  Then SimilarityToRecentSolutions( $\vec{a}$ )
  Then  $diffOrder(\vec{a})$ 
19: end if

```

the distances between the importance of given issues, denoted by "diff importance" in Table 2. For example, the result for the pair $\langle \text{Basketball court, Patio} \rangle$ is $3 - (-4) = 7$. The higher the sum, the better candidate it is for selection in the partial solution in order to allow trade-offs. Figure 2 demonstrates the proposition of a trade-off based on these two issues. This example is based on a negotiation in which the parties exchanged several offers and did not come to an agreement. The mediator's offer

Issue	$\lambda_A(issue)$	$\lambda_B(issue)$	diff importance
Garbage	2	2	0
Basketball court	1	4	3
Noise	3	3	0
Patio	5	1	-4
Parking lot	4	5	1

TABLE 2

An example of deciding the trade-off between issues.

also included a message stating that it believes that they can achieve higher scores if they are willing to compromise on these issues.

Another consideration implemented in *AniMed** is a simple argumentation mechanism to try to convince the parties why the proposed solution suggested by *AniMed** should be considered. When *AniMed** proposes a solution it attaches a predefined text message stating that if the negotiators make the suggested trade-off based on the issues previously discussed, they can achieve higher scores (the text is slightly different if the suggested solution includes issues that were previously agreed upon by the parties or simply discussed).

An additional approach incorporated in *AniMed** relates to its presence during negotiations. In order to compel people to listen to the mediator’s proposals, whenever it proposes a solution, it takes over the entire screen so people cannot conceal or ignore it. Moreover, the mediator was implemented as an English speaking avatar (translation to native language was given in the text) (see Figure 2), using a commercial text-to-speech engine in order to convey a more “humanized” appearance and a less distant and computerized one.

6 EXPERIMENTS

The two negotiators were given the task of negotiating a beneficial solution. Five different experiments were conducted using the same domain, in order to compare the proficiency of *AniMed**. Each subject played in only one of the experiments in order not to bias the results. One experiment involved matching people without any mediator. In another experiment we matched two people with a simple automated mediator, *AutoMed*, suggested by Chalamish and Kraus [4]. They demonstrated that this mediator enables the negotiators to achieve more satisfactory




Fig. 2. *AniMed** avatar example.

solutions in environments where only messages are exchanged.

Finally, we matched people in a setting which included our proposed mediator, *AniMed**, as well as its two reduced functionalities versions – *AniMed* and *AniMed⁻*.

6.1 Experimental Methodology

The human negotiators accessed the negotiation interface via a web address. The negotiation itself was conducted as follows: Using a video conference the two players negotiated face-to-face about the different relevant issues. Since the focus of the research was on the strategy of the automated mediator, natural language processing (NLP) was beyond the scope of our research, and thus we required the negotiators to submit their proposals also using the negotiation system (this might also resemble real life negotiations in which people talk and only propose official solutions once in a while). This allowed the information to be processed by the automated mediator. Nonetheless, the negotiation itself was not constrained and was employed via a face-to-face video conference. The acceptance or decline of each solution was also done via the user interface. The mediator sent proposed solutions to the parties via the animated avatar and the negotiation system.

We tested our agent against human subjects, all of whom are computer science undergraduates and graduate students. 130 human subjects participated in the experiments (65 pairs). A total of five sub-

		Alex	Tyler	SW	Outcome Satisfaction	Mediator's Helpfulness
<i>AniMed*</i>	Avg	735	669	1404	3.31	2
	Stdev	52	45	59		
<i>AniMed</i>	Avg	723	665	1388	3.29	1.08
	Stdev	92	69	75		
<i>AniMed⁻</i>	Avg	674	609	1283	3.23	0.96
	Stdev	104	111	164		
<i>AutoMed</i>	Avg	651	590	1241	3.11	0.35
	Stdev	80	103	145		
<i>No Mediator</i>	Avg	645	595	1240	3.19	N/A
	Stdev	130	121	150		

TABLE 3

Average utility scores, social welfare (SW), standard deviations and satisfaction levels (with 0 being the lowest and 4 the highest) in the different experiments.

experiments were conducted, 13 pairs in each sub-experiment. The subjects did not know any details regarding the automated mediator with which they were matched. The outcome of each negotiation was either they reached a full solution, they opted out or the deadline was reached.

Prior to the experiments, the subjects were given oral instructions and were shown an instructional video regarding the experiment and the domain. Handouts with the negotiation protocol, specification and their scoring function were handed to the subjects. The subjects were instructed to play based on their score functions and to achieve the best possible solution for them.

6.2 Experiment Results

To verify the proficiency of *AniMed** we compared the final utility results in all experiments, as well as the number of proposals exchanged between the negotiators in each experiment. Lastly, we administered questionnaires inquiring about the satisfaction of the negotiators from the outcome and their view on the helpfulness of the automated mediator.

Table 3 summarizes the results of the individual utilities and the social welfare, measured by the sum of utilities of the negotiating parties (in our domain they are denoted as Alex and Tyler). First we examined the final utility values of all the negotiations for each role and the social welfare, measured by the sums of the final utility values. We used *ANOVA* in order to test whether there is a significant difference when using the different types of automated mediator. Using *Scheffe's test*

as post hoc testing of the ANOVA, we obtained that using *AniMed** or *AniMed* allows achieving significantly higher social welfare (1388 and 1404 for *AniMed** and *AniMed*, respectively) than using *AniMed⁻* (1283 with $p < 0.009$ and $p < 0.02$ as compared to *AniMed** and *AniMed*, respectively), *AutoMed* (1241 with $p < 0.0075$ and $p < 0.004$) or no mediator at all (1240 with $p < 0.0065$ and $p < 0.004$). The utility of both sides was also statistically significantly different when *AniMed** was involved compared to when *AniMed⁻*, *AutoMed* or no mediator was involved.

While *AniMed** achieved better results than *AniMed* the difference was not statistically significant, yet *AniMed** was able to achieve significantly better results than *AniMed⁻* in both roles, while *AniMed* did not.

Finally, we gathered the satisfaction levels of the negotiators from the final outcome they reached and their perception of how helpful the mediator was in reaching this outcome (see Table 3). The satisfaction levels ranged from 0 (lowest) to 4 (highest). The results significantly demonstrate that the negotiators perceived *AniMed** as more helpful than *AutoMed* (using Kruskal-Wallis test with $p < 0.0001$). Surprisingly, the negotiators were content with the final outcome in every experiment, and though the satisfaction level was slightly higher when *AniMed** was the mediator the difference was not statistically significant. This is in contrast to the fact that the negotiators achieved *higher* utilities, both individually and combined, when *AniMed** was involved, compared to the other experiments. These results support our belief in the need and benefits of using mediators in negotiation settings when people are involved. It is also interesting to note that when *AniMed⁻* was used, it received the lowest satisfaction and helpfulness ratings, compared to *AniMed** and *AniMed*.

7 CONCLUSIONS

This paper presents *AniMed**, a novel automated mediator capable of proficiently interacting with people. The success in proficiently interacting with people has great implications on the outcome of the negotiations and allows the negotiating parties to maximize their revenues.

Experiments with 130 people demonstrated the benefits of *AniMed** compared to other automated

mediators and to settings without any mediator. In essence, the main contribution of *AniMed** compared to the other mediators relies on its vivid animated English speaking avatar representation, as well as its strategy mechanism, mainly being domain-independent, proposing solutions that are in the context of the current negotiation state and its ability to propose partial solutions. The fact that *AniMed** can be employed in any setting with the requirement of knowing only the structure of the preference relation between the issues, reflects on its generality and its prospects of becoming widespread and useful in numerous settings.

Future research will involve validating the results in additional scenarios, including ones with non-linear utility functions and ones with a larger number of issues. We will also investigate other negotiation protocols, in which the negotiation issues are not necessarily listed a-priori but rather elicited during the negotiation itself. We will also extend *AniMed** to present the negotiators with threats and the ability to enforce solutions and penalties. Argumentation is an important issue as well, and future work will also extend the argumentation mechanism implemented by *AniMed**.

ACKNOWLEDGMENTS

This research is based upon work supported in part under NSF grant 0705587 and by the U. S. Army Research Laboratory and the U. S. Army Research Office under grant number W911NF-08-1-0144. We thank Michele Gelfand and Elizabeth Salmon for suggesting the domain and the fruitful discussions and comments.

REFERENCES

- [1] D. Roebuck, *Ancient Greek Arbitration*. Holo Books, 2001.
- [2] R. Lin, S. Kraus, D. Tykhonov, K. Hindriks, and C. M. Jonker, "Supporting the design of general automated negotiators," in *Proceeding of the Second International Workshop on Agent-based Complex Automated Negotiations (ACAN)*, 2009.
- [3] E. Salmon and M. J. Gelfand, "The oakland apartment dispute," unpublished case, 2011, university of Maryland, College Park.
- [4] M. Chalamish and S. Kraus, "AutoMed - an automated mediator for multi-issue bilateral negotiations," *Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 2011.
- [5] R. Mailler and V. Lesser, "Using cooperative mediation to solve distributed constraint satisfaction problems," in *Proceedings of the Third International Joint Conference on Multiagent Systems (AAMAS'04)*, 2004, pp. 446–453.
- [6] D. Sarne, "Competitive comparison-shopping mediated markets," in *the 2009 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT'09)*, 2009, pp. 387–394.

- [7] M. Bratu, J. M. Andreoli, O. Boissier, and S. Castellani, "A software infrastructure for negotiation within inter-organisational alliances," in *Proceedings of the Fourth International Workshop on Agent-Mediated Electronic Commerce (AMEC-IV)*, Berlin, Germany, 2002, pp. 161–179.
- [8] N. Yee, J. N. Bailenson, and K. Rickertsen, "A meta-analysis of the impact of the inclusion and realism of human-like faces on user experiences in interfaces," in *Proceedings of Computer/Human Interaction (CHI'07)*, 2007.
- [9] C. Nass and Y. Moon, "Machines and mindlessness: Social responses to computers," *Journal of Social Issues*, vol. 56, no. 1, pp. 81–103, 2000.



Raz Lin received the B.Sc. degree (summa cum laude) in mathematics and computer science, the M.Sc. degree (magna cum laude) in computer science, and the Ph.D. degree from the Bar-Ilan University, Ramat-Gan, Israel, in 2001, 2002, and 2008, respectively. He is currently a Postdoctoral Fellow with the Department of Computer Science, Bar-Ilan University, where he investigates issues of automated negotiations, personalization, training, and learning, and where he also received a four-year Presidents scholarship for outstanding students.



Yehoshua Gev is currently working toward an M.Sc. degree in computer science at Bar-Ilan University, Israel, where he has received a B.Sc. degree in computer engineering in 2009. His research interests include security of IP networks and defenses against denial-of-service attacks.



Sarit Kraus (Ph.D. Computer Science, Hebrew University, 1989) is a Professor of Computer Science at Bar-Ilan University and an Adjunct Professor at the Institute for Advanced Computer Studies, University of Maryland. She has been working extensively in the following areas: the development of intelligent systems, negotiation and cooperation in mixed open environments (including people), personalization, learning and clustering, optimization of complex systems and security of physical systems. In 1995 Kraus was awarded the IJCAI Computers and Thought Award (the premier award for a young AI scientist). In 2001 she was awarded the IBM Faculty Partnership Award and in 2002 she was elected as AAAI fellow. In 2007 she was awarded the ACM SIGART Agents Research award and her paper with Prof. Barbara Grosz was a winner of the IFAAMAS influential paper award (joint winner). In 2008 she was elected as ECCAI fellow and in 2010 she was awarded the EMET prize. She has published over 270 papers in leading journals and major conferences and is an author of the book *Strategic Negotiation in Multiagent Environments* (2001) and a co-author of a book on *Heterogeneous Active Agents* (2000); both published in MIT Press. Kraus is an associate editor of the *Annals of Mathematics and Artificial Intelligence Journal*, and on the editorial board of the *Journal of Autonomous Agents and Multi-Agent Systems*, the *Journal of Applied Logic*, the *Journal of Philosophical Logic* (JPL) and the *ACM Transactions on Intelligent Systems and Technology*.